Dissecting Darknets: Measurement and Performance Analysis

Xiaowen Chu, Xiaowei Chen
Hong Kong Baptist University

Adele L. Jia, Johan A. Pouwelse, and Dick H.J. Epema
Delft University of Technology

BitTorrent (BT) plays an important role in Internet content distribution. Because public BTs suffer from the free-rider problem, Darknets are becoming increasingly popular, which use Sharing Ratio Enforcement to increase their efficiency. We crawled and traced 17 Darknets from September 2009 to Feb 2011, and obtained data sets about over 5 million torrents. We conducted a broad range of measurements, including traffic, sites, torrents, and users activities. We found that some of the features of Darknets are noticeably different from public BTs. The results of our study reflect both macroscopic and microscopic aspects of the overall ecosystem of BitTorrent Darknets.

1. INTRODUCTION

BitTorrent (BT) is currently the dominant Peer-to-Peer (P2P) content sharing protocol [Zhang et al. 2011]. It changes the way of content distribution from traditional Client/Server mode to P2P mode, which greatly decreases the burden on file servers. The tracker (together with its communities), which is one of the core components of the BT protocol, plays an important role in the course of content distribution by periodically providing updated peer lists to connected peers. In practice, a peer accesses website-based BT communities to obtain torrents and then to join the content distribution swarms. The BT communities can be divided into two categories: public (a.k.a. public BT) and private (a.k.a. Darknet, private tracker) communities [Zhang et al. 2010]. There are numerous public BT communities, such as ThePirateBay, Mininova, and ISOHunt, which anyone can access. However, there is a major problem with public BTs. Although BT has implemented the Tit-for-Tat (TFT) algorithm as an incentive mechanism, public BTs still suffer from the well-known free-rider problem [Andrade et al. 2005]. First, a peer in a public BT may stop uploading immediately after the download task finishes. Second, peers usually set a limit on their upload bandwidth [Zhang et al. 2010]. As a result, high upload bandwidth doesn't necessarily guarantee good download performance, and peers have no incentive to continue seeding after downloading. Therefore, the supply side of public BT swarms is limited.

In recent years, Darknets have become more and more popular, and there are now almost 1,000 communities on the Internet [Sharky 2011]. Collectively, many more torrents are shared on Darknets than on the existing public BTs [Sharky 2009a]. Our measurements reveal that some of these Darknets account for huge amounts of Internet traffic. For example, Rutracker.org alone has more than 2 Peta-bytes of content (over 1.3 million torrents) and 12.9 million users that can generate as high as several tens of GB/s of Internet traffic. One main reason for popularity of Darknets is that users can usually achieve much faster download speeds than those in public BTs. This can be attributed to the Darknets' strict policies for controlling membership (i.e., Darknets are closed communities and an outsider can only become a registered user through invitation codes) and their adoption of the Sharing Ratio Enforcement (SRE) mechanism as an auxiliary incentive to overcome the free-riding issue. SRE forces registered users to keep their sharing ratios (upload-to-download ratios) higher than the predefined threshold (e.g., 0.8). In this paper, we use the terms "registered user" and "member" interchangeably. In general, a registered user will be banned from the Darknet if his or her sharing ratio is lower than the given threshold. On the contrary, a user with a high sharing ratio will be rewarded, e.g., he is eligible to invite someone else to join the Darknet, or he can use more IP addresses to download/seed. There are two typical ways of achieving a high sharing ratio: uploading as fast as possible and prolonging the seeding time of torrents. As a consequence, Darknet users are motivated to use up their uplink bandwidth and seed as long as possible.

The main differences between the public BTs and Darknets are summarized as follows.

— User Access: Public BTs are open to everyone, but Darknets are only available to registered users.

— Incentive Mechanisms: Public BTs rely on the BitTorrent TFT algorithm, while Darknets implement an additional SRE mechanism to incent users to contribute as much content as possible.

— Traffic Counting: Public BTs do not count users' traffic during content distribution, whereas Darknets accurately record users' total amounts of upload data (T_u) and download data (T_d). Each user's ***sharing ratio*** is then calculated as T_u/T_d. If a user's sharing ratio is lower than a predefined threshold after an initial transition period, he or she will be banned from that Darknet community.

— Downloading Performance: Because of the SRE mechanism, users in Darknets are encouraged to provide a high uploading bandwidth and long seeding times. Therefore Darknets usually have much better downloading performance than public BTs.

Since Darknets have attracted a large population of users and generate huge Internet traffic, it becomes important to understand the characteristics of Darknets to help design better mechanisms and build a more sustainable environment for BT content distribution. This paper presents the results of a measurement study and provides analysis of the collected data sets. Our main contributions to the literature are summarized as follows.

— We performed large-scale measurements covering 17 Darknets, 2 public BitTorrent communities, and 1 BitTorrent search engine from September 28, 2009 to Feb 28, 2011. We obtained 35 data sets covering over 5 million torrents. We made our data sets publicly available to the research community [Chu 2013].

— We conducted detailed analysis of the collected data sets at the community-level and the torrent-level, including traffic, sites, torrents, and user activities. We found that some features of the Darknets are noticeably different from the public BTs. We showed that the SRE is an effective mechanism for encouraging users to contribute as much as possible. However, the SRE alone leads to the phenomenon that many users have difficulties in maintaining the required sharing ratio. Darknets have to rely on freeleech promotion or other artificial mechanisms to help users improve their sharing ratios. E.g., if a torrent is tagged as freeleech by the Darknet administrator, it will not count against the users' download amount, but will count positively towards the upload amount.

The remainder of this paper is organized as follows. Section 2 provides an overview of Darknets. Section 3 describes our measurement methodology and provides an overview of the collected data sets. Section 4 presents the detailed analysis of the community-level data sets, while section 5 analyzes the torrent-level data sets. Related research is presented in Section 6, and Section 7 concludes the paper.

2. OVERVIEW OF DARKNETS

2.1 Taxonomy and Terminology

In this section, we present a taxonomy of Darknets based on the roles and operational processes of the communities. There are two types of roles in Darknet communities, as shown in Table 1. In general, Darknets are established and operated by the Darknet owner, who then uses the system's high downloading performance and newly updated contents to attract users. Outsiders must follow specific protocols to become registered users [Sharky 2009b]. Registered users can download torrents from the Darknet's website and then join the underground BT swarm using some authorized BT client software. All registered users must keep their sharing ratios above a minimum requirement after an initial transition period.

The terminology used in the various Darknets is not standardized. For the purpose of clarity, Table 1 lists some of the terms that are widely used in Darknet communities and we define these terms one by one.

Table 1. Taxonomy of Darknets Based on Role

Role	Operational Process	Related Terms
Darknet user	registration	invitation code
	content sharing	sharing ratio, passkey, user class, freeleech, seeding, leeching, wait time
	community contribution	credit/point mechanism
Darknet owner	building and publishing	codebase
	ranking	pre-time

Definition 2.1 (Invitation Code). The invitation code is a typical registration method in closed or semi-closed Darknets. Contributing users who meet specific requirements are eligible to invite their friends to join the Darknet by sending an invitation code.

Definition 2.2 (Sharing Ratio). Darknets calculate the sharing ratios of their users, which comprise the total amount of data a user has uploaded, divided by the total amount he or she has downloaded.

Definition 2.3 (Passkey). A passkey is a unique identity assigned to each registered user of a Darknet. Usually composed of a 32 bit hexadecimal string, the passkey is appended to the announcing URL in the .torrent file that is dynamically generated for each user. This prevents private torrents from being uploaded to public BTs. Users are prohibited from leaking their passkeys and announcing URLs to other users or public trackers.

Definition 2.4 (User Class). Most Darknets deploy a ranking system to categorize their users based on each user's contribution to the community. For example, users may be ranked on a scale such as newcomers, users, power users, and elite users. Users' contributions normally refer to their download and upload volumes, sharing ratios, and membership time. Different classes of users have different privileges.

Definition 2.5 (Seeding). A seed(er) is a client on the BT network that has a complete copy of the content. Once a user finishes downloading, he or she can choose to stay in the BT swarm to upload data to other users. This is known as being a seed or seeding.

Definition 2.6 (Leeching). A leech(er) is a client on the BT network that does not have a complete copy of the content yet. Whenever a user begins downloading, he or she becomes a leecher until the entire content has been downloaded.

Definition 2.7 (Snatched). This indicates that a torrent has been completely downloaded by a member.

Definition 2.8 (Freeleech). When a torrent is flagged as a freeleech, leeching that torrent will not affect a user's download volume nor decrease his or her sharing ratio. On the other hand, seeding the torrent will continue to increase the user's upload volume. There are different levels of freeleech promotion, e.g., 50%, 30%, or free, which means that only 50%, 30%, or none of the download volume is counted, respectively, and the upload volume is not affected. Freeleech is mainly used by Darknets to help users increase their sharing ratios.

Definition 2.9 (Wait Time). Wait time is the period of time that a user must wait before he or she can start to download. This delay in downloading will only affect users with a low sharing ratio, and is controlled at tracker level.

Definition 2.10 (Credit/Point Mechanism). Many Darknets incorporate a "credit" or "point" mechanism along with the SRE mechanism. Credits or points can be earned by maintaining a good sharing ratio, continuing to upload torrents, and offering donations, and can be exchanged for upload traffic or invitation codes.

Definition 2.11 (Codebase). Instead of developing their own community website, Darknets typically use or modify open source code (or codebase) such as Gazelle and XBT Tracker (XBTT).

Definition 2.12 (Pre-time). Pre-time refers to the time taken by a Darknet to release new and popular torrents. In general, shorter pre-time will result in better reputation of the Darknet [Chen et al. 2010a].

2.2 Operational Principles

Darknets implement strict rules to control user eligibility and content quality. Users gain membership to Darknets through an invitation system. By using a unique passkey, each user is given a unique announcing URL to perform content distribution. Darknets also build their own User Class system based on the SRE mechanism, whereby users can be automatically or manually promoted or demoted. Many Darknets enforce another visiting frequency policy, which requires registered users to log in to their accounts and access the Darknet website at least once within a time period, e.g., 6 weeks. A typical Darknet operational structure is illustrated in Fig. 1.

