
Delft University of Technology
Parallel and Distributed Systems Report Series

**From user experience to strategies: how to survive in a
private BitTorrent community**

Adele L. Jia, Xiaowei Chen, Xiaowen Chu, Johan A. Pouwelse
adele.lu.jia@gmail.com, {xwchen, chxw}@comp.hkbu.edu.hk, J.Pouwelse@tudelft.nl

Completed September 2011.

Report number PDS-2011-004



ISSN 1387-2109

Published and produced by:
Parallel and Distributed Systems Section
Faculty of Information Technology and Systems Department of Technical Mathematics and Informatics
Delft University of Technology
Zuidplantsoen 4
2628 BZ Delft
The Netherlands

Information about Parallel and Distributed Systems Report Series:
reports@pds.twi.tudelft.nl

Information about Parallel and Distributed Systems Section:
<http://www.pds.twi.tudelft.nl/>

Abstract

Many private BitTorrent communities employ Sharing Ratio Enforcement (SRE) schemes to incentivize users to contribute their upload resources. It has been demonstrated that communities that use SRE are greatly oversupplied, i.e., they have much higher seeder-to-leecher ratios than communities in which SRE is not employed. The first order effect of oversupply under SRE is a positive increase in the average downloading speed. However, in this paper we show that the oversupply induces severe side-effects and under SRE 1) users are forced to seed for extremely long times to maintain adequate sharing ratios to be able to start new downloads, and 2) many users have seeded for very long times but are still with low sharing ratios, due to the counter-intuitive phenomenon that long seeding time does not necessarily lead to large upload amount. We propose strategies for peers to gain sharing ratios more efficiently. We further analyze the existing community strategies beyond SRE, which are initially designed to further incentivize contribution. We show that some strategies have limited or even negative effects on the system performance and we propose our remedies.



Contents

1	Introduction	4
2	Methodology	4
3	A general view: the rich is rich and the poor is poor	5
4	Long seeding time: the expense of high downloading speed	6
4.1	Long seeding time, even for poor peers	6
4.2	Possible reasons?	6
5	Why the poor is poor and how to become rich?	9
5.1	Community level	9
5.1.1	The influence of the upload speed	9
5.1.2	The influence of the age	9
5.2	Torrent level	10
5.2.1	Different upload amounts in the very same swarm	10
5.2.2	Possible reasons and how to gain more?	10
6	Basic requirements for private communities	13
7	Effect of existing community strategies and our remedies	14
7.1	Free leech and double seeding bonus	14
7.2	Priority for rich peers	14
7.3	Bonus for publishing new contents	14
7.4	Counting seeding time	14
7.5	Against potential strategic user behavior	15
8	Conclusion	15
9	Appendix: Results from two other communities	16

List of Figures

1	The CDF of user's sharing ratio.	5
2	The CDF of user's seeding and leeching time.	6
3	Upload amount vs. download amount.	7
4	The upload amount vs. seeding time.	7
5	The upload speed vs. seeding time.	8
6	Oversupply in CHDBits swarms	8
7	The CDF of user's average upload speed in CHDBits.	9
8	The CDF of user's age in CHDBits.	10
9	The influence of user age	11
10	The CDF of user's upload amount in one swarm.	11
11	Upload amount vs. seeding time in one swarm.	12
12	Upload amount vs. upload speed in one swarm.	12
13	Upload amount vs. time to start seeding in one swarm.	12
14	The dynamic demand and supply in one swarm.	13
15	The CDF of user's sharing ratio in HDStar.	16
16	The CDF of user's seeding time and leeching time in HDStar.	17
17	The upload amount vs. download amount in HDStar.	17
18	The upload amount vs. seeding time in HDStar.	17
19	The upload speed vs. seeding time in HDStar.	18
20	Oversupply in HDStar swarms	18
21	Oversupply in HDTV swarms	18
22	The CDF of user's average upload speed in HDStar.	19
23	The CDF of user's upload amount in one swarm in HDTV.	19
24	Upload amount vs. seeding time in one swarm in HDTV.	19
25	Upload amount vs. upload speed in one swarm in HDTV.	20
26	Upload amount vs. time to start seeding in one swarm in HDTV.	20
27	The dynamic demand and supply in one swarm in HDTV (HDTV).	20

List of Tables

1 Introduction

BitTorrent is a popular Peer-to-Peer (P2P) protocol for file distribution. A key of its success lies in its Tit-For-Tat (TFT) incentive policy, which works reasonably well in fostering cooperation among downloading peers¹ (also known as *leechers*). However, TFT does not provide any incentive for peers to remain in the system after the download is complete, in order to *seed* the entire file to others. Therefore, peers are free to engage in “Hit and Run” behavior, the scenario under which a peer leaves immediately upon completing a download.

In recent years, there has been a proliferation of so-called *private* BitTorrent communities aimed at incentivizing seeding. These communities employ a private tracker based method that maintains centralized accounts and records the *sharing ratio* of each peer, i.e., the ratio between its total amount of upload and download. Community administrators specify some *threshold* above which all members are required to maintain their sharing ratios. This mechanism is known as *Sharing Ratio Enforcement* (SRE). Community members whose sharing ratios drop below the threshold are warned and then banned from downloading, or even expelled from the community. In this way, it is guaranteed that each peer provides a certain level of contribution to the community.

The main motivation for implementing SRE is to close the gap between bandwidth demand and supply as observed in public communities, where there is significantly more demand than supply [12]. Thus, the basic design goal of SRE is to achieve higher system-wide downloading speed by increasing the bandwidth supply. Several measurement studies have shown that SRE is very effective in increasing supply [5, 11, 12, 16, 8]. For instance, [12] reports seeder-to-leecher ratios that are at least 9 times higher in private communities than in public ones, while downloading speeds are measured to be 3-5 times higher. However, do the users in private community always enjoy the high downloading speed without any adverse effects?

