Enhancing Peer-to-Peer Traffic Locality through Selective Tracker Blocking

Haiyang Wang, Feng Wang, Jiangchuan Liu

Enhancing Peer-to-Peer Traffic Locality Through Selective Tracker Blocking

Haiyang Wang, Feng Wang, and Jiangchuan Liu

{hwa17, fwa1, jcliu}@cs.sfu.ca
School of Computing Science, Simon Fraser University,
British Columbia, Canada

Abstract. Peer-to-peer (P2P) applications, most notably BitTorrent (BT), are generating unprecedented traffic pressure to the Internet Service Providers (ISPs). To mitigate the costly inter-ISP traffic, P2P locality, which explores and promotes local peer connections, has been widely suggested. Unfortunately, existing proposals generally require that an ISP control the neighbor selection of most peers, which is often not practical in a real-world deployment given there are *noncooperative trackers*.

In this paper, for the first time, we examine the characteristics and the impacts of these noncooperative trackers through real-world measurements. We find that tracker blocking has the potential to address this noncooperation problem, and help the ISPs to control more peers for traffic locality. Yet, how to guarantee torrents' availability at the same time remains a significant challenge for the ISPs. To this end, we model the tracker blocking problem coherently with torrent's availability, and address it through a novel selective tracker blocking algorithm, which iteratively improves traffic locality with a given availability threshold. Our trace-driven evaluation shows that our solution successfully reduces the cross-ISP traffic in the presence of noncooperative trackers and yet with minimal impact to torrents' availability.

Key words:
BitTorrent, Locality, Peer control, Tracker blocking

1 Introduction

Peer-to-Peer (P2P) networking has emerged as a successful architecture for content sharing over the Internet. BitTorrent (BT), the most popular P2P application, has attracted attentions from network operators and researchers for its wide deployment. BitTorrent, however, also generates a huge amount of cross-ISP traffic. Since the ISPs typically pay their peering or higher-level ISPs for global connectivity, the traffic between different ISPs is costly and presents significant network engineering challenges.

In order to alleviate this cross-ISP traffic, P2P locality has been widely suggested. Different with caching or blocking the P2P traffic, this method explores the access of existing localities to reduce the long-haul traffic. It is well known that modifying the trackers[1] is the one of the most efficient ways to deploy P2P locality. This method proposes to manipulate the neighbor selection of peers via a biased tracker deployment.

In particular, the biased trackers will select the majority, but not all, of the neighbors for the BT peers within the same ISP.

In this paper, we explore the problem that many trackers are not owned by the ISPs and thus can hardly be modified for traffic locality. We take a first step towards the characteristics and the impacts of these *noncooperative trackers* (the trackers that cannot be modified for locality) via extensive measurements. We show that the existence of noncooperative trackers will greatly reduce the efficiency of traffic locality. In particular, we show that an ISP will lose the control of 53% ($67,132$ out of $126,133$) peers due to the existence of some Pirate Bay trackers[1]

The well known tracker blocking approach has the potential to address the problem. Yet, this approach will also reduce torrents' availability and unavoidably decay peers' downloading experience. Therefore, we discuss the main challenge in this design: "*How to control the neighbor selection of more peers and minimize the impact to torrents' availability at the same time?*". Fortunately, the great popularity of multiple tracker configuration [2] gives us the opportunity to minimize the impact to torrents' availability. We thus formulate the tracker blocking problem coherently with the torrents' availability under this scenario. The problem is then addressed through a novel selective tracker blocking algorithm, which iteratively improves traffic locality with a given availability threshold. Our evaluation shows that it can successfully reduce the cross-ISP traffic in the presence of noncooperative trackers and yet with minimal impact to torrents' availability.

The rest of this paper is organized as follows: In section 2, we illustrate the related works. We discuss the existence of noncooperative trackers and our motivation in section 3. In order to address the problem, we formalize the problem and proposed a selective tracker blocking approach in section 4, and Section 5 presents our trace-based evaluation. We further discuss some piratical issues in Section 6 and conclude the paper in Section 7.