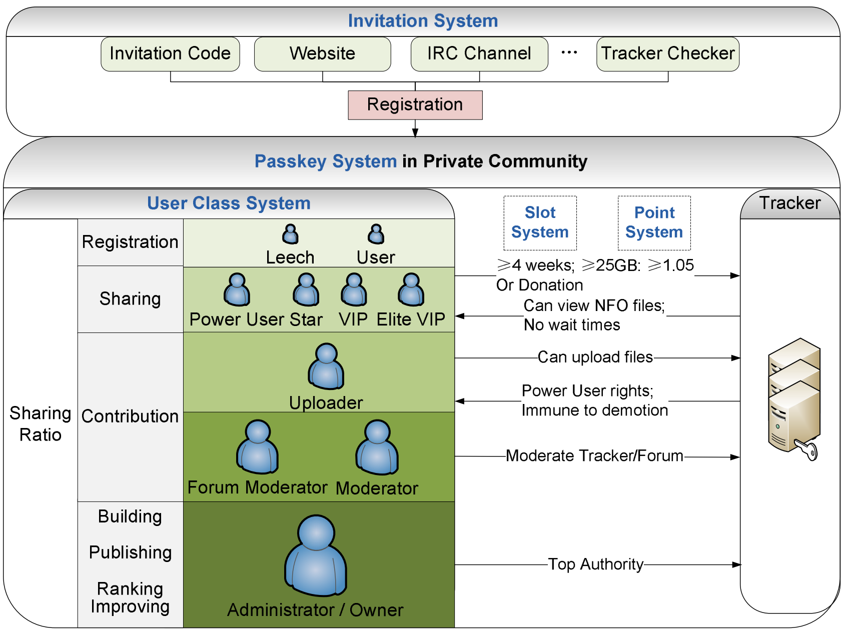

Fig. 1. A Typical Darknet Operational Structure

Remark: NFO file is a special text file that contains release information about the media content. The forum moderators are usually the owners of the forums.

2.3 Incentive Mechanisms

The incentive mechanisms used in Darknets are classified in Table 2. To the best of our knowledge, all existing incentive mechanisms can be classified into one of these categories. The SRE mechanism is the main incentive mechanism. Because the SRE mechanism alone will lead to a new problem that many users cannot achieve the required sharing ratio, two other methods (i.e., credit/point mechanism and ratio free mechanism) are being used to maintain the stability of Darknets.

Table 2. Taxonomy of Darknets based on Incentive Mechanisms

Category	Description	Examples
SRE mechanism	Only uses the sharing ratio. A minimum sharing ratio is required.	RevolutionTT, TorrentLeech, ILoveTorrents, Bitsoup.
SRE mechanism and Credit/Point mechanism	The sharing ratio and credit/point mechanisms are used together. These Darknets normally have minimum sharing ratio and credit/point requirements.	HD-Torrents, HDChina, CHDBits, HDStar.
Ratio Free mechanism	No minimum upload requirements. Normally have minimum requirements for seeding time or credit/point.	RuTracker, FtN, PtN, BTN.

The *Sharing Ratio Enforcement (SRE) mechanism* is the default incentive mechanism in Darknets. As discussed in Section 1, SRE incents users to seed for as long as possible to improve the whole community's download performance. However, as will be shown in Section 4, the supply (i.e., the number of seeders) is far more than the demand (i.e., the number of leechers), and many users have difficulties in maintaining the required sharing ratio, especially for those with limited uploading bandwidth.

A number of Darknets use the *Credit/Point mechanism* alongside the SRE mechanism to resolve the above problem. Although there are various credit systems, their basic principles are the same. In general, users can gain credits by seeding more torrents, increasing their seeding period, seeding old torrents, and seeding large torrents. The credits can be exchanged for additional upload traffic amounts. As a result, the credit/point mechanism can serve to increase users' sharing ratios, and boost their enthusiasm to seed more torrents, and to seed for a longer time, etc.

The *Ratio Free (or No Ratio) mechanism* is another incentive mechanism used in some Darknets. "Ratio Free" Darknets aim to promote higher levels of voluntary seeding and increased numbers of long-term seeds, without enforcing mandatory minimum upload data requirements. Users are not obligated to reach the sharing ratio requirement. Instead, they generally need to seed torrents for between 36 and 72 hours after the download.

3. MEASUREMENT SETUP

3.1 Statistics of BitTorrent Communities

Since Darknets are underground closed communities, there is no accurate estimation on the number of Darknets over the world. The research of [Sharky 2011] has reported the existence of nearly 1,000 Darknets. The Darknets only expose their inside information to their registered users, so the first challenge of carrying out a measurement study is to join different Darknets. According to [Sharky 2011] and [Torrentking.org 2012], there are two general ways to gain membership to Darknets. One is to get an invitation code from a friend who is already a registered user of a Darknet. Notice that the invitation codes are usually offered as rewards based on a user's contribution to the Darknet and the number of invitation codes given to a registered user is strictly limited. The other way of gaining membership is to sign up during the “open registration period”. Some Darknets occasionally open for membership registration on specific holidays, and they only accept a very limited number of new users. We managed to obtain membership to 17 Darknets by invitation code or open registration. After obtaining these precious memberships, we joined the content distribution of the Darknet to increase our sharing ratios above the pre-defined threshold during the specified 1-3 week period, and then continue to remain active in the Darknets. Otherwise, we would have been banned from the Darknets for having a low sharing ratio. Some Darknets provide additional channels that allow users to make donations to exchange for invitation codes or upload traffic to increase the sharing ratios.

We developed a crawler to collect useful information from the websites of Darknets, which are only available to members. Because different Darknets are built from different codebases, the kind of information we can obtain from each Darknet is very different. Commonly available information include the number of torrents, the number of registered users, and the list of torrents. Some Darknets also release torrent-level information to members, such as the number of seeders, the number of leechers, seeder-to-leecher ratio (i.e., the ratio of the number of seeders to the number of leechers in a BT swarm), etc. A few Darknets also release user profile information, including each user’s download amount, upload amount, seeding time, and sharing ratio. We summarize our data sets in Table 3, which gives the overview of 17 Darknets, along with one public BT search engine and 2 public BTs.

Since the crawler is located in Hong Kong, all the time and date information reported in this paper are in UTC+8. The crawled data is recorded periodically by scanning each Darknet’s webpages.

To have a better understanding on the properties of Darknets, we analyze the data at two different levels: community-level and torrent-level, as summarized in Table 4. At the community-level, we divide the data sets into the following: community related, user related, and sharing ratio related. At the torrent-level, we divide the data sets into user related and torrent related.

Table 3. List of BitTorrent Communities and Related Information

Category		Name	# of Torrents	# of Users	Torrent-level Statistics	Torrent List	Trace Duration (mm/dd/yy)
BT Search Engine		ISOHunt	6,851,730	N/A	√	√	12/29/09 – 02/28/11
Public	General	ThePirateBay	3,697,031	4,939,460	√	(Part)	12/29/09 – 02/28/11
		TorrentPortal	2,494,726	1,218,022	√	√	02/02/10 – 02/28/11
Private	General	Demonoid	355,299	N/A	N/A	√	01/17/10 – 02/28/11
		ILoveTorrents	10,486	N/A	N/A	√	09/28/09 – 02/28/11
	Scene Release	RevolutionTT	46,554	N/A	N/A	√	12/26/09 – 02/28/11
		TorrentLeech	41,333	N/A	N/A	√	01/15/10 – 02/28/11
		TorrentVault	40,032	15,000	√	√	03/02/10 – 02/28/11
		DigitalHive	29,282	25,693	√	√	03/02/10 – 02/28/11
		Bitsoup	14,604	N/A	N/A	√	10/16/09 – 02/28/11
	High Definition	HDChina*	39,281	26,078	√	√	02/28/10 – 02/28/11
		CHDBits*	33,450	31,350	√	√	09/28/09 – 02/28/11
		HDStar*	17,561	23,045	√	√	10/05/09 – 02/28/11
		HD-Torrents*	16,660	26,562	√	√	01/23/10 – 02/28/11
	Foreign	RuTracker.org+	816,835	8,445,600	√	N/A	10/01/09 – 02/28/11
	Music	DimeaDozen	40,814	114,998	√	√	02/02/10 – 02/28/11
	TV	TheBox.bz	70,608	N/A	N/A	√	02/02/10 – 02/28/11
	DVD	AsianDVDClub	37,252	35,408	√	√	02/02/10 – 02/28/11
	e-Learning	ElBitz	12,789	N/A	N/A	√	03/16/10 – 02/28/11
	Adult	PureTNA	66,893	631,015	√	√	02/02/10 – 01/01/11

Data updated until 00:00 EST, Feb 28, 2011, except for PureTnA (until Jan 1, 2011)

N/A indicates the community does not provide the corresponding data.
PureTnA was shut down Jan 1, 2011.
DimeaDozen is the only community that provides the statistics of each torrent's transfer speed in the Table.
* indicates the PTs use SRE mechanism with Credit/Point mechanism; + indicates the PT uses Ratio Free mechanism, others use SRE mechanism.

Table 4. Information in the Data Set of Selected BitTorrent Communities

Data Set Category		Included meta-information
Community Level	community related	page view per user, pre-time, seeder-to-leecher ratio, bandwidth
	user related	region, user class, user's upload amount, download amount, seeding time, leeching time, sharing ratio, join date
	sharing ratio related	sharing ratio distribution
Torrent Level	user related	upload amounts, download amounts, average upload speed and downloading speed in a swarm, seeding time and leeching time in a swarm
	torrent related	the number of seeders, the number of leechers, the number of finished downloads, published time duration of a torrent

3.2 Challenges during Crawling

We encountered two major challenges in the course of data collection: the incomplete torrent list of ThePirateBay (public tracker) and information loss due to the crawling time intervals.

ThePirateBay only displays the most recent 100 pages for each torrent category (36 categories in total). Each page lists the records of 30 torrents, and we crawled 108,000 torrent records in total. Although we are unable to gain a complete picture of ThePirateBay, we crawled all of the available torrents. Moreover, the hidden torrents are not recent torrents, and they only can be accessed by the administrator of ThePirateBay.

The crawl results from the two public trackers and the BitTorrent search engine are used to analyze and compare with the Darknets in terms of seeder-to-leecher ratios (ISOHunt, TorrentPortal and ThePirateBay), seeder distribution (TorrentPortal), and active torrent rate (TorrentPortal and ThePirateBay). ThePirateBay's crawled information is therefore only used in two aspects of the analysis. Furthermore, because ISOHunt and TorrentPortal have complete torrent lists, ThePirateBay results are used for reference only. Accordingly, the incomplete list of torrents on ThePirateBay will not affect the final conclusion. On the other hand, we successfully collected complete torrent lists from the selected Darknets.