To answer this question, in this paper we consider more performance metrics besides the downloading speed, i.e., the seeding time, the upload speed, the evolution of sharing ratio, etc. These metrics are highly related to user experience, but have never been considered before in previous works. The contributions of this paper are as follows:

- We perform a measurement study on three private communities that provide user-level information including each user’s upload amount, download amount, seeding time, leeching time, and sharing ratio. Among the tens of existing communities, they are the only ones that provide such detailed information.
- We show that while users achieve very high downloading speed, to achieve the adequate sharing ratios they have to seed for much longer times (compared to their downloading times), though most of the time their seedings are not very productive due to the oversupply induced by SRE.
- For users who intend to increase their sharing ratios, our analysis shows that seeding for longer duration is not as effective as increasing the upload speed. To increase the upload speed, besides upgrading the internet connection, users can also try to avoid the extreme oversupply by joining a swarm earlier.
- Given the possible user strategies against SRE, we further analyze the existing strategies adopted by administrators in private communities, which are initially designed to further incentivize contribution. We show that some strategies have limited or even negative effects on the performance and we propose our remedies.

2 Methodology

The private community tracker collects information that is periodically reported by the BitTorrent clients of its users, which it displays in the form of HTML pages available only to its users. The information provided by the trackers are varied across different communities. To support our analysis, we examined over 20 private communities and only 3 of these communities provide the information detailed enough for our analysis i.e., CHDBits, ChinaHDTV, and HDStar. For these three communities, we crawled their trackers and the data sets we obtained are divided into following categories:

¹In this paper, we use *user* and *peer* alternatively to refer to the individual in a private community.




Figure 1: The CDF of user’s sharing ratio.

- *Community-level user profile*: in this data set, we crawl the profile page of each community user and obtain the information of its upload and download amount, seeding and leeching time, sharing ratio, join date, etc.

It should be noted that the seeding time of each user recorded by the tracker is swarm-based instead of user-based: after a user has seeded in 2 swarms for 10 hours, its seeding time will be added $2 \times 10 = 20$ hours. In the later sections, when we calculate the average upload speed of a user, we calculate its per swarm average upload speed, i.e., the total upload amount divided by the swarm-based seeding time. In this way, we get a rough estimation of a user’s seeding time and upload speed. Though more accurate calculation of the seeding time and upload speed would be better, to the best of our knowledge, until now no private communities provide such information and it is impossible to deploy a client and contact every user individually to get the community-level user profile information.

- *Community-level torrent profile*: in this data set, we collect the information of each community torrent, including its seeder and leecher number, the number of finished downloads, and the time this torrent was published.
- *Torrent-level user activity*: in this data set, we choose a newly released torrent with the largest number of participants within the first 24 hours. We follow this torrent for 7 days and record the activity of each user who has participated or is participating in it. The collected information includes each user’s upload amount, download amount, seeding time, and leeching time in this swarm, as well as the time it joins and leaves the swarm.

We analyzed the three communities we measured in detail. Due to the limit of space, we only demonstrate the results of CHDBits. More results can be found in our technical report [cite]. It should be noted that the measurement data sets of the other two communities also demonstrate similar performance. We collect the data during May 2011. At the time of our measurement, CHDBits had 31,547 registered users, 40,040 torrents and total download amount of 24.303 PB.

3 A general view: the rich is rich and the poor is poor

Fig. 1 shows the CDF of user’s sharing ratio in CHDBits. We see that around 15% users have sharing ratios less than 1 (considered as the *poor*), while around 18% users have sharing ratios larger than 5 (considered as the *rich*). The behavior of accumulating a large sharing ratio may be triggered by various motivations, such as altruism, a desire to be part of the rich elite of the community, or a habit of depositing sharing ratio for the future. The rich peers have little worry about staying in the community, since their sharing ratios are far beyond the SRE threshold, which for CHDBits is 0.7. On the other hand, poor peers are at the risk of being expelled from the community due to the relatively small sharing ratios. As a consequence, they need

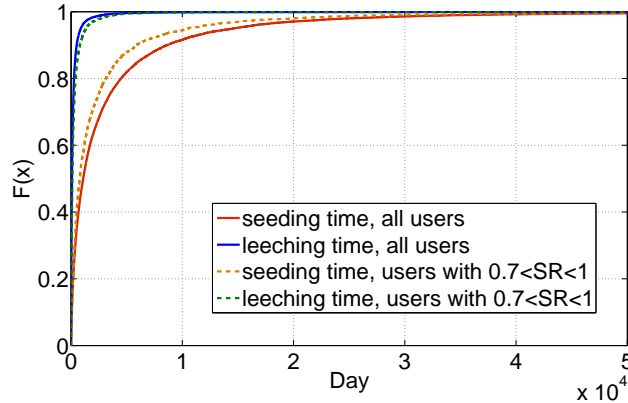


Figure 2: The CDF of user’s seeding and leeching time.

to concern a lot about their movement, such as whether to download new files they really desire but might reduce their sharing ratios to a more risky level.

One may argue that the poor peers are the free-riders, who intend to keep low and risky sharing ratios that are just enough to stay in the community. However, the highly restricted membership in private communities, especially in CHDBits where new members can only join by invitation, makes it very difficult to get a new membership. Hence, we conjecture that not all poor peers are strategic and psychologically strong enough to face the expel from the community due to insufficient sharing ratios. Interestingly, as we will show in the following sections, the poverty is induced by the fact that the poor peers are not strategic enough.

4 Long seeding time: the expense of high downloading speed

Many previous studies have shown that under SRE users seed for long durations [12, 11, 16, 5]. They consider this as a positive effect of SRE since long seeding duration leads to high downloading speed. However, in this section we argue that the long seeding duration can also be a negative effect, especially for poor peers.

4.1 Long seeding time, even for poor peers

Fig. 2 shows the CDF of user’s total seeding time and total leeching time in CHDBits. We see that in general the seeding time is much longer than the leeching time: while 50% users have leeched only for less than 70 days, 50% users have seeded for more than 1100 days. As we pointed out earlier, the seeding time and leeching time are swarm-based: if for each user the number of parallel leechings and the number of parallel seedings are similar, this means that roughly users seed 15 times longer than they leech.

Intuitively, long seeding times for rich peers (in terms of high sharing ratios) are to be expected, since the rich peers are depositing sharing ratios by seeding. However, we observe from Fig. 2 that the poor peers also seed much longer than they leech. While intuitively poor peers should be the ones that are not “hard-working” enough, why some of them seed for long durations but are still with low sharing ratios?

4.2 Possible reasons?

One may argue that the long seeding times for poor peers are due to the fact that users who contribute more also consume more, hence they seed for long durations but are still with low sharing ratios. This argument is partially true. Andrade *et al* [2] has shown and we also observe from our measurement (Fig. 3) that, the individual upload amount (contribution) roughly increases with the corresponding download amount (consume).