2 Related Works

There have been numerous studies on the analysis, optimization, and implementation, of the BitTorrent system[3]. P2P locality has recently attracted particular attention following the pioneer work of Karagiannis et al. [4]. For example, Blond et al. [5] showed through a controlled environment that high locality values (defined by [4]) yield up to two orders of magnitude savings on cross-AS traffic. Xie et al. [6] suggested cooperation between peer-to-peer applications and ISPs by a new locality architecture, namely, P4P, which can reduce both the external traffic and the average downloading time. Choffnes et al. also proposed Ono, a BitTorrent extension that leverages a CDN infrastructure, which effectively locates peers that are close to each other. A recent study from Ren et al. [7] also confirms the possible benefit of a topology-aware and infrastructure-independent BitTorrent protocol.

[1] The Pirate Bay is a Swedish website that indexes BitTorrent contents. It has been involved in a number of lawsuits generally due to the violation of copyright laws. Meanwhile, the ISPs can hardly modify the Pirate Bay trackers for traffic locality.

On the other hand, many studies also addressed some pitfalls of the locality mechanism. Piatek et al. [8] shown that a ”win-win” outcome is unlikely to obtain for all the ISPs during the locality; the reason is that reducing inter-domain traffic reduces costs for some ISPs, while it also reduces revenue for others. Cuevas et al. [9] further investigated the maximum transit traffic reduction as well as the ”win-win” boundaries across the ISPs. These studies indicate that the ISPs are more likely to be selfish during the locality deployment, especially given the considerable gain of traffic locality.

Our work extends these studies through an Internet-wide measurement. We show that the peer control is a very important issue to reduce the cross-ISP traffic. The existence of noncooperative trackers may greatly reduce the efficiency of traffic locality. Given by the popularity of multiple tracker configuration, we propose a selective tracker blocking approach that can effectively help the ISPs to control more traffic while minimizing the impact to the torrents' availability.

Table 1. Top-10 Most Popular Trackers

Rank	Peers	Torrents*	Tracker Sites (URLs)
1	607987	19915	open.tracker.thepiratebay.org
2	593205	16724	trackeri.rarbg.com
3	560580	23386	denis.stalker.h3q.com
4	509140	15308	tpb.tracker.thepiratebay.org
5	504173	12117	vtv.tracker.thepiratebay.org
6	442708	12821	vip.tracker.thepiratebay.org
7	414095	10019	eztv.tracker.prq.to
8	262991	6079	tracker.prq.to
9	184843	3016	tk2.greedland.net
10	142220	3114	www.sumotracker.org

*Note that the torrent level popularity is obtained from the metainfo files which can include multiple trackers.

3 Motivation: Problem of Peer Control

To investigate the deployment of traffic locality, we conduct a measurement over 3-month and collect the information for more than 9 million BT peers. Our measurement configuration and the raw dataset (including the torrents information) can be found at: http://netsg.cs.sfu.ca/n_tracker.htm.

It is well known that the trackers play a very important rule for the deployment of traffic locality. Table 1 presents the site information of the Top-10 most popular trackers

in our measurement. We can see that many of them are belonging to Pirate Bay and etc., which are involved in a series of lawsuits, as plaintiffs or as defendants. In terms of traffic locality, unless the copyrights and other related problems are well solved, we can hardly expect to organize these tracker sites together for traffic locality optimization. Therefore, we use the term *noncooperative tracker* to refer to them; intuitively, if the ISPs cannot modify these trackers, they will also fail to control the peers that managed by the trackers.

In order to quantify their impact in deployment, we investigate the tracker blocking problem in AS#3352 as an example. This is the most popular AS with $126,133$ peers in our measurement; these peers are distributed in $6,065$ torrents that managed by 384 trackers.

In this case study, we pick out a set of trackers that are managed by private organizations. In particular 4 trackers belong to Demonoid and 4 trackers belong to Pirate Bay. Except these noncooperative trackers, we assume that all other trackers have already been modified by the ISPs for traffic locality. Figure 1 shows the probability that the peers in AS#3352 will connect to the modified biased trackers (the probability that their traffic will be optimized). The detailed data can be found in Table 2. (where in Case A, all the trackers can be modified for traffic locality; Case B, 4 noncooperative trackers from Pirate Bay cannot be for traffic locality; and Case C, 8 noncooperative trackers from Pirate Bay and Demonoid cannot be for traffic locality).

We can see that in Figure 1(a) when 4 noncooperative trackers (from Pirate Bay) cannot be modified, only $60,331$ peers in this ISP will be optimized by traffic locality for sure. We find that, 53% peers ($67,132$ out of $126,133$) will be affected by the noncooperative trackers; none of these peers will be optimized for sure and the peers will connect to the biased trackers with relatively low probability. Moreover, as shown in Figure 1(b), if more trackers are not willing to cooperate (8 in this case), the ISP will lose the control of more peers.