The frequency of the data collection needs to be carefully tailored to achieve smooth crawling. It is impossible for trackers to display all of the torrent information on one page. The torrent information covers multiple pages (typically more than 200 pages), with each page displaying 30-50 torrent records. Note that almost all trackers deploy a minimum time interval $T_{tracker}$ (generally 30 seconds minimum) between two page visits to prevent malicious access. If our time interval for crawling each webpage (T_{crawl}) was too short (i.e., T_{crawl} is shorter than $T_{tracker}$), then the tracker would temporarily block our IP. Alternatively, if the time interval T_{crawl} was too long (i.e., T_{crawl} is much longer than $T_{tracker}$) then the status of each torrent would change too much, which would interfere with the analysis results based on a snapshot of the complete torrent list.

To resolve this issue, we estimated the time interval, T_{crawl}, of each of the 17 Darknets to prevent information loss related to the time interval $T_{tracker}$, the total number of torrents in a Darknet, the number of torrents displayed on one page, the total pages of a Darknet, and the changing frequency of torrents, etc. We configured our crawler to obtain complete torrent lists every 4 hours and to limit the variation of the status of each torrent, such as the number of seeders and leechers. For sites with too many torrents such as ISOHunt and TorrentPortal, we use multi-threading to crawl torrent related information so that we can crawl the complete torrent list within 4 hours. Therefore, the complete torrent lists reflect the overall populations of torrents in the trackers. The “Torrent-level Statistics” in Table 3, which are shown on the index page of some of the Darknets and are collected every hour (T_{stat}), reflect the macro-scale tracking. However, the time interval T_{stat} should not be too short, because the Darknets do not update these statistics in real-time. According to our investigation, setting T_{stat} to a small value is wasteful as statistics are updated at a lower frequency. In our measurement study, we set T_{stat} to 1 hour as a good tradeoff between data accuracy and crawling traffic. Furthermore, we traced the status (i.e., the numbers of seeders and leechers, and the snatched times at one moment) of three popular torrents every 10 minutes to study the life span of the torrents, which will be further discussed in Section 5.3.

4. COMMUNITY-LEVEL MEASUREMENT AND ANALYSIS

This section provides the results of the community-level measurements of the Darknets, such as user activities, traffic amount and sharing ratios, and compares these results with those of the public BitTorrent communities when possible. As mentioned in Section 3.1, different Darknets provide different information, and hence our discussion in each of the following sub-sections may be based on different Darknets. Specifically, many HD Darknets in China adopt the code base (i.e., NexusPHP powered by p2pnow.net) which can provide abundant statistics for us to analyze, while other Darknets outside China rarely use this code base and provide less statistics data.

4.1 Community Related Data

1) Page Views per User

Using the traffic rankings provided by [Alexa 2012], we compare the page views per user during a period of 6 months (until Jan 1, 2011) for our selected tracker sites. Page views measure the number of pages viewed by each site visitor. Multiple page views of the same page by the same user on the same day are counted only once. Therefore, the page views per user in Fig. 2 are the average number of unique pages viewed per user per day in the period of 6 months. When calculating the average page views per user, only those users who actually visited the website are counted. Besides ISOHunt, thePirateBay, TorrentPortal, we select some other famous public trackers (TorrentZ, Bitsnoop, Mininova, TorrentBox, BTJunkie) from [Torrentking.org 2012] sorted by highest hits, and compare them with our measured Darknets in Fig. 2. In general, the Darknets have higher page views per user than those public BTs, which indicates that Darknets have more active user engagement than public trackers. PureTNA provided adult content and received the highest page views per user among all sites we investigated.

One main reason for this is that the Darknets have a better download performance than public trackers [Zhang et al. 2010]. Only some of the most famous public trackers can compete with Darknets. In addition, the SRE mechanism in combination with the User Class ranking system guarantees user engagement, which helps to further increase user activity. In addition, users prefer to download content as soon as it has been released [Andrade et al. 2005]. This phenomenon reflects the users’ need for esteem and self-actualization, which is in accordance with Maslow’s “hierarchy of needs” theory [Maslow 1954]. Maslow's hierarchy of needs is a theory in psychology proposed by Abraham Maslow. Maslow used the terms Physiological, Safety, Belongingness and Love, Esteem, and Self-Actualization needs to describe the pattern that human motivations generally move through. SRE mechanism guarantees the Safety need for Darknet users to keep staying in communities, and the users can achieve Esteem and Self-Actualization needs in terms of high user class rank and very quick downloading rate. Darknets also provide very fast pre-times, whereas the users of public trackers such as ThePirateBay have to endure the “trickle-down” effect, where they have to wait until a torrent is leaked from one of the major Darknets.

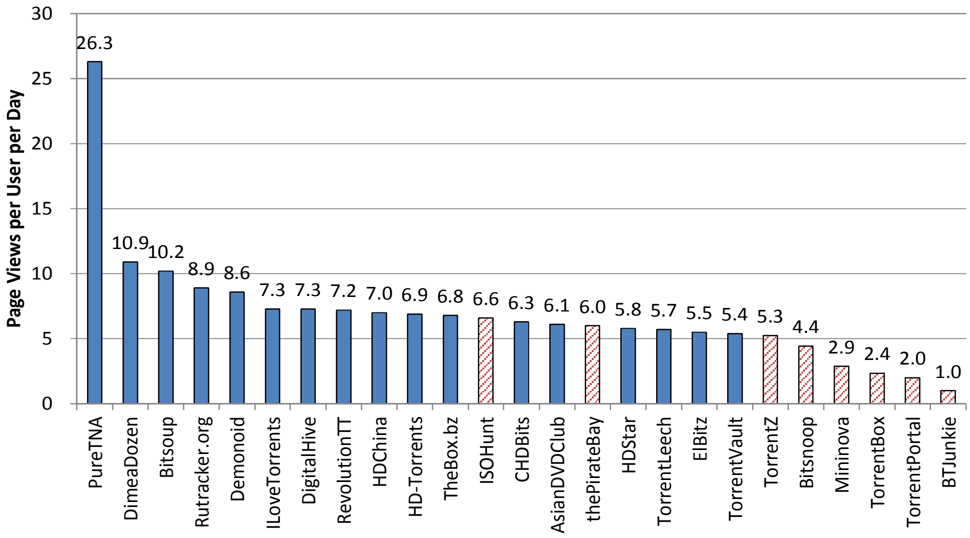

Fig. 2. Page Views per User per Day in Popular Tracker Communities (until Jan. 1, 2011)

2) Seeder-to-Leecher Ratio

Fig. 3 shows the average Seeder-to-Leecher Ratios (SLRs) of five Darknets over a 260-day period. As a comparison, we also show the average SLRs of five public trackers (inside the red box and shown as dashed line). The SLRs of the Darknets (shown as solid lines) are clearly much higher than those of the public trackers. This is indeed a major nice property of Darknets. The high SLRs of Darknets are strong evidence of the effectiveness of the SRE mechanism in incenting Darknet users to serve as seeds as long as possible and as many torrents as possible. On the contrary, the users of public BTs have little incentive to be a seeder after they have finished downloading. A high SLR usually implies high content availability and download performance, and hence makes the Darknets more healthy and attractive.

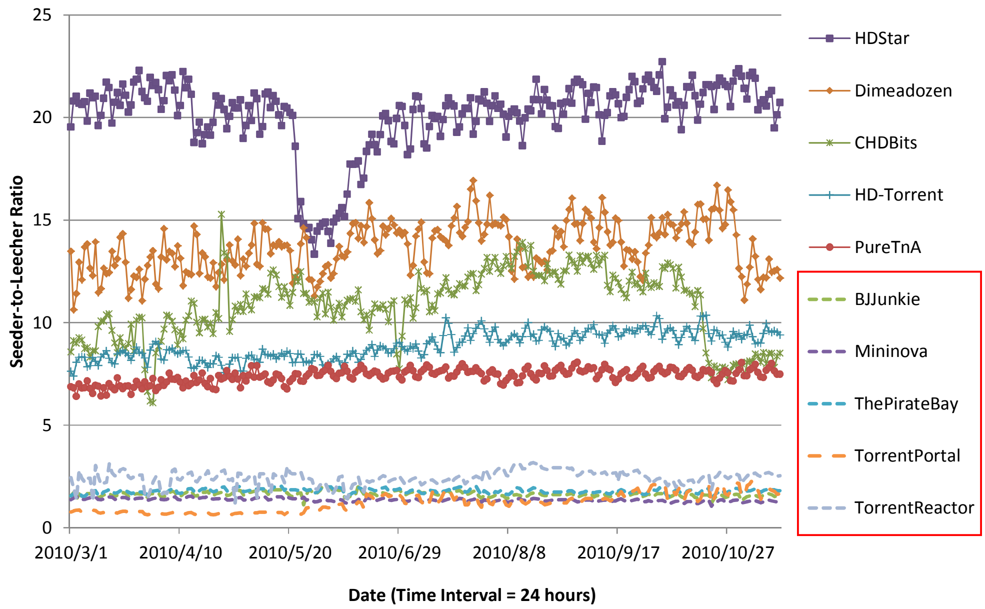

Fig. 3. Seeder-to-Leecher Ratio in Eight Darknets and Five Public Tracker Communities

These data also reveal an interesting phenomenon. Although the free-riding problem exists in the public trackers, the SLRs of the public trackers are still above 1 most of the time (between 0.6 and 3.2). This is because a small fraction of publishers are responsible for 67% of the published contents on public trackers, e.g., the distribution of software updates. As these publishers are driven more by the desire for financial gain other than altruism, they may significantly affect the popularity and content availability of these BT public trackers [Cuevas et al. 2010].

3) Traffic

In this section, we show how much traffic is generated by Darknets. Since the total upload traffic equals the total download traffic within a Darknet, we show the growth of total upload traffic per user in CHDBits and HDStar in Fig. 4, which is calculated as the ratio of the total upload traffic of the Darknet to the total number of users. In the two months period, on average each user uploads 300 to 400GB of data, which is quite enormous. It's also interesting to see that the two Darknets have very similar slopes in Fig. 4. This is because CHDBits and HDStar are located in China with very similar taste of contents (High-definition movies and TV series), and as a matter of fact, they share a lot of common contents. Once a new content is released in one of them, the same content will appear in another one very soon. Since there is a strong correlation between the increase of traffic per user and the increase of new contents, these two Darknets exhibit similar slopes. The sharp increase of the CHDBits curve around 28 February 2010 is because CHDBits removed nearly 8,000 users who couldn't meet the requirement in that day.