However, this doesn’t necessarily mean that the heavy contribution induces long seeding time, nor does it mean that long seeding time leads to heavy contribution. Quite counter-intuitively, as shown in Fig. 4(a), a




Figure 3: Upload amount vs. download amount.




Figure 4: The upload amount vs. seeding time.

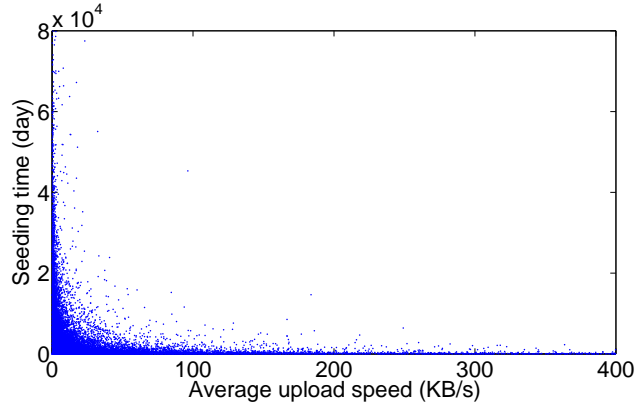
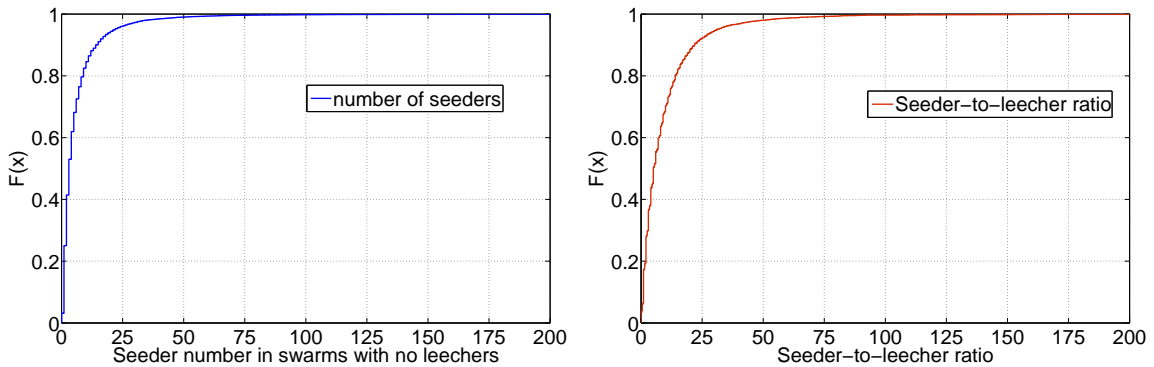


Figure 5: The upload speed vs. seeding time.



(a) The CDF of seeder number in swarms with no leechers (b) The CDF of seeder-to-leecher ratio in swarms with at least one leecher

Figure 6: Oversupply in CHDBits swarms

peer’s upload amount has little relation to its seeding time: many peers seed for long durations but only have uploaded relatively small amount of data, and for some other peers they seed for relatively short durations but have successfully achieved large upload amounts. The same argument is also applicable to poor peers (Fig. 4(b)). This interesting phenomenon implies that, for poor peers who intend to increase their upload amount to become rich, seeding for longer durations may not be an effective method, even if intuitively it seems so.

Though there is no strict relationship between a peer’s seeding time and its upload amount, we do observe that a peer’s seeding time is related to its average upload speed, regardless of its upload amount. As shown in Fig. 5, most of the long seeding durations happen to the peers with relatively small upload speeds, and for peers who have high upload speeds, their seeding times are normally short. We conjecture the possible reason is that for peers in private communities, they need to seed at least for a certain amount to keep adequate sharing ratios to stay in the community, and while $upload\ amount = time \times upload\ speed$ always hold, a higher upload speed means a smaller seeding time for the compulsory upload required by SRE.

The most intuitive reason for a low upload speed is a limited internet access. However, we argue that the limited internet access is not the only reason for the low upload speed. At the time point when we crawled the site, CHDBits has 33041 active swarms (with at least one leecher or one seeder), among which 26402 swarms (79.9%) have no leechers at all. As shown in Fig. 6(a), 40% of the swarms with no leechers still have at least 5 seeders, and 5% of these swarms even have more than 20 seeders. For swarms with at least 1 leecher, the seeder-to-leecher ratio (SLR) is quite high: as shown in Fig. 6(b), 50% of these swarms have a SLR larger than 6, and 5% of these swarms even have a SLR larger than 30! We see clearly that a majority of the swarms in CHDBits are heavily oversupplied. In such swarms, seeders are not able to

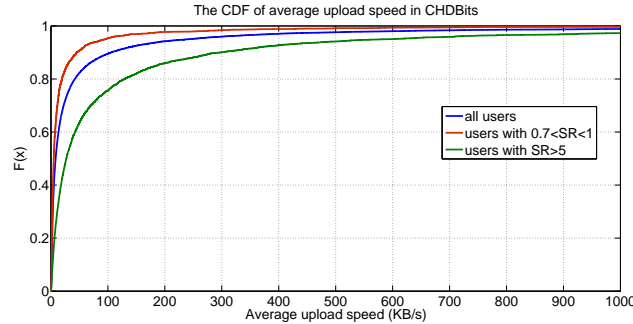


Figure 7: The CDF of user’s average upload speed in CHDBits.

perform any actual uploads due to the insufficient demand and unsatisfied supply. We term this situation as *unproductive seeding*. As a consequence, users have to seed for extremely long durations to achieve the sharing ratio required by SRE.

While a low upload speed mainly leads to a long seeding time, in the next section we show its influence on a user’s status. We analyze the reasons for the poor being poor and we propose strategies for users to become rich efficiently.

5 Why the poor is poor and how to become rich?

As sharing ratio is defined as the ratio between a peer’s upload and download amount, two possible reasons for a peer being poor are that it has downloaded too much or has uploaded not enough. The download amount depends on user’s interests in files. We do not suggest users to download less so as to become rich, since the fundamental user experience should be guaranteed by communities is that users should not need to limit their download needs. Following this argument, in this section we focus on user’s upload activity and analyze why some users have uploaded not enough (hence is poor) and how they can improve it (to become rich).