We can see that even a small number of noncooperative trackers can bring noticeable damage to the ISPs; a large number of noncooperative trackers many easily ruin the deployment of traffic locality. However, control the neighbor selection of numerous peers one by one is even more troublesome. It is thus important to see whether we can enhance the peer control for the ISPs.

Table 2. Number of peers that will choose the modified biased trackers in certain probability (In AS#3352)

Pr	**0.1-0.2**	**0.3-0.4**	**0.5-0.6**	**0.7-0.8**	**0.9-1.0**
Case A	0	0	0	0	126133
Case B	2360	7317	28685	23279	60331
Case C	11751	19672	32448	7469	47347

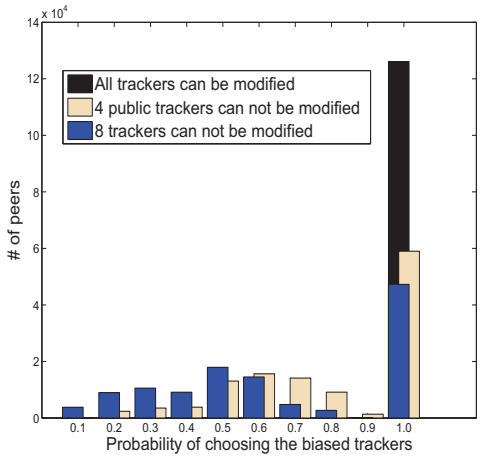

Fig. 1. The impact of noncooperative trackers

4 Tracker Blocking Problem

Intuitively, the ISPs can prevent their peers to connect to these noncooperative trackers via tracker blocking. Yet, this approach will also reduce the torrents' availability and the experience of P2P users. In our study, we find that block all the noncooperative trackers will greatly reduce the availability of torrents (see detailed discussions in Section 5). Fortunately, the latest BitTorrent metainfo file can include multiple tracker sites stored in the *announce-list* section [10]. In particular, unless we block all the trackers in torrents' announce-lists, their availability will not be affected.

In our measurement, we also record the announce-list of the torrents, pick out the cited trackers and then compute the number of trackers that have been used by the torrents. Figure 2 confirms that more than 80% torrents have specified at least two trackers for the load balance (or backup) purpose, and a few torrents even have announce-lists of several hundred trackers. This is much higher than an earlier measurement in 2007 [11] (observed multi-trackers in 35% of the torrents), and thus suggests that the multiple tracker configuration has been quickly recognized and deployed in the BitTorrent community. We thus formulate the tracker blocking problem coherently with the torrents' availability under the scenario of multiple tracker configuration.

4.1 Problem Formulation

We now give a formal description of the tracker blocking problem in BitTorrent networks. We use $\aleph$ to denote all the ASes on the Internet, $\Re$ to denote the set of existing torrents and use $\Im$ to denote the set of trackers. We define three variables A, S and T; A takes on values over the set of ASes $a \in \aleph$; S takes on values over the set of torrents $\Re$,and T takes on values over the set of trackers $\Im$ that managing the torrents.

Based on the above components, two relationships can be learnt from the measurement: (1) The relationship between S (torrents) and T (trackers); (2) The relationship

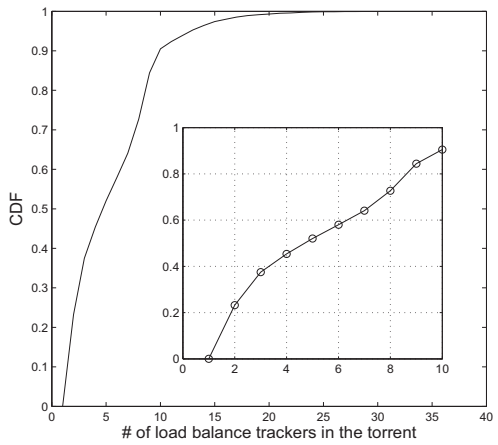

Fig. 2. Popularity of multiple tracker configuration

between A (ASes) and S (torrents). We use binary matrix $R^{tor,tra}$ to define the relationship between S and T, and this matrix is learnt directly from the metainfo (.torrent) files; each element of $R^{tor,tra}(s,t)$ is of a binary value, indicating whether torrent s includes tracker t in its metainfo file (1-Yes, 0-No). On the other hand, we use matrix $R^{tor,as}$ to refer to the frequency table of S and A. This matrix is learnt from the IP/AS information of the probed BT peers; each element $R^{tor,as}(s,a)$ is an integer which refers to the number of peers (in torrent s) that belong to AS a. The peers' AS information is learnt by the 'whois' command on the Linux system, and most replies are from 'whois.cymru.com'. A local view of $R^{tor,as}$ is shown in Figure 3.