Different from the previous two Darknets, RuTracker.org provides the information of aggregated data transfer rate on its website. In theory, the aggregated data transfer rate can be computed by the trackers because they have recorded the detailed information of data upload/download for each user. RuTracker.org is the only Darknet that provides this information to the members. From Fig. 5, we can see the total traffic on RuTracker.org was as high as 15-35 GBps, with nearly 700,000 living torrents. This tremendous volume of traffic may place a heavy burden on the ISPs [David and Fabian, 2008].

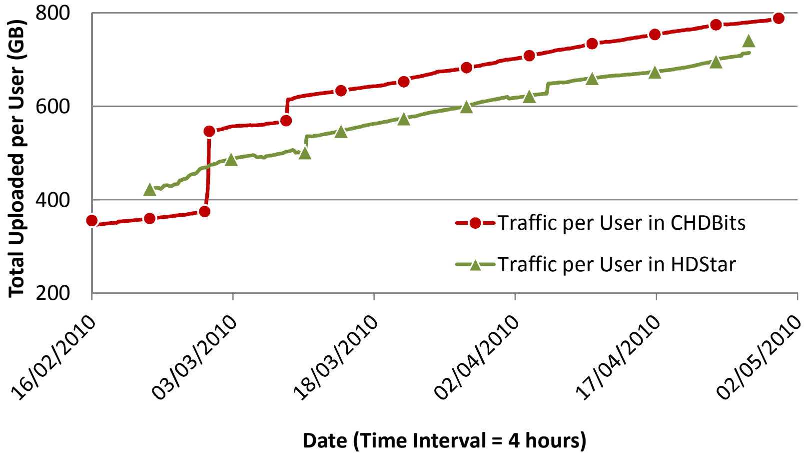

Fig. 4. Growth of total upload traffic per user in CHDBits and HDStar

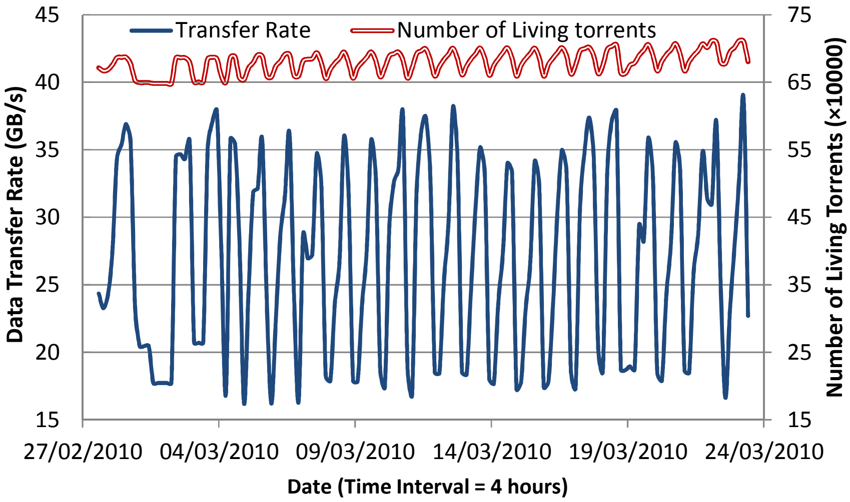

Fig. 5. Data Transfer Rate in RuTracker.org

4.2 User Related Data

1) User Class Distribution

Most Darknets use User Class ranking systems to categorize their users based on their contributions, such as sharing ratios, and upload and download traffic. Table 5 illustrates the User Class system in CHDBits on June 2010. The ranking systems used in the other 16 Darknets follow the same pattern as CHDBits [CHDBits FAQ 2012]. The typical Darknet operational process is as follows. Each Darknet implements a strict set of rules to control member eligibility and content quality. Each user is given a passkey to perform content distribution and their sharing ratio is recorded by the Darknet. The use of a predefined User Class system together with the SRE mechanism enables users to be automatically or manually promoted or demoted.

We closely examine the activity of users by analyzing the number of users in each User Class in TorrentVault, CHDBits, and HDStar. Fig. 6(b) shows the measured data obtained over a 3-month period. The total number of users remains relatively stable, which is typical of many Darknets. This stability is mainly dependent on two factors: the capacity of the Darknet's tracker, and the administrator's maintenance of the Darknet. In Darknets, DHT is disabled and users rely on trackers to find peers. Trackers in Darknets also need to record each user's downloading and uploading behaviors. So the tracker limits the number of users that can be supported by the Darknet.

Table 5. User Classes in CHDBits Darknet on June 21, 2010

<table>
<tr><th>User Class</th><th>Explanation</th><th>Rules</th><th>Number</th></tr>
<tr><td>Peasant</td><td>Demoted users who must improve their ratio within 20 days or they will be banned.</td><td rowspan="2">Demoted to this class if:
Downloaded>10GB & ratio<0.3
Downloaded>50GB & ratio<0.4
Downloaded>100GB & ratio<0.5
Downloaded>200GB & ratio<0.6
Downloaded>400GB & ratio<0.7</td><td>442</td></tr>
<tr><td>User</td><td>It is the default class of new users,
Can upload subtitles and delete subtitles.
Cannot view NFO file.</td><td>20,004</td></tr>
<tr><td>Power User</td><td>Can view NFO file, Can view user list,
Can ask for reseed, and can send invitation,
Can view Top 10 and view other's torrents</td><td>Be a member ≥ 5 weeks;
Downloaded ≥ 50GB & ratio ≥ 1.05;
Auto demoted if ratio<0.95.</td><td>2,343</td></tr>
<tr><td>Elite User</td><td>Elite User or above would never be deleted if parked.</td><td>Be a member ≥ 10 weeks;
Downloaded ≥ 120GB & ratio≥1.55;
Auto demoted if ratio<1.45.</td><td>1,832</td></tr>
<tr><td>Crazy User</td><td>Can view other users' history comments and forum posts.
Can be anonymous when seeding/leeching.</td><td>Be a member≥15 weeks;
Downloaded≥300GB & ratio≥2.05;
Auto demoted if ratio<1.95.</td><td>652</td></tr>
<tr><td>Insane User</td><td>Can upload torrents without permission/review of admins.</td><td>Be a member ≥ 20 weeks;
Downloaded ≥ 500GB & ratio≥2.55;
Auto demoted if ratio<2.45.</td><td>280</td></tr>
<tr><td>Veteran User</td><td>Veteran User or above would never be deleted whether parked or not.</td><td>Be a member ≥ 25 weeks;
Downloaded≥750GB & ratio≥ 3.05;
Auto demoted if ratio<2.95.</td><td>121</td></tr>
<tr><td>Extreme User</td><td rowspan="3">Can update outdated external information.</td><td>Be a member ≥25 weeks;
Downloaded ≥ 1TB & ratio≥ 3.55;
Auto demoted if ratio<3.45.</td><td>60</td></tr>
<tr><td>Ultimate User</td><td>Be a member ≥ 30 weeks;
Downloaded≥1.5TB & ratio≥4.55;
Auto demoted if ratio<3.95.</td><td>26</td></tr>
<tr><td>Nexus Master</td><td>Be a member ≥ 30 weeks;
Downloaded ≥ 3TB & ratio≥4.55;
Auto demoted if ratio < 4.45.</td><td>11</td></tr>
</table>

Different Darknets adopt different User Class policies, although their internal structures remain very similar. Fig. 6(a) shows the member distribution by User Class for three different Darknets. We can see that the majority of members in three Darknets belong to the lowest user class. Note that the Sharing Ratios in lowest user rank is close to but less than 1. Members will not be banned if they maintain Sharing Ratios at this level. This phenomenon indicates that the majority of members have low Sharing Ratios. Meanwhile, a small amount of members (usually those with high upload bandwidth and those who upload original contents) can obtain high Sharing. In a closed Darknet system, the seeming unbalanced member's ranking distribution just demonstrates the relative balance of total traffic between upload and downloading.

We take a closer look at this activity of users by analyzing the number of users in each class of CHDBits. As shown in Fig. 6(b), most members belong to the "User" rank. The number of members in the "User" rank is decreasing with time, while the numbers of users in the "Power User" and "Elite User" ranks are increasing. This shows the effectiveness of the Darknet ranking systems, as members are encouraged to gain promotion to a higher user class. Fig. 6(b) also shows that many members move to a higher level, even though the total number of members remains stable. This drift towards the higher ranks reflects the free-leech promotions or credit/point mechanisms the Darknet administrators use to help users increase their sharing ratios. These methods also help the Darknets to maintain their membership.

(a) Member Distribution by User Class

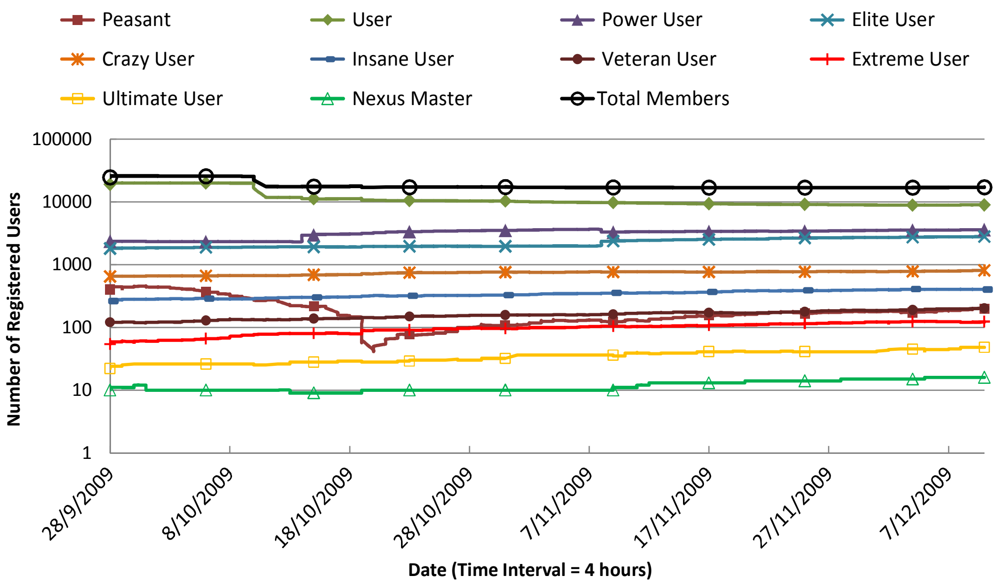

(b) Number of Users of each "User Class" in CHDBits

Fig. 6. "User Class" Distribution

2) User Activity

To be sustainable, the Darknets need to encourage their users to be as active as possible. To this end, the Darknets use the SRE mechanism along with the User Class ranking system to stimulate user activity. The SRE mechanism encourages low bandwidth members to continue uploading to maintain their membership, while the User Class ranking system encourages all members to join the Darknet's content distribution from time to time, which further increases the users' activity.