5.1 Community level

5.1.1 The influence of the upload speed

Inspired by the previous analysis that the seeding time has little influence on the upload amount but the upload speed does, we conjecture that the upload speed will further influence whether a user is rich or poor. Our conjecture is validated by the measurement result shown in Fig. 7. In general, rich peers (in our case peers with $SR \geq 5$) have much higher upload speeds than poor peers ($SR \leq 1$). For example, 80% of the poor peers upload at a speed less than 20 KB/s, while at least 40% rich peers can upload at a speed larger than 50 KB/s. Together with previous result, we conclude that instead of seeding for longer durations, peers who intend to become rich should seed with higher upload speeds. And to seed with a higher upload speed, a user could either upgrade its internet access, or choose a swarm less oversupplied.

5.1.2 The influence of the age

While CHDBits has existed for many years, one may argue that the age of a peer may also influence its status. One possible argument could be, for the old peers, they spend time long enough in the community, hence they have the time and chance to improve their status so that they may be rich. Another possible argument could be, for the young peers, they are blank-paper-like and hence they might be rich since they don’t have much bad history yet. Both arguments are reasonable but they understand the problem from different angles. Fig. 8 shows the measurement result of this problem.

From the measurement result we see that, in general old peers are richer: 70% of the richers peers are less than 500 day old while 70% of the poor peers are less than 650 day old. This result implies that as the users getting older, they contribute more and are with relatively higher sharing ratios. But one can also

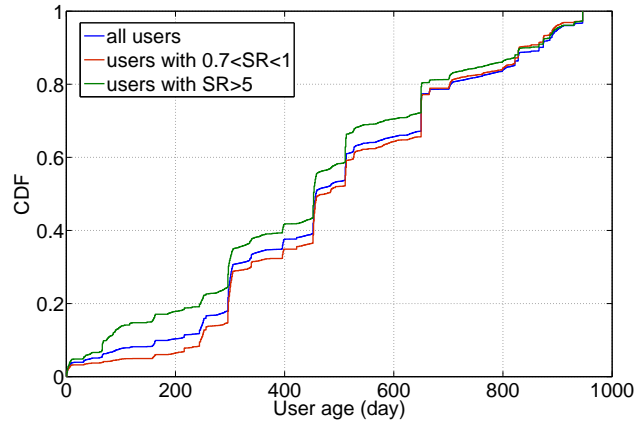


Figure 8: The CDF of user’s age in CHDBits.

understand this from the fact that, for those peers with low sharing ratios, the community first gives them the opportunity to catch up, then after a certain time, they are expelled if they failed. This explains the fact that very poor peers that are young can still stay in the community, but it’s rare to see poor peers that are very old since most of them are already expelled from the community.

Meanwhile, we see from Fig. 8 that there are several sudden burst in population on certain days, like around day 300, 450, 500, etc. We believe these are the few periods that CHDBits allows open registration. For the rest of the days, CHDBits only admits registration by very limited invitation.

In line with the discussion in this subsection, we show three further results related to the influence of user age in Fig. 9.

One may argue that the above analysis is based on the community-level download and upload activities, which might be influenced by the number of swarms a user has participated in. In the following subsection, we focus on a single swarm and demonstrate the torrent-level user experience and strategies to become rich.

5.2 Torrent level

5.2.1 Different upload amounts in the very same swarm

Fig. 10 shows the CDF of user’s upload amount in one single swarm, from which we observe a huge difference: while most of the users are only able to upload very little, a small fraction of users upload a lot. For example, 60% of the users have uploaded less than 10 GB, which is less than the amount they have downloaded (11.7 GB). On the other hand, 5% of the users have uploaded more than 50 GB. Apparently, the users who managed to upload more will become richer. While these users have participated in the very same swarm, why some managed to gain a lot while some didn’t?

5.2.2 Possible reasons and how to gain more?

One intuitive reason for a small upload amount is a short seeding time. However, similar to previous analysis, again we find a counter-intuitive result: in one single swarm a peer’s upload amount is not related to its seeding time (Fig. 11). On the other hand, it is influenced by its upload speed. As shown in Fig. 12, most of the small upload amounts happen to the peers with relatively small upload speeds, and for peers who have high upload speeds, their upload amounts are normally high.

When we organize the peers according to the time they start to seed, we find another interesting phenomenon: peers that start to seed earlier normally have uploaded more (Fig. 13). The same phenomenon has also been observed by Kash *et al* in [10]. One may argue that the peers who start to seed earlier can seed for longer durations, hence they upload more. However, in Fig. 11 we already show that the upload amount is not related to the seeding time. Then why peers that start to seed earlier upload more?

Again we conjecture that this is due to the unbalanced demand and supply in the swarm. As shown in Fig. 14(a), after the burst at the beginning two hours since the file was published, the peer arrival rate