For a given ISP x (generally includes n ASes $A_x = \{a_1, a2, ..., a_n\}$, $n \geq 1$) that managed by k trackers $T_x = \{t_1, t2, ..., t_k\}$; We use the set ${T_x}^{modified}$ to refer the trackers that are modified for traffic locality, set ${T_x}^{blocked}$ to refer the trackers that are blocked by the ISPs, and set ${T_x}^{unchanged}$ to refer the trackers that are unchanged (neither modified nor blocked) during the locality deployment. Note that in this definition, the set of noncooperative trackers ${T_x}^{noncoop} = {T_x}^{blocked} \cup {T_x}^{unchanged}$.

Note that when we decide to block a noncooperative tracker, this decision will bring ISPs certain benefits and costs. In particular, the benefit is that some peers will switch to their alterative trackers that can be modified for traffic locality; the cost, on the other hand, can be qualified by the number of peers that cannot find an alterative trackers and thus fail to connect to the BT networks. The tracker blocking problem is that for ISP x (with a known set of noncooperative trackers ${T_x}^{noncoop}$), how to division the tracker set T_x into three parts of ${T_x}^{blocked}$, ${T_x}^{unchanged}$, and ${T_x}^{modified}$. This division should maximize the total benefit; meanwhile, the *torrents' availability* (total number of peers that cannot connect to the BT networks due to the tracker unavailability) should also be bounded by a given threshold β; where $\beta \in [0,1]$ refers to the percentage of peers that become unavailable. We formulate this problem step by step as follows:

For any division, the peers in a given torrent (for example torrent s_1) will be optimized by the locality mechanism with the probability of:

$$P_{s_1} = \sum_{t \in T_x{}^{modified}} R^{tor,tra}(s_1, t) / \sum_{t \in T_x} R^{tor,tra}(s_1, t) \tag{1}$$

This probability is computed by the ratio of modified trackers and the total number of trackers that the torrent cite in its metainfo file. Therefore, a probability distribution of each torrents S can be given by $P_s = \{P_{s_1}, P_{s_2}, ..., P_{s_i}\}$. This distribution describes the probability that the peers(in each torrents) will connect to the bias trackers. On the other hand, for ISP x, the distribution of peer population across the torrents is given by:

$$D_s = \sum_{a \in A_x} R^{tor,as}(S, a) \tag{2}$$

Based on these two distributions, we use a normalized expectation to qualify the total benefit. The tracker blocking problem is therefore to maximize:

$$E(x) = \sum_{S} (D_s \cdot P_s) / \sum_{S} D_s \tag{3}$$

$$s.t. \qquad \frac{V(x)}{\sum_S D_s} \geq \beta \tag{4}$$

where $V(x) = \sum_S \left\{ \lceil [\sum_{a \in T^*} R^{tor,tra}(S, t)]/k \rceil \cdot D_s \right\}$ indicating the total number of peers that can at least connect to one available tracker. $T^* = T_x - T_x^{blocked}$ refers to the trackers that are not blocked. The physical meaning of $E(x)$ is the average probability that the peers in ISP x will be optimized by the traffic locality; meanwhile, the constraint in eqn.4 helps us to bound the torrent availability during the tracker blocking.

4.2 Which Trackers Should be Blocked?

In this part, we will discuss the details of tracker blocking algorithm. It is easy to see that this problem can be transformed into a restricted napsack problem which known as NP complete. Note that in the real-world deployment, both $R^{tor,tra}$ and $R^{tor,as}$ cloud be dynamic (we will further discuss this impact in Section 5). Therefore, instead of finding an optimal solution, we are focusing on a more efficient algorithm for real-world implementations.

To this end, we design an algorithm that can improve the solution quality iteratively. As shown in , in each liberation, we compute the blocking benefits B and the costs H for all the unblocked noncooperative trackers (as discussed in eqn.3 and eqn.4). If the total cost is less than β, then we block the tracker with maximal W ($W = B - H$) in this liberation. After this blocking, we recompute the blocking benefits/costs (B and H) for the remaining trackers and goto next round. Otherwise, if the availability constraint was violated or all the noncooperative trackers are blocked, stop and exit. In next section, we will show that our solution can bring considerable benefits for the real-world ISPs.