Among the 17 Darknets we studied, only CHDBits and HDStar provide detailed statistics of users' browsing behavior and tracker activities. Active users can partake in two activities: as active content distributors, they are involved in BT content distribution (uploading/downloading); while as active browsing users, they simply browse the Darknet website and participate in forum discussions. Fig. 7 shows the active user ratio (active users to total users) statistics per day and per week in CHDBits and HDStar. Although the two Darknets in Fig. 7 have different numbers of users, their active user ratios are surprisingly very close, with the per day ratio remaining around 50% and the per week ratio around 80%.

We then select a one-week period (October 13, 2009 to October 20, 2009) to compare the change in the ratios of active browsing users and active users connecting to the tracker (i.e., active tracker users) in CHDBits and HDStar. Fig. 8 shows that the active browsing users in the two trackers present the same time trends. The wave peak occurs at 22:00 each day and the sub-peak is at 10:00 (Time zone

UTC+08:00), which means that most of the browsing activities occur around these two time points. The wave hollow is located at 06:00, which is quite reasonable.

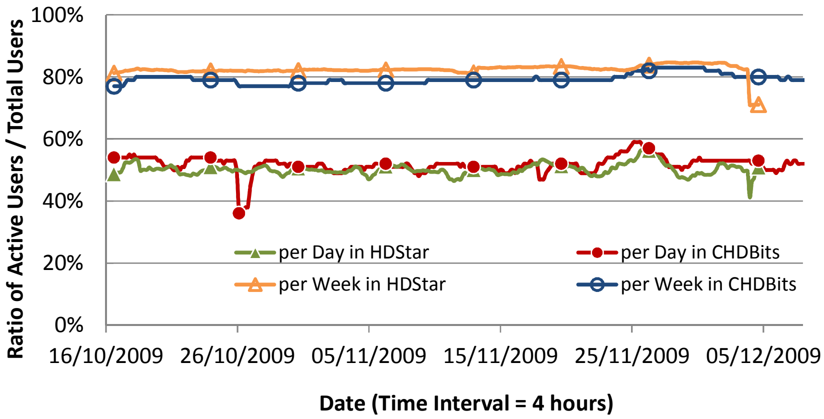

Fig. 7. Ratio of Active Users/Total Users per Day and per Week

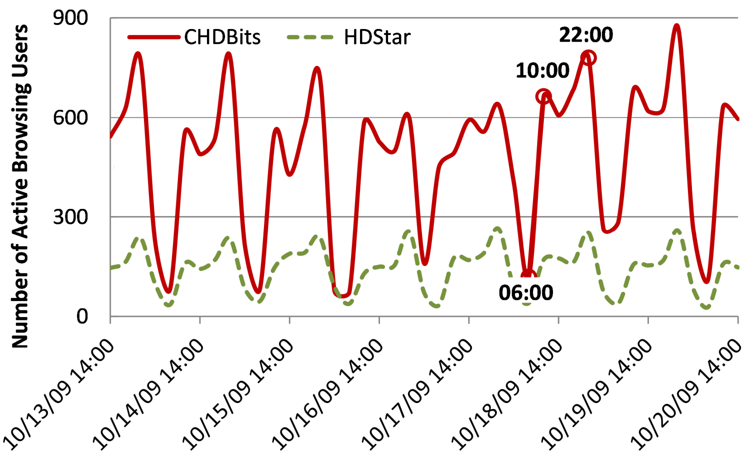

Fig. 8. Numbers of Active Browsing Users

Fig. 9(a) depicts the ratio of active tracker users to active browsing users. It is interesting to notice that the wave peak shifts to 06:00, and the respective wave hollows to 10:00 and 22:00, which is exactly the opposite of the curves in Fig. 8. This is because the number of active tracker users is very stable during the whole day. Although most users are not browsing at 6:00AM, their BT software keeps connected to the tracker. Accordingly, the ratio of active tracker users to active browsing users reaches peak status around this time.

Fig. 9(b) shows the number of active tracker users and the number of active browsing users. We can see the fact that both Darknets show regular patterns everyday. Most Darknet users remain connected to tracker to keep downloading but not browsing the Darknet website.

The high member activity indicates the effectiveness of SRE mechanism and User Class ranking system. The SRE mechanism pushes low bandwidth members to continue uploading to maintain their identities, and the User Class ranking system encourages all of members to join Darknet's content distribution from time to time, which further increase the member activity.

The irregular wave changes (around 22:00, October 17, 2009 to 10:00, October 18, 2009) for CHDBits in Figs. 8 and 9 are a result of the tracker suffering a non-periodical DDoS attack (as announced by the Darknet administrator).

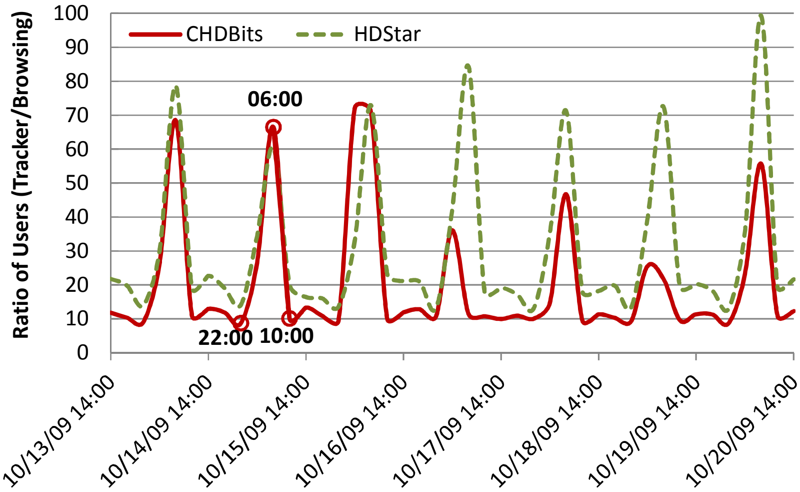

(a) Ratio of Active Tracker Users/Active Browsing Users

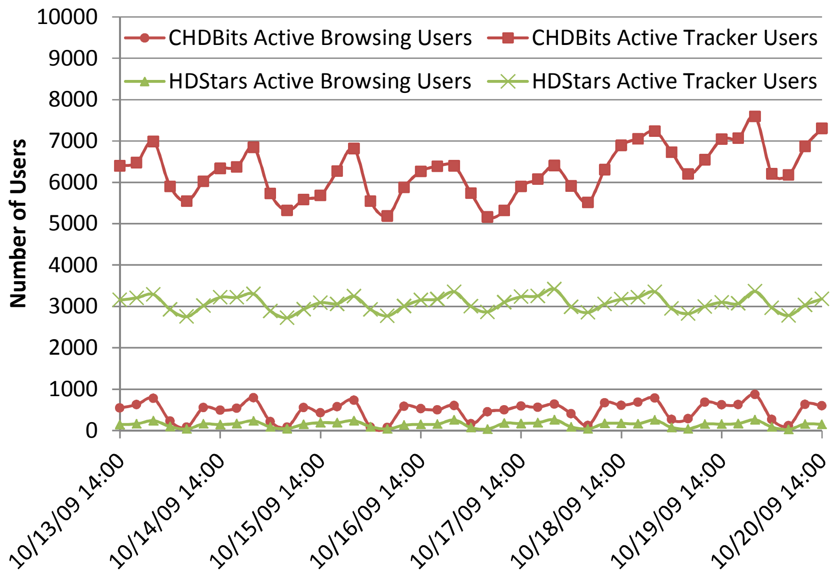

(b) Number of Active Tracker users and Number of Active Browsing Users

Fig. 9. Statistics of active tracker users and active browsing users

3) Membership Life Span

Besides motivating members to be more active, it is also important to retain existing members for as long as possible to maintain the sustainability of Darknets. Here "active" means a user visits a Darknet and joins a BT swarm at least once in the time period, and the user's sharing ratio and visiting frequency must obey Darknet's policy. Fig. 10(a) shows the average membership life span in two Darknets, HDChina and DimeaDozen. About 53% (the black spot in Fig. 10) of members were active during the 334-day measurement period on the two Darknets. Around 80% of the members of HDChina and DimeaDozen were active for 572 days and 1512 days, respectively. Note that HDChina and DimeaDozen had been operating for 754 days and 2265 days, respectively, prior to December 23, 2010. This indicates that the Darknets have very high member engagement. Most Darknets have their own content "taste", so users normally come to register out of interest. Membership of a Darknet provides users with constant access to what they want to download. This feature cannot be guaranteed in public BTs. The life span data shows that the SRE mechanism and the User Class system can enhance member engagement.

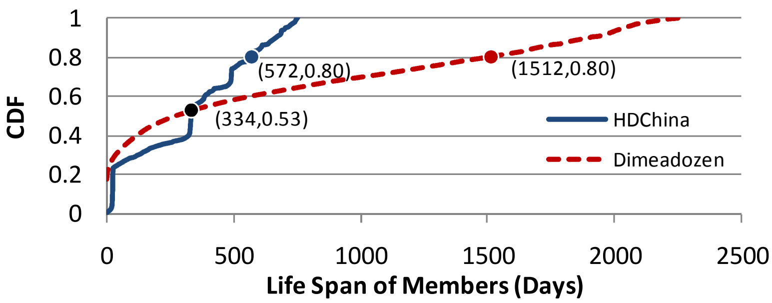

Fig. 10. Life Span of Membership in HDChina & Dimeadozen

4.3 Sharing Ratio Related Data

CHDBits provides the sharing ratio information of all its members. We took a snapshot on all members' sharing ratios and show the CDF in Fig. 11. Around 23% of users have sharing ratios less than 1, and around 19% have sharing ratios higher than 5. Here, we define a "poor" sharing ratio as lower than 1, and a "rich" sharing ratio as higher than 5. Note that the SRE threshold in CHDBits is 0.7.