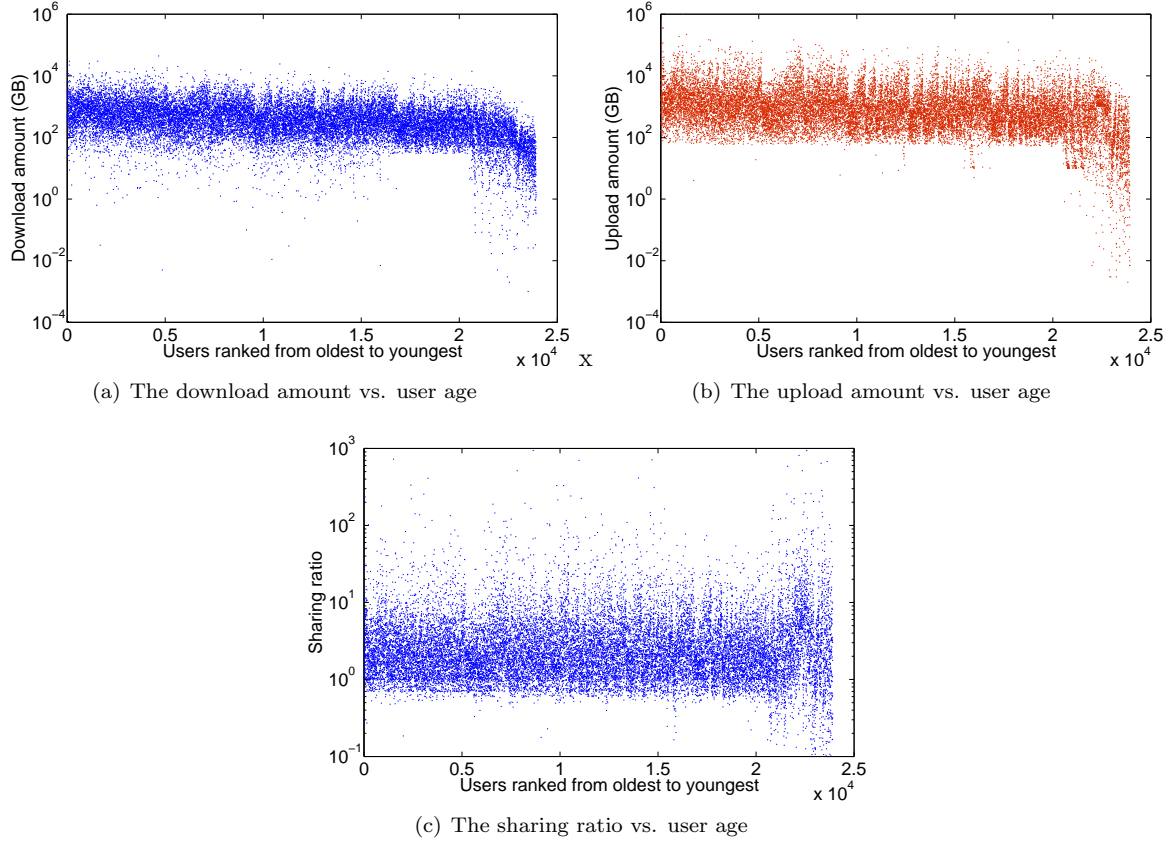


Figure 9: The influence of user age




Figure 10: The CDF of user's upload amount in one swarm.

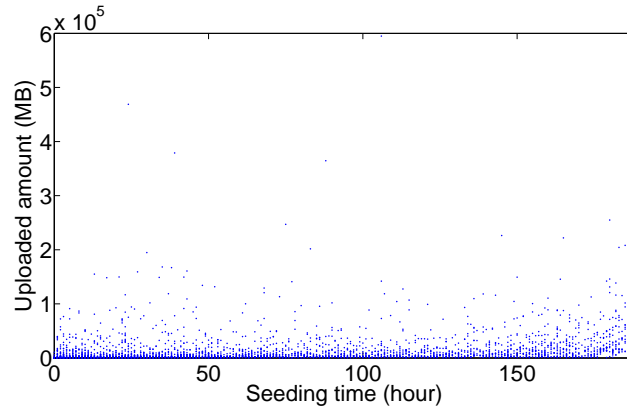


Figure 11: Upload amount vs. seeding time in one swarm.

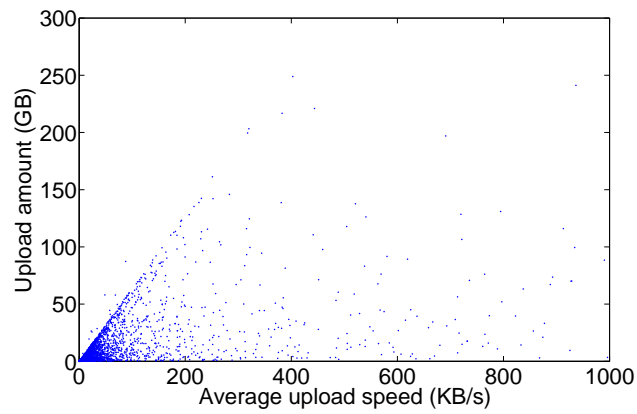


Figure 12: Upload amount vs. upload speed in one swarm.

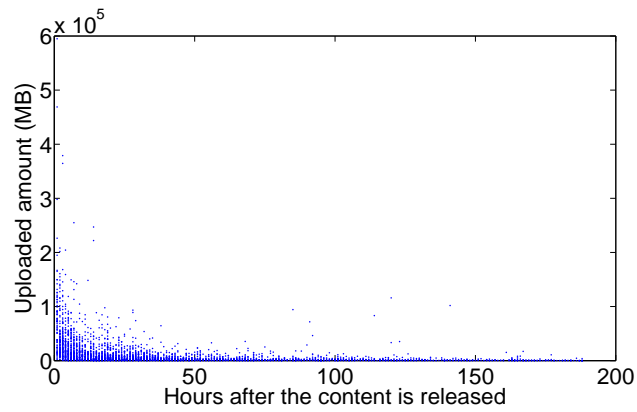


Figure 13: Upload amount vs. time to start seeding in one swarm.




Figure 14: The dynamic demand and supply in one swarm.

decreases dramatically. On the other hand, the number of seeders (supply) increases quickly at the first 60 hours, then decreases with a much smaller rate (Fig. 14(b)). In general, the number of leechers is negligible compared to the number of seeders. As a consequence, peers who join late have to compete with a large number of seeders for uploading, which leads to a low upload speed, and hence a small upload amount. Therefore, for peers who intend to become richer, they should join the swarm in the early state when it is still not extremely oversupplied.

6 Basic requirements for private communities

In previous sections we have shown that SRE induces significant negative effects as the expense of the high downloading speed, such as unproductive seeding and long seeding time, which significantly deteriorate the user experience. Should the download speed be the unique and ultimate design goal for private communities? What are the basic requirements for private communities and what are the fundamental user experience should be guaranteed? We summarize our ideas as follows.

High downloading speed: With no doubt, high downloading speed is the fundamental design goal for private communities, and right until now it has been achieved very well.

Fairness towards all peers that put effort in the community: According to the conservation law, while some users achieve high sharing ratios, there must be some other users with relatively small sharing ratios. However, as we demonstrated, users in private communities that are not strategic enough or with limited capacities have difficulties in achieving high sharing ratios or even maintaining the required one, though they behave well and seed all the files all the time.

One may argue that communities have the rights to discriminate against these users, “kick them out”, and let only the strategic users with high capacities stay. However, we argue that when a majority of users are strategic, the system performance will be deteriorated. For example, if many users know it is easier to gain sharing ratios in the early state of a swarm and hence join the swarm immediately after a new content is published, then, 1) many users will download something they don’t want, but only for gaining sharing ratios; 2) the downloading speed in the early state of a swarm will be very low, because a large number of strategic users joining simultaneously makes the swarm heavily *flash-crowded*²; and 3) it will be more difficult to perform any actual uploads after the early state, since only a few non-strategic users will join the swarm during that period.

Good file availability: For some swarms sharing less popular files, users can hardly perform any downloads due to the insufficient number of seeders. It is important to increase the file availability of those swarms.