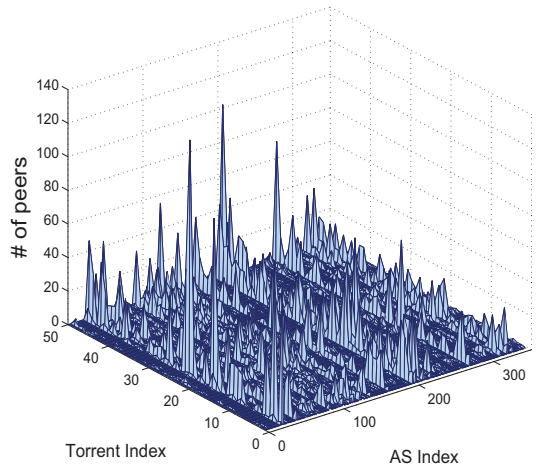

Fig. 3. Local view of matrix $R^{tor,as}$

5 Evaluation

In this section, we will evaluate the performance of our tracker blocking algorithm based on real trace. First, we will discuss the discuss the improvement of peer control for the ISPs. After that, we will further investigate the traffic saving as well as the possible degrading of peers' availability.

We evaluate the tracker blocking problem in a real ISP with multiple ASes. This ISP includes AS#3352, AS#6461, AS#3301, AS#3243, and AS#2847 with $402,496$ peers that managed by 428 trackers. Among all these trackers, we find 104 noncooperative trackers that can hardly cooperate with the ISPs (based on the hosting of copyright contents). We set $\beta = 90\%$ in this simulation, which indicates the ISP need to guarantee the access of 90% peers to BT networks.

Table 3. Number of peers that will choose the modified biased trackers in certain probability

Pr	0.1-0.2	0.3-0.4	0.5-0.6	0.7-0.8	0.9-1.0
Case1	30494	82496	137533	24100	52952
Case2	20151	51943	108906	52819	154257

Based on the our algorithm, 10 noncooperative trackers are selected (blocked) in this ISP. Following two cases are discussed to evaluate its performance: Case #1. All the noncooperative trackers are neither modified nor blocked by the ISPs; Case #2. A

Algorithm Tracker Blocking()

1: **while** the tracker list is not empty and the total availability cost is equal
2: or less than β
3: **for** $\forall$ tracker $t \in T^{noncoop}$
4: compute E of blocking this tracker based on $eqn.3$
5: compute V of blocking this tracker based on $eqn.4$
6: find the tracker with maximal $W = B - H$
7: **if** the tracker is not checked and the total cost
8: after blocking this tracker is equal or less than β
12: **then**
9: remove this tracker from tracker list (this tracker is blocked)
10: update this information in $R^{tor,tra}$ and $R^{tor,as}$.
12: **else**
13: mark this tracker as checked and **goto** $6:$
14: **end if**
15: **end for**
16: **end while**

Fig. 4. The selective tracker blocking algorithm

selective set of noncooperative tracker is blocked by the ISP while other noncooperative trackers are unchanged (and not blocked); Note that Case#2 is the proposed method; the comparison of these two cases indicates the benefit of our selective blocking approach.

Figure 5 shows the cooperation of these two cases. In particular, the $E(x)$ values stress the gain is over 50% where the $E(x)$ in Case#1 is 0.4422 and $E(x)$ in Case#2 is 0.6809 (recall that this value refers to the probabilities that the peers will be optimized in each cases). In this figure, we can see that if we simply ignore these noncooperative trackers (Case#1), only 6.45% (25952 out of 402496) peers will be optimized for sure and most peers will be optimized with the probabilities no greater than 0.5. On the other hand, with our selective blocking approach, 38.22% (153833 out of 402496) peers will be optimized for sure. The detailed data of this figure can be found in Table 3.

We also compute the case when all the noncooperative trackers are blocked. In this case, only 28.18% ($113,419$ out of $402,496$) peers are able to connect to the BT networks. This less than our β value which is set to 90%. The result indicates that the over-blocking of noncooperative trackers will greatly harm the users' experiences and can hardly be appreciated.