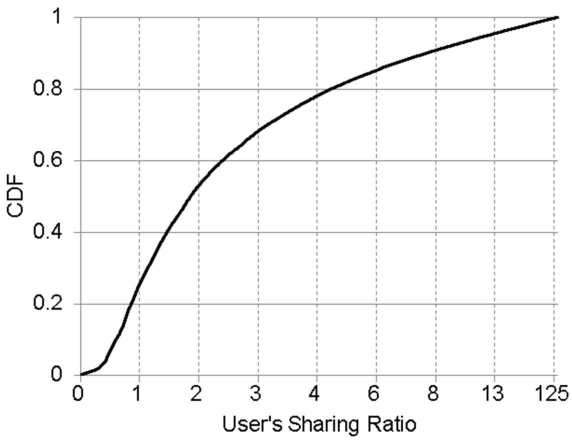

Fig. 11. The CDF of User's Sharing Ratios in CHDBits

Only 4% of members' sharing ratios are lower than the threshold. This is reasonable because the Darknets have very strict membership rules. Those with sharing ratios lower than the pre-defined threshold will first be warned, and then eventually be banned from the community if their ratios don't meet the threshold in time. However, Fig. 11 also indicates that the sharing ratios of the majority of members are far beyond the threshold. Notice that the average sharing ratio can be larger than 1 because of the freeleech promotion scheme. A higher sharing ratio increases the safety of the user's membership.

Users with a high sharing ratio, who upload a lot of data to the community, create an excessive resource supply. This inequity phenomenon (unbalanced supply and demand) is induced by the SRE mechanism. Previous studies have also shown that members tend to seed for long durations, which leads to Darknet's ultra-high downloading performance [Andrade et al. 2005; Andrade et al. 2009]. However, we should take notice of all of the consequences (pros and cons) of these high sharing ratios. When a new member joins a swarm, he or she can no doubt enjoy a sound downloading performance. However, the user may also fear that this will reduce their sharing ratio, especially in light of CHDBits' use of "the last position eliminating principle" (LPEP), which means the person ranked last will be removed from the Darknet. LPEP mechanism sets several minimum thresholds that members must achieve, e.g., uploads of 80 GB, downloads of 60 GB, and a sharing ratio of 0.7. Those who do not meet these requirements will be warned or banned from the community.

Fig. 12 shows the CDFs of users' total seeding time and total leeching time in CHDBits. Notice that the seeding and leeching times are swarm-based, not user-based. For example, if a user is seeding simultaneously in three different swarms for an hour, then his seeding time will be calculated as three hours. From Fig. 12, it is noticeable that the seeding time is much longer than the leeching time; around 50% of members leeched for less than 70 days, and around 50% seeded for more than 1100 days, which means that users seed for around 15 times longer than they leech. This result is consistent with our previous findings [Chen et al. 2010a; Chen et al. 2010b] and the theoretical results of [Jia et al. 2011a, Jia et al. 2011b].

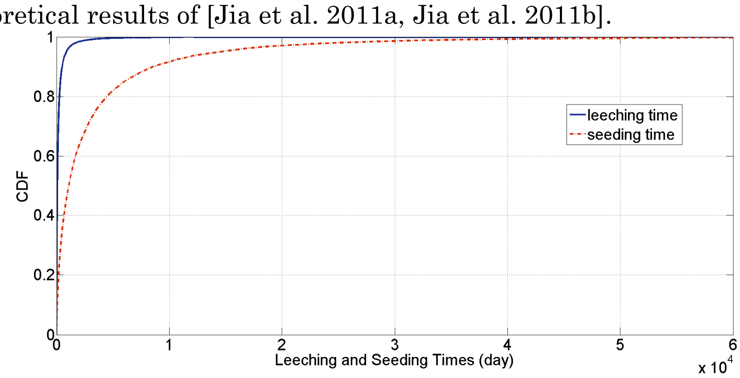

Fig. 12. The CDFs of users' leeching and seeding times in CHDBits

There is an intuition that long seeding durations can bring about high sharing ratios. This is true in a balanced supply and demand environment, because the long seeding durations will result in greater amounts of content being uploaded to the community, and a high sharing ratio. However, in many Darknets supply exceeds demand, which induces very high Seeder-to-Leecher Ratios. Therefore, the real scenario is that users with high sharing ratio increase their sharing ratios by uploading content to new users, but the new users may have difficulty in uploading content to others because they do not have sufficient data blocks for upload. To keep their membership alive, the users with poor sharing ratio have to stay in swarms and hope that the long seeding times can bring about higher ratios.

Fig. 13 shows the upload amount versus download amount of each CHDBits member. It is obvious that each user's upload amount (contribution) increases with the corresponding download amount (consumption), which is similar to the findings of [Andrade et al. 2009]. However, this doesn't necessarily mean that heavy contributions induce long seeding times, nor does it mean that long seeding times lead to heavy contributions.

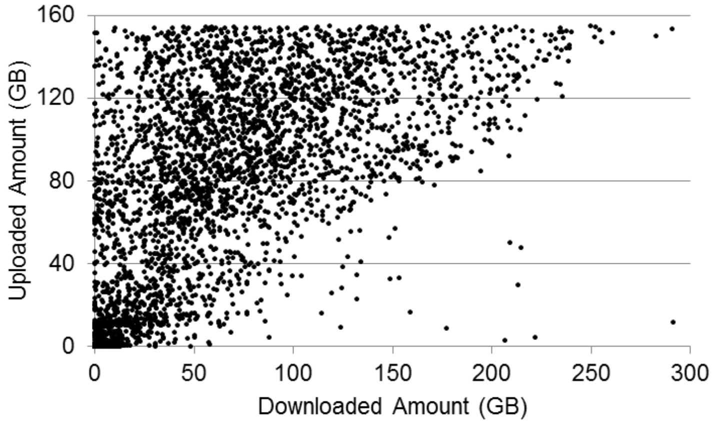

Fig. 13. Upload Amount vs. Download Amount per Member in CHDBits

From the above discussion, we can see a majority of the swarms are heavily oversupplied. In such swarms, seeders are not able to perform any actual uploads due to the insufficient demand. We term this situation *unproductive seeding*. As a consequence, users have to seed for excessively long durations to achieve the sharing ratio required by SRE. As a result, there are currently 6 to 23 times more seeding peers than leeching peers in Darknets, as shown in Fig. 3. To alleviate this imbalance of supply and demand, many Darknets use freeleech promotion strategy and credit systems as additional incentive mechanisms.

5. TORRENT-LEVEL MEASUREMENT AND ANALYSIS

5.1 Active Torrents Ratio

Active torrents refer to those torrents with at least one seed. The active torrents ratio is defined as the ratio of the number of active torrents to the total number of torrents in the community, and is widely used as a key performance indicator to evaluate the activity level and ranking of communities [Zhang et al. 2010, Chen et al. 2012, Jia et al. 2013a, Jia et al. 2013b, Jia et al. 2014]. Fig. 14 compares the active torrents ratio of 16 Darknets and 2 public BTs over 60 days. RuTracker.org is not included because it does not provide a full list of torrents.

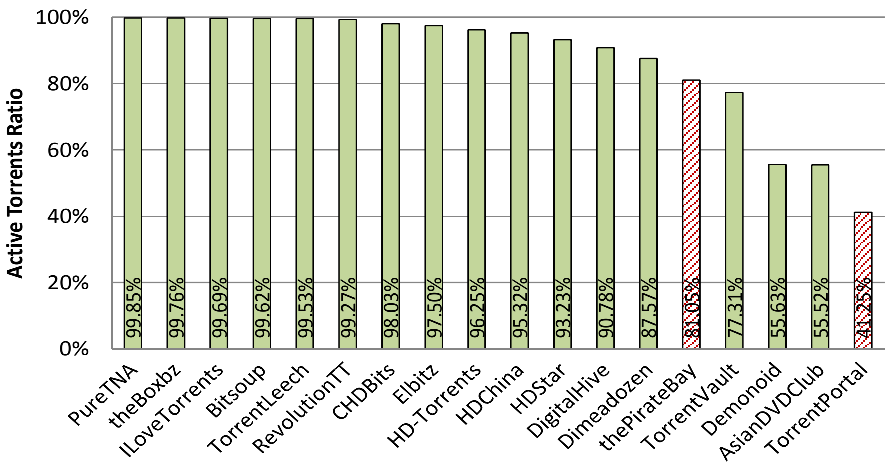

Fig. 14. Active Torrents Ratio in 16 Darknets and 2 public BitTorrent communities

From Fig. 14, we can see most of Darknets have higher active torrent ratio than public BTs. The high availability of contents can attract users to stay longer in Darknets.

The administrators of the Darknets periodically delete unpopular or seedless torrents to allow new torrents to be made available. More importantly, few fake/malicious/wrong torrents exist in Darknets, which is difficult to achieve in public BitTorrent communities. It is almost impossible for public BitTorrent communities to maintain the high levels of content availability and quality found in Darknets. Public BitTorrent communities contain too many torrents for the administrators to maintain and there are no incentive mechanisms for users to upload after they finish downloading. Only a small fraction of the public BitTorrent communities can achieve relatively high content availability.

5.2 Torrent Life Span

We measured all of our Darknets to examine the age of active torrents and their distribution. Fig. 15 (a) shows the torrent age CDF of each single Darknet, while Fig. 15 (b) shows the torrent age CDF of all torrents as a whole. In Fig. 15(a), there are about 53% of torrents that have been alive for more than 1 year. From Fig. 15(b), we can see that about 40% of torrents were created in 96 days and 80% of torrents were created within 388 days. Surprisingly, some torrents that were created 3 years before were still active. The two figures reveal that 40% torrents in Darkents provide fresh contents (no more than 3 months), and half of torrents or more have long life span (more than 1 year). The freshness and long life span of torrents make contents of Darknets updated and available for a long time, which is important to attract users to stay in these communities.