Given these requirements, in the following section we evaluate the existing strategies adopted by private communities, and we propose our remedies for the ones with limited or even negative effects.

²We refer a swarm to be flash-crowded if at a particular time point it has much more leechers than seeders.

7 Effect of existing community strategies and our remedies

7.1 Free leech and double seeding bonus

From time to time private communities adopt *free leech* or *double seeding bonus* for some swarms [13, 4, 6, 7]. Free leech means that a user can download a particular file for free, i.e., that file will not be added to its community-level download amount. Double seeding bonus means that a user's community-level upload amount will be added twice of the actual amount it has uploaded in that swarm. Both strategies are designed to circulate the credit and have been show to be effective in breaking the crash or crunch in the system [14].

While these strategies are performed manually, administrators should be careful and not adopt it for long time, otherwise strategic users might wait and not download anything until the files are for free. In our previous work, we propose self-organized strategy named *SRE with supply-based price* [9] that prevents this potential manipulation of strategic users. The basic idea is that the price for downloading one unite of data should go with the supply. If a swarm is oversupplied, i.e., with a large seeder-to-leecher ratio, instead of manually adopting free leech (i.e., zero price) by the administrators, SRE with supply-based price automatically decreases the price. Once the supply goes tight again, SRE with supply-based price will automatically increase the price. In this way, the administrators do not need to observe the swarms all the time and try to find out when to stop the free leech. The same argument is also applicable to the case of double seeding bonus.

7.2 Priority for rich peers

Some private communities try to further incentive contribution beyond SRE and they propose that rich peers have the priority to access the newly published contents, or they can seed more files in parallel [1]. However, as discussed previously, joining early in a new swarm and seeding more files will help the users, especially the poor users, gain sharing ratios more efficiently. By giving the priority to the rich peers, administrators are basically taking the opportunities away from the poor peers for gaining sharing ratios. Unless the administrators intend to let the rich be richer and the poor be poorer (which will lead to a more intense competition and a potential deterioration of performance as discussed previously), we suggest administrators to remove these restrictions, and further educate their users that they benefit from seeding more files in parallel. While users are seeding more files in parallel, the file availability in the community will be improved as well.

7.3 Bonus for publishing new contents

To increase the diversity of the files, some private communities provide bonus for publishing new contents. Normally they have specific requirements for the new content being published. For example, it is required to be of high quality, the metadata should be well organized, .etc. In general there are two kinds of bonus for publishing new contents, i.e., *one-time bonus* and *sale-volume-based bonus* [15]. For the first case there is no incentive for users to publish popular files, since publishing any file will gain them the same amount of bonus. For the latter case users are not encouraged to publish good and yet unpopular contents like classical music, since publishing unpopular files will not gain them much bonus. We propose to combine these two bonus modes together, i.e., a constant bonus plus a sale-volume-based bonus, which we believe will help increase the file diversity.

7.4 Counting seeding time

Our previous work [9] has quantitatively shown the performance improvement when taking user's seeding time into account. Many private communities record user's seeding time [4, 6, 3]. However, like CHDBits, they turn a user's seeding time into virtual credit that can be used in some other use, such as giving invitation to friends. On the other hand, we propose that administrators take the seeding time into account when decide whether to ban a user or not: if it has performed a sufficiently long time of seeding, it should be given the opportunity to stay.

7.5 Against potential strategic user behavior

Though altruistic users always exist, we conjecture that most users in private communities are *selfish*. Their initial goal in a community is to download all the files that they are interested in. To achieve this, while not limiting their download needs, they always try to increase their sharing ratios if it is possible. The strategies they'd apply mainly optimize their own benefit, without considering the *social welfare*, i.e., the performance of other users. As we discussed previously, the potential strategies they could apply to gain sharing ratios include 1) seeding all the files, if communities count the seeding time; 2) always joining free-leech swarms; 3) joining new swarms earlier.

The strategies adopted by private communities should take the above potential user behaviors into account. For example, once the administrators of private communities notice that a majority of the users try to join new swarms immediately after they are published, to alleviate the negative effects they should adopt some methodologies to reduce the benefit of joining early, e.g. free leech for old swarms (hence users may join late since they know it will be cheaper). Or administrators can just use the SRE with supply-based price, and let it adjust the price automatically.

8 Conclusion

While previous work only focus on showing the effectiveness of SRE in increasing the supply, in this paper we provide a deeper understanding of private communities by analyzing both the positive and negative effects of SRE. We show that swarms in private communities are greatly oversupplied. Users achieve very high downloading speeds, but at significant expense including extremely long seeding times and very low upload speeds. As a consequence, users with small sharing ratios are forced to limit their downloading needs so as to keep adequate sharing ratios to stay in the community. For those who intend to increase sharing ratios, our analysis shows that seeding for longer durations is not as effective as increasing the upload speed. If it is not realistic for the users to upgrade their internet access, we suggest them to join swarms earlier or to join less oversupplied swarms. We further show that some of the existing strategies adopted by administrators have limited or even negative effects. We suggest the administrators to educate their users the strategies to increase sharing ratios efficiently, give them the opportunity to contribute, and be fair towards all peers that have put effort in the communities. Further, some strategies like SRE with supply-based price that help close the gap between bandwidth supply and demand are also effective in improving user experience.

References

- [1] <http://www.bitsoup.org/>. 14
- [2] N. Andrade, E. Santos-Neto, F. Brasileiro, and M. Ripeanu. Resource demand and supply in bittorrent content-sharing communities. *Computer Networks*, 53, 2008. 6
- [3] BitHDTV. <http://www.bit-hdtv.com>. 14
- [4] CHDBits. <http://chdbits.org/>. 14
- [5] X. Chen and X. Chu. Measurements, analysis and modeling of private trackers. In *Proceeding of IEEE P2P*, 2010. 4, 6
- [6] ChinaHDTV. <http://www.chinahdtv.org>. 14
- [7] HDStar. <http://www.hdstar.org/>. 14
- [8] A.L. Jia, L. D'Acunto, M. Meulpolder, and J.A. Pouwelse. Modeling and analysis of sharing ratio enforcement in private bittorrent networks. In *Proceeding of IEEE ICC*, 2011. 4
- [9] A.L. Jia, R. Rahman, T. Vinko, J.A. Pouwelse, and D.H.J. Epema. Fast download but eternal seeding: the reward and punishment of sharing ratio enforcement. In *Proceeding of IEEE P2P*, 2011. 14