The improvement of peer control is encouraging especially considering the overhead of blocking 10 trackers. To further investigate the saved cross-ISP traffic, we perform another simulation using the discrete-event BitTorrent simulator developed by Stanford University [12] as [1] did; we summary the key network settings as follows:

The network contains 13 ASes and $1,600$ BT peers, where the peers are managed by 80 logic trackers and 75 of them are noncooperative trackers. The peers are skewedly distributed among these trackers as we observed in our measurement. When one tracker is blocked, the peers will randomly switch to an alternative tracker as described in the standard BT protocol. Note that in this simulation, each AS from an individual

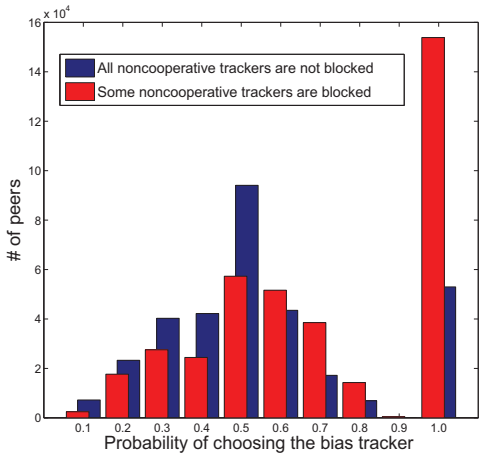

Fig. 5. Improvement of peer control

ISP. These ISPs will all benefit from the tracker blocking approach, and we compute the overall inter-AS (ISP) traffic to quantify this benefit. For other detailed configurations, all peers inside the ISPs are modeled as nodes behind cable modem and DSl, and have asymmetric upload/download bandwidth. The upload bandwidth of these peers is $100k$bps and downloading bandwidth is 1Mbps. Considering the peer arrival/departure, most peers are joining the network at once, i.e. the flash crowd scenario. We focus on this feature since it is the most challenging for ISPs to handle. For each torrent, there is one original seeder that will always stay online (with 400Kbps uplink bandwidth), and other peers will leave the BT network forever as soon as they finish downloading.

We run multiple simulations to average the randomness and the results are shown in Figure 6 and Figure 7 with different β values. Figure 6 present the traffic saving of our algorithm. We can see that when no trackers are blocked with $\beta = 1.0$ (only 5 out of 80 trackers are modified for traffic locality), the ratio of cross-ISP traffic is around 90%. As we decrease the β value and block more noncooperative trackers, more peers will switch to the modified trackers, and the cross-ISP traffic will noticeable decreased. Compare to Figure 7, we can see that the blocking of 20 trackers (β changed from 0.93 to 0.84) will generally reduce the 17% traffic across the ISPs. It is also worth noting that when we block more trackers, the peers' availability is linearly decreasing. However, the cross-ISP traffic is more slowly decreasing with increased standard deviation. This result also indicates the inefficiency of over-blocking (see details in Table 4).

6 Further Discussions

This paper takes a first step towards the tracker blocking problem for the traffic locality. There are still many piratical issues that can be further explored.

First, the trackers are dynamic, which can be considered in the locality deployment. We are currently probing the availability of more than 700 trackers (we have already

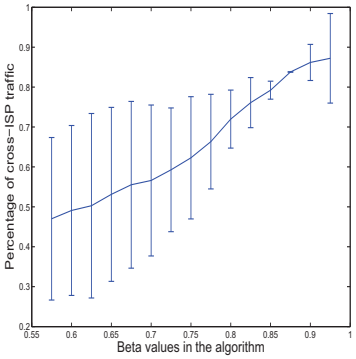

Fig. 6. Improvement of traffic locality

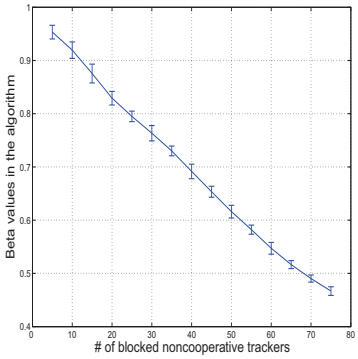

Fig. 7. Degrade of peer availability

Table 4. Facts of views (in Figure 6)

Cross traffic	$\beta = 0.9$	$\beta = 0.8$	$\beta = 0.7$	$\beta = 0.6$
Max	0.887	0.837	0.784	0.738
Min	0.809	0.627	0.441	0.341
Median	0.833	0.677	0.456	0.368
Mean	0.861	0.719	0.566	0.491
Std	0.045	0.072	0.189	0.212

obtained the data for over 8 months), as shown in Figure 8. We cluster the trackers in three classes: 1) Highly available trackers (with the total online time more than 7 months); 2. Normally available trackers (with the total online time more than 0 and less than 7 months); 3. Not available trackers, as shown in Figure 9. We can see that more than 30% trackers have very good availability. These trackers are more eligible to be blocked (if noncooperative) during the locality deployment. We are currently trying to add this information and further improve our model.