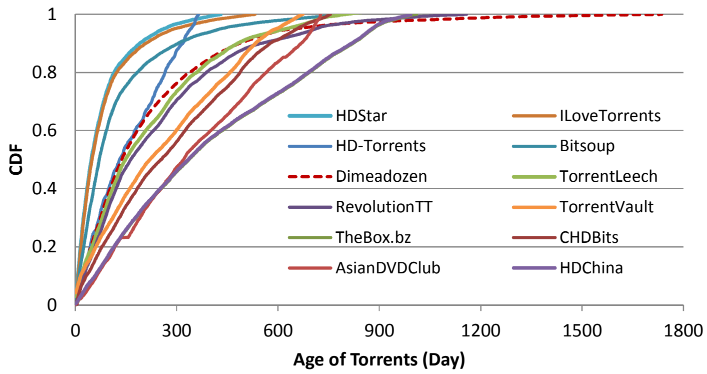

(a) Age distribution in each individual Darknet

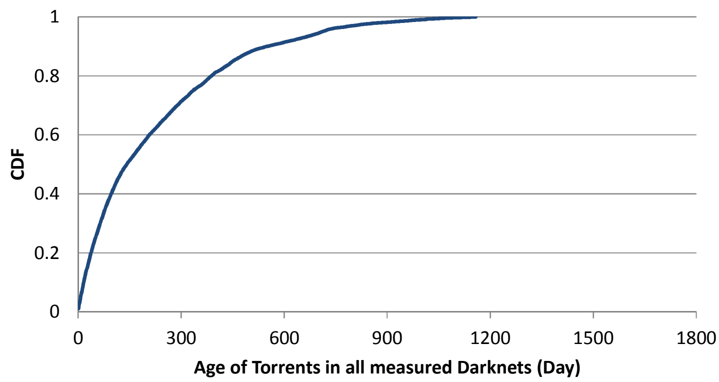

(b) Age distribution of all torrents

Fig. 15. Age Distribution of Torrents in All Measured Darknets

5.3 Torrent Activity

In order to observe the evolution of a torrent swarm at finer time granularity, we traced the status (i.e., the number of seeders, the number of leechers, and the snatched times) of three popular torrents since their birth with a time interval of 10 minutes. Fig. 16 presents the evolution of two torrents of one randomly selected video (one with 720P resolution and another one with 1080P resolution) in CHDBits. The main finding here is that most of the time the number of seeders is significantly greater than the number of leechers. The number of seeders and the number of leechers grow fast at the beginning, and then drop down gracefully. After one to two weeks, the number of leechers drops to a very low level. Fig. 17 shows the results for one torrent in HD-Torrents at a different time, which displays similar features.

From Section 4.2, we know that the majority of Darknet users are at the lowest User Class rank. This could be caused by their low uploading bandwidth, or the short seeding time. The membership of the Darknets surges during the early stages of content distribution, as is the case in public tracker sites. As a large proportion of members subsequently fail to achieve the safe sharing ratio threshold, they continue to stay in the swarm as seeders. Unlike in public tracker sites, where many users leave the swarm as soon as they have finished downloading, Darknet members are eager to upload content to those who join the content swarm later.

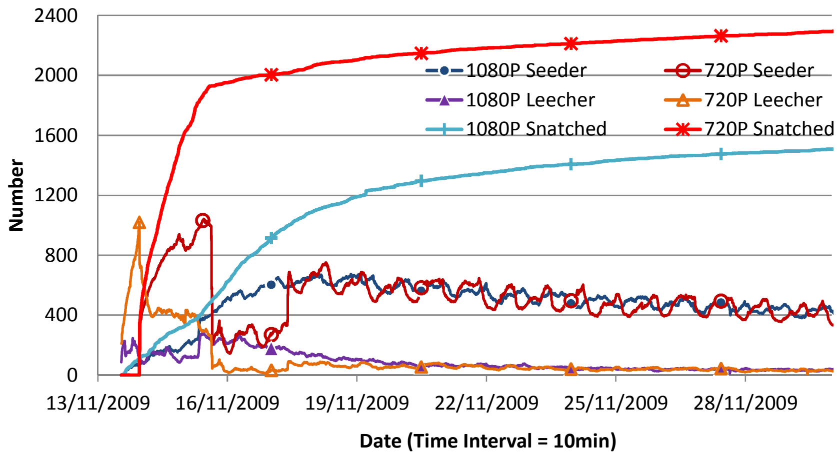

Fig. 16. Number of Seeders, Number of Leechers, and Snatched times of Two Torrents in CHDBits

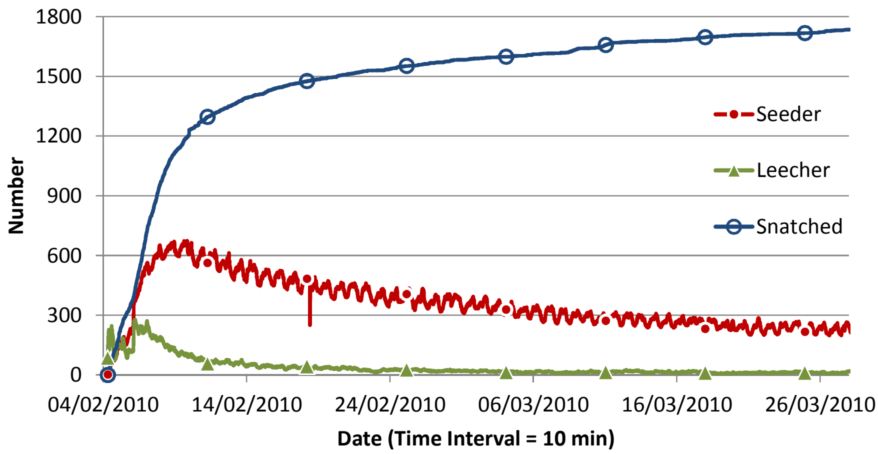

Fig. 17. Number of Seeders, Number of Leechers, and Snatched times of a Torrent in HD-Torrents

5.4 Seeding Time and Upload Amount

Figs. 18(a) and (b) depict the relationship between users' seeding times and upload amounts at the torrent-level for a group of 132 users who have joined a swarm within the first 24 hours since the swarm was born. We took a snapshot on each user's sharing ratio, seeding time duration and upload amount at the end of the 24th hour. Fig. 18(a) plots all users' seeding times and upload amounts, and Fig. 18(b) plots the subgroup of 11 "poor" users whose sharing ratios are between 0.7 and 1.

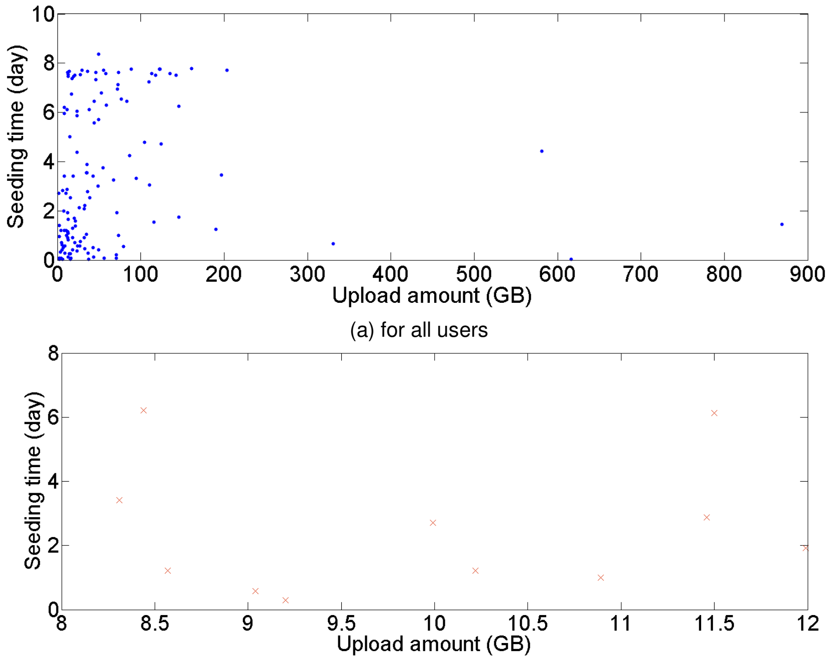

(a) for all users

(b) for users whose sharing ratios are between 0.7 and 1

Fig. 18. Seeding Times vs. Upload Amounts in CHDBits

As shown in Fig. 18(a), the relationship between seeding time and upload amount (contribution) is not obvious. Some users seed for long durations but only have uploaded relatively small amounts of data, while other users seed for relatively short durations but have successfully achieved large upload amounts. The same argument is also applicable to poor users, as shown in Fig. 18(b). These observations are in accordance with the torrent evolution shown in Fig. 16 and Fig. 17. As a matter of fact, users who join the torrent earlier have much higher chance to upload than users who join the torrent later because the number of leechers drops quickly after the peak. Hence it may be more efficient for poor users to closely monitor the release of new torrents and join them as early as possible.

In order to study the long-term impact of seeding time on the upload amount, we took another snapshot at a much later stage when the swarm has 3246 users. We divide these users into four groups based on their upload amount: Group 1's upload amount is less than 300GB; Group 2's upload amount is between 300GB and 1000GB; Group 3's upload amount is between 1000GB and 3000GB; and Group 4's upload amount is more than 3000GB. We show the CDF of seeding time for each group in Fig. 19. It is quite obvious that users with more upload amount are more likely to have longer seeding time. In general, seeding for a long time leads to increase in the sharing ratio.

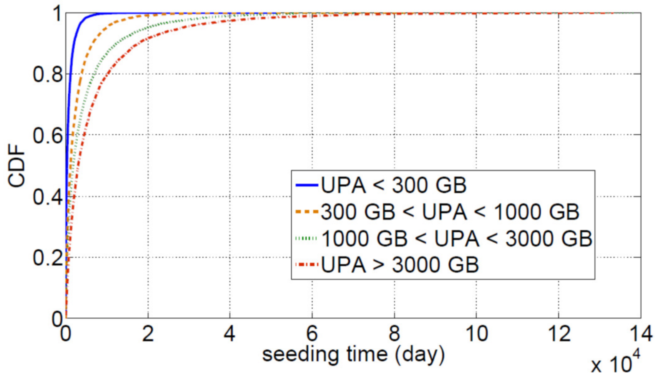

Fig. 19 Distribution of seeding time for different groups of users (UPA: upload amount)

6. RELATED WORK

Most of the existing research on BitTorrent focuses on public trackers. These measurement studies can be classified into two categories. First, some studies use crawling to collect information from BitTorrent systems [Pouwelse et al. 2005; Piatek et al. 2008; Menasche et al. 2009; Siganos et al. 2009; Zhnag et al. 2011]. The crawling techniques exploit the BitTorrent protocol by periodically contacting clients participating in the torrents from the tracker to obtain detailed information. Second, other studies use the log traces from public trackers [Bellissimo et al. 2004; Izal et al. 2004; Guo et al. 2005; Guo et al. 2007], which does not actively interfere with the system. The crawling techniques can be further divided into two categories. [Guo et al. 2007] and [Menasche et al. 2009] found that most torrents are short-lived because of an exponentially decreasing peer arrival rate. The results of these studies draw on the hypothesis that the content is unavailable when there are no seeders present in the swarm. [Kaune et al. 2010] investigated unavailability in BitTorrents and confirmed the conclusion of previous studies that seeders have a significant influence on performance and availability. However, they found that the presence of seeders is not the sole factor determining file availability and that unavailability usually occurs in cyclic periods of intermittent availability. All of these studies explored the ecosystems of BitTorrent public trackers and revealed promising findings.