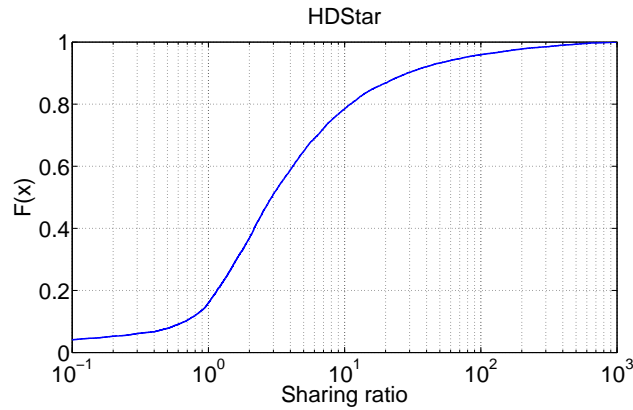


Figure 15: The CDF of user's sharing ratio in HDStar.

- [10] I.A. Kash, J.K. Lai, H. Zhang, and A. Zohar. Economics of bittorrent communities. In *Proceeding of NetEcon 2011*, 2011. 10
- [11] Z. Liu, P. Dhungel, D. Wu, C. Zhang, and K.W. Ross. Understanding and improving incentives in private p2p communities. In *Proceeding of ICDCS*, 2010. 4, 6
- [12] M. Meulpolder, L. D'Acunto, M. Capota, M. Wojciechowski, J.A. Pouwelse, D.H.J. Epema, and H.J. Sips. Public and private bittorrent communities: A measurement study. In *Proceeding of IPTPS*, 2010. 4, 6
- [13] PolishTracker. <http://polishtracker.net/>. 14
- [14] R. Rahman, D. Hales, T. Vinko, J.A. Pouwelse, and H. J. Sips. No more crash or crunch: Sustainable credit dynamics in a p2p community. In *Proceeding of HPCS*, 2010. 14
- [15] TVTorrents. <http://www.tvtorrents.com/>. 14
- [16] C. Zhang, P. Dhungel, Z. Liu Di Wu, and K.W. Ross. Bittorrent darknets. In *Proceeding of IEEE INFOCOM*, 2010. 4, 6

9 Appendix: Results from two other communities

Fig.17 to 27 show the results from two other communities, HDStar and HDTV, which are ordered in the same story line as the ones for CHDBits.

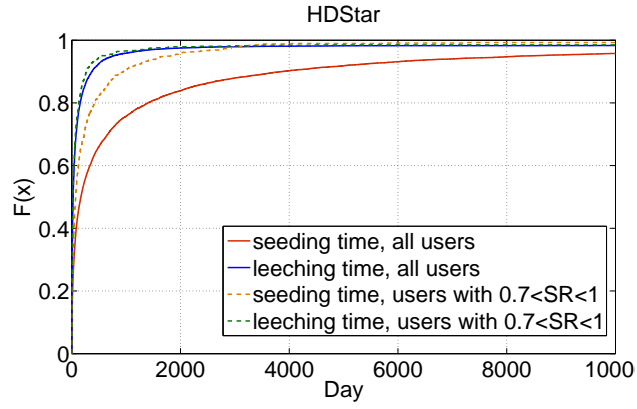


Figure 16: The CDF of user's seeding time and leeching time in HDStar.

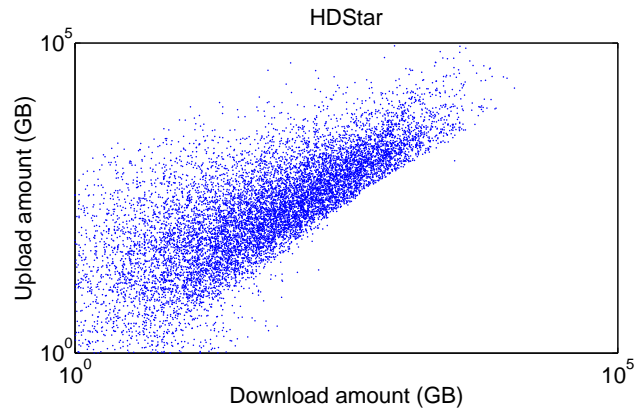


Figure 17: The upload amount vs. download amount in HDStar.

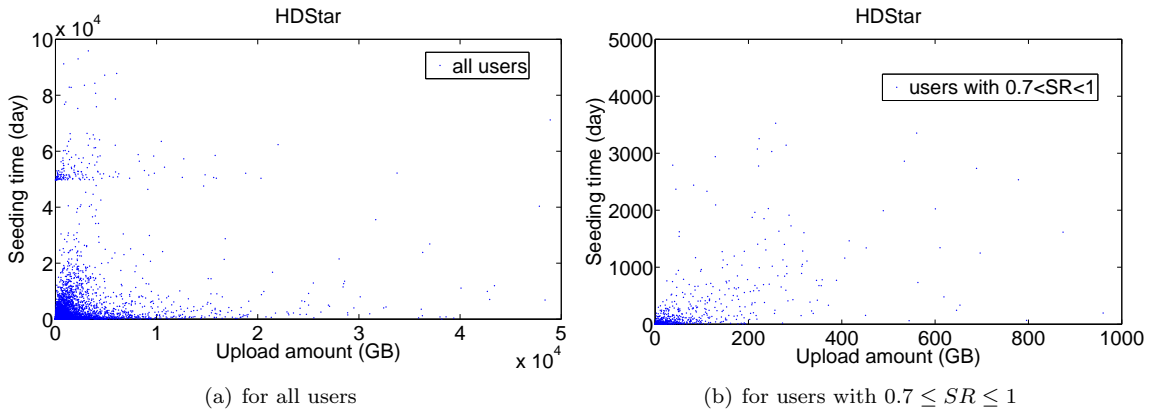


Figure 18: The upload amount vs. seeding time in HDStar.