Second, our model is based on the measurement information of $R^{tor,tra}$ and $R^{tor,as}$. Therefore, an inefficient dataset may potentially reduce the benefit. Our ongoing work is to use the relationship between ASes and trackers to enhance our model. Note that this relationship can be easily computed/probed by the ISPs. We find that this relationship is quite consistent over time. More importantly, it can also be used to infer $R^{tor,tra}$ and $R^{tor,as}$ when either of which is missing.

7 Conclusions

In this paper, we studied the deployment traffic locality when some tracker sites cannot be modified by the ISPs. Due to the existence of these noncooperative trackers, the ISPs will lose the control of many peers and unavoidably reduce the efficiency of traffic locality.

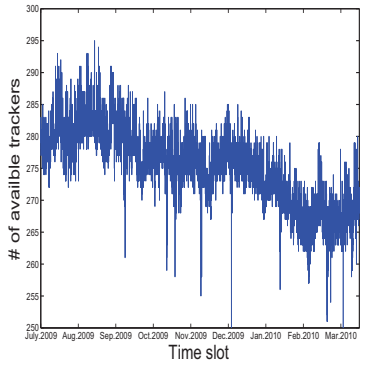

Fig. 8. # of available trackers over time

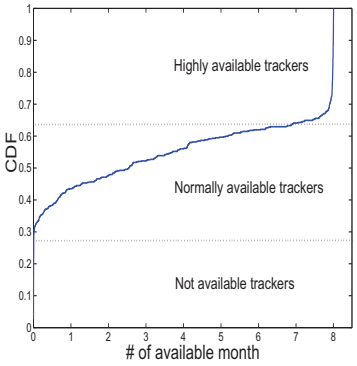

Fig. 9. CDF of tracker availability

We show that a selective tracker blocking approach can well address this problem, and thus formulate the tracker blocking problem coherently with the torrents' availability under the scenario of multiple tracker configuration. Our trace-based evaluation show that this solution successfully reduces the cross-ISP traffic in the presence of non-cooperative trackers and yet with minimal impact to torrents' availability.

References

1. 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. IEEE ICDCS, 2006*.
2. J. Liu, H. Wang, and K. Xu, "Understanding Peer Distribution in Global Internet," *IEEE Network Magazine, 2010*.
3. D. Qiu and R. Srikant, "Modeling and Performance Analysis of Bit Torrent-Like Peer-to-Peer Networks," in *Proc. ACM SIGCOMM, 2004*.
4. T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?" in *Proc. ACM/USENIX IMC, 2005*.
5. S. L. Blond, A. Legout, and W. Dabbous, "Pushing BitTorrent Locality to the Limit," *INRIA Tech, Rep., 2008*.
6. H. Xie, R. Y. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4p: Provider Portal for Applications," in *Proc. ACM SIGCOMM, 2008*.
7. S. Ren, E. Tan, T. Luo, S. Chen, L. Guo, and X. Zhang, "TopBT: A Topology-aware and Infrastructure-independent BitTorrent Client," *Proc. IEEE INFOCOM 2010*.
8. M. Piatek, H. V. Madhyastha, J. P. John, A. Krishnamurth, and T. Anderson, "Pitfalls for ISP-friendly P2P Design," *Proc. ACM HOTNETS, 2009*.
9. R. Cuevas, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep Diving into BitTorrent Locality," *Telefonica Research, Tech, Rep., 2009*.
10. BitTorrent Multi-tracker Specification. [Online]. Available: http://www.bittornado.com/docs/multitracker-spec.txt
11. G. Neglia, G. Reina, H. Zhang, D. Towsley, A. Venkataramani, and J. Danaher, "Availability in BitTorrent Systems," in *Proc. IEEE INFOCOM, 2007*.
12. BT-SIM. [Online]. Available: http://theory.stanford.edu/simcao/btsim-code.tgz