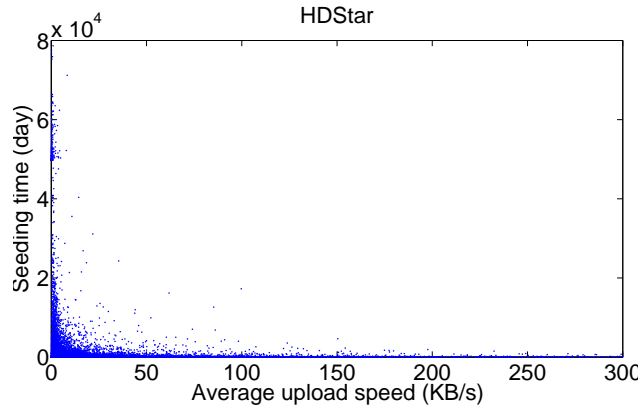




Figure 19: The upload speed vs. seeding time in HDStar.



(a) The CDF of seeder number in swarms with no leechers (b) The CDF of seeder-to-leecher ratio in swarms with at least one leecher

Figure 20: Oversupply in HDStar swarms



(a) The CDF of seeder number in swarms with no leechers (b) The CDF of seeder-to-leecher ratio in swarms with at least one leecher

Figure 21: Oversupply in HDTV swarms

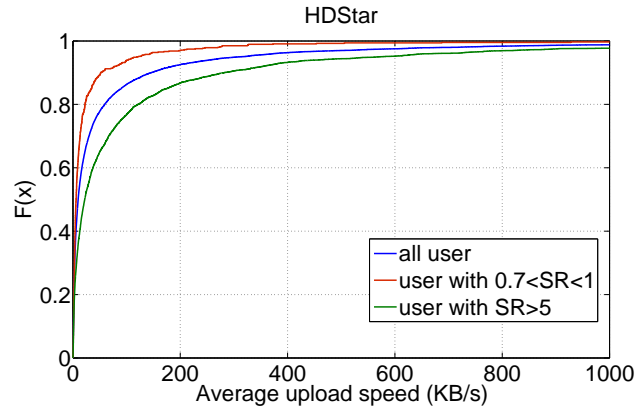


Figure 22: The CDF of user's average upload speed in HDStar.

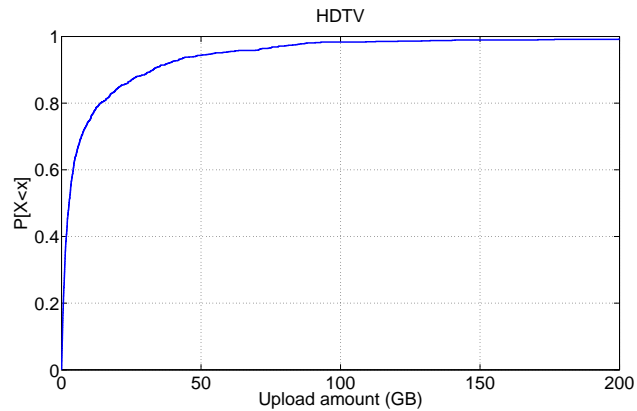


Figure 23: The CDF of user's upload amount in one swarm in HDTV.

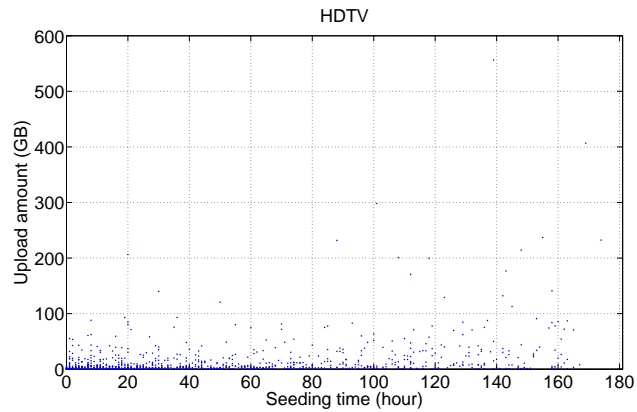


Figure 24: Upload amount vs. seeding time in one swarm in HDTV.

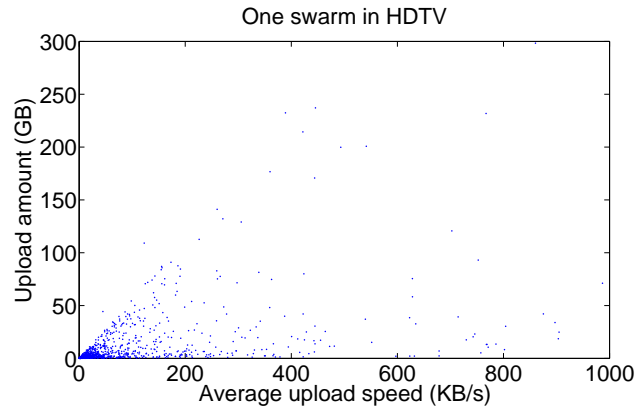


Figure 25: Upload amount vs. upload speed in one swarm in HDTV.




Figure 26: Upload amount vs. time to start seeding in one swarm in HDTV.

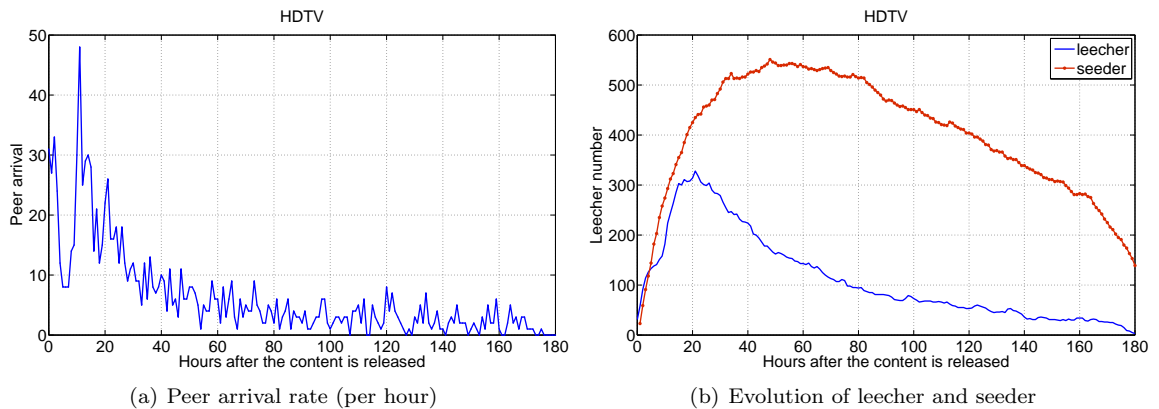


Figure 27: The dynamic demand and supply in one swarm in HDTV (HDTV).