However, there is limited research on private trackers (Darknets), possibly because they operate underground and are hard to access without membership. Studies examined a typical Darknet (bitsoup.org) in relation to two other BitTorrent tracker sites in terms of resource demand and supply [Andrade et al. 2005; Andrade et al. 2009; Hales et al. 2009]. They show that users typically try to increase their contribution levels by seeding for longer and not by providing more bandwidth to the system. If peers upload more than they download they should gain more credit. However, the Tribler team discovered that Darknets may suffer from a “credit squeeze” phenomenon, whereby a lack of credit significantly reduces the efficiency of the system [Liu et al. 2010]. Moreover, as they focused on a single Darknet HDChina, their findings cannot be generalized to other Darknets. Zhang studied more than 800 Darknets based on geography and content distribution, and further crawled 4 Darknets at medium-scopic level. In addition, they conducted a comprehensive comparison of public tracker sites and Darknets. Their studies provided insights into Darknets and public tracker sites from the macroscopic, mesoscopic, and microscopic perspectives [Zhang et al. 2010; Zhang et al. 2011]. [Meulpolder et al. 2010] conducted extensive measurements of two public trackers and three Darknets, and revealed a number of significant features of Darknets that are different from public trackers.

As compared with these previous studies, we provided a much larger scale of measurement on 17 Darknets for more than 1 year. The measured Darknets cover the world's representative private tracker communities at community-level and torrent-level, which increase the credibility of our analysis results. Existing works focus on exposing the landscape of Darknets, but we demonstrated both the positive and negative effects of SRE. We also discussed the effectiveness of well-adopted community strategies and their effects against strategic user behavior.

7. CONCLUSION

This paper provided a taxonomy of Darknets. We crawled and compared the system behaviors of 17 Darknets, 2 public trackers, and 1 BitTorrent search engine for more than 1 year. We then conducted an in-depth investigation of the Darknets at the community- and torrent-levels. We have shown that Darknet users have higher user engagement than public BTs. The seed-to-leech ratios in Darknets are significantly higher than that in public BTs. We also showed that SRE is an effective mechanism for encouraging users to seed as much as possible. However, SRE alone is not stable because many users have difficulties in maintaining the required sharing ratio. We further discussed the credit/point mechanism and the ratio free system which are used to address the instability of SRE mechanism. In summary, the results of this study extend our previous research and provide a comprehensive picture of the inside workings of Darknets.

ACKNOWLEDGEMENT

We sincerely thank the anonymous reviewers who provided many useful suggestions to improve the quality of this paper. This work is supported by a Hong Kong GRF grant HKBU 210412.

REFERENCES

Alexa. 2012. Alexa: The Web Information Company. http://www.alexa.com/.

N. Andrade, M. Mowbray, A. Lima, G. Wagner, and M. Ripeanu. 2005. Influences on cooperation in BitTorrent communities. In *Proceedings of the Third Workshop on the Economics of Peer-to-peer Systems, ACM SIGCOMM 2005*. Philadelphia, PA.

N. Andrade, E. Santos-Neto, F. Brasileiro, and M. Ripeanu. 2009. Resource demand and supply in BitTorrent content-sharing communities. *Computer Networks*, 53, 4, 515-527.

A. Bellissimo, B.N. Levine, and P. Shenoy. 2004. Exploring the use of BitTorrent as the basis for a large trace repository. *Technical Report*. 04-41, UMASS Amherst.

X. Chen, X.-W. Chu, and J. Liu. 2010a. Unveiling popularity of BitTorrent Darknets. In *Proceedings of IEEE Global Communications Conference (Globecom'10)*. IEEE. Miami, FL.

X. Chen, Y. Jiang, and X.-W. Chu. 2010b. Measurements, analysis, and modeling of private trackers. In *Proceedings of the 10th IEEE International Conference on Peer-to-Peer Computing (P2P'10)*. IEEE, Delft, Netherlands.

X. Chen, X.-W. Chu, and Z. Li. 2011. Improving sustainability of private P2P communities. In *Proceedings of the 20th International Conference on Computer Communication Networks (ICCCN'11)*. Maui, HI.

X. Chen, X.-W. Chu, A.L. Jia, and J.A. Pouwelse. 2012. Inequity of sharing ratio enforcement in Darknet: Measurement and improvement. In *Proceedings of the 14th International Conference on High Performance and Communications (HPCC'12)*. Liverpool, UK.

Chdbits. 2012. FAQ. http://chdbits.org/faq.php.

X.-W. Chu. 2014. Darknet data sets. http://www.comp.hkbu.edu.hk/~chxw/darknet.

R. Cuevas, M. Kryczka, A. Cuevas, S. Kaune, C. Guerrero, and R. Rejaie. 2010. Is content publishing in BitTorrent altruistic or profit-driven? In *Proceedings of the 6th International Conference on emerging Networking Experiments and Technologies (ACM CoNEXT'10)*. Philadelphia, PA.

D.R. Choffnes and F.E. Bustamante. 2008. Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems. In *Proceedings of the ACM SIGCOMM 2008*. ACM, New York, NY, USA, 363-374.

R. Gibbons. 1992. A primer in game theory. *Harvester*. New York, NY.

L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. 2005. Measurements, analysis, and modeling of BitTorrent-like systems. In *Proceedings of the 12th Internet Measurement Conference (IMC'05)*. Berkeley, CA.

L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. 2007. A performance study of BitTorrent-like Peer-to-Peer systems. *IEEE Journal on Selected Areas in Communications (JSAC), Special Issue on Peer-to-Peer Communications and Applications*. 25, 1, 155-169.

D. Hales, R. Rahman, B. Zhang, M. Meulpolder, and J.A. Pouwelse. 2009. BitTorrent or BitCrunch: Evidence of a credit squeeze in BitTorrent? In *Proceedings of the 5th Collaborative Peer-to-Peer Systems Workshop (COPS'09)*. Tarragona, Spain.

M. Izal, G. Urvoy-Keller, E.W. Biersack, P.A. Felber, A. Al Hamra, and L. Garcés-Erice. 2004. Dissecting BitTorrent: Five months in a torrent's lifetime. In *Proceedings of the 5th Passive & Active Measurement Workshop (PAM'04)*. Antibes Juan-les-Pins, France.

A.L. Jia, X. Chen, X.-W. Chu, and J.A. Pouwelse. 2011a. From user experience to strategies: How to survive in a private BitTorrent community. *Technical Report*. http://www.pds.ewi.tudelft.nl/fileadmin/pds/reports/2011/PDS-2011-004.pdf.

A.L. Jia, R. Rahman, T. Vinkó, J.A. Pouwelse, and D.H.J. Epema. 2011b. Fast download but eternal seeding: The reward and punishment of sharing ratio enforcement. In *Proceedings of the IEEE 11th International Conference on Peer-to-Peer Computing (P2P'11)*. IEEE, Kyoto, Japan.

A.L. Jia, R. Rahman, T. Vinkó, J.A. Pouwelse, and D.H.J. Epema. 2013a. Systemic risk and user-level performance in private P2P communities. *IEEE Transactions on Parallel and Distributed Systems*. 24, 12, 2503 - 2512.

A.L. Jia, X. Chen, X.-W. Chu, J.A. Pouwelse, and D.H.J. Epema. 2013b. How to survive and thrive in a private BitTorrent community. In *Proceedings of 14th International Conference on Distributed Computing and Networking (ICDCN'13)*. Mumbai, India.

A.L. Jia, X. Chen, X.-W. Chu, J.A. Pouwelse, and D.H.J. Epema. 2014. User behaviors in private BitTorrent communities. *Computer Networks*. 60, 34-45.

S. Kaune, R.C. Rumín, and G. Tyson. 2010. Unraveling BitTorrent's file unavailability: measurements and analysis. In *Proceedings of 10th IEEE International Conference on Peer-to-Peer Computing (P2P'10)*. IEEE, Delft, Netherlands.

I. Kash, J. Lai, H. Zhang, and A. Zohar. 2011. Economics of BitTorrent communities. In *Proceeding of the 6th Workshop on the Economics of Networks, Systems, and Computation (NetEcon'11)*. San Jose, CA.

Z. Liu, P. Dhungel, D. Wu, C. Zhang, and K.W. Ross. 2010. Understanding and improving incentives in private P2P communities. In *Proceedings of the 30th International Conference on Distributed Computing Systems (ICDCS'10)*. Genoa, Italy.

A. Maslow. 1954. Motivation and Personality (3rd Edition). *Harper and Row Publishers*, 236-237.

D.S. Menasche, A.A. de A Rocha, L. Bin, D. Towsley, and A. Venkataramani. 2009. Content availability and bundling in swarming systems. In *Proceedings of the 5th International Conference on emerging Networking Experiments and Technologies (CoNEXT'09)*. Rome, Italy.

M. Meulpolder, M. Capotă, M. Wojciechowski, J.A. Pouwelse, D.H.J. Epema, and H.J. Sips. 2010. Public and private BitTorrent communities: A measurement study. In *Proceedings of the 9th International Workshop on Peer-to-Peer Systems (IPTPS'10)*. San Jose, CA.

M. Piatek, T. Isdal, A. Krishnamurthy, and T. Anderson. 2008. One hop reputations for peer to peer file sharing workloads. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI'08)*. San Francisco, CA.

J.A. Pouwelse , P. Garbacki , D.H.J. Epema, and H.J. Sips. 2005. The BitTorrent P2P file-sharing system: Measurements and analysis. In *Proceedings of the 4th International Workshop on Peer-to-Peer Systems (IPTPS'05)*. Ithaca, NY.

Sharky. 2009a. Private Trackers with the most torrents. http://filesharefreak.com/2009/09/19/private-trackers-with-the-most-torrents/.

Sharky. 2009b. The Essential guide to getting into Private Trackers. http://filesharefreak.com/2009/06/18/the-essential-guide-to-getting-into-private-trackers/.

Sharky. 2011. Trackers list. http://filesharefreak.com/trackers-list/.

G. Siganos, J.M. Pujol, and P. Rodriguez. 2009. Monitoring the BitTorrent monitors: A bird's eye view. In *Proceedings of the 10th Passive and Active Measurement Conference (PAM'09)*. Seoul, South Korea.

Torrentking.org. 2012. 1196 Torrent tracker list. http://www.torrentking.org/

C. Zhang, P. Dhungel, D. Wu, Z. Liu, and K.W. Ross. 2010. BitTorrent Darknets. In *Proceedings of IEEE INFOCOM 2010*. IEEE. San Diego, CA.

C. Zhang, P. Dhungel, D. Wu, Z. Liu, and K.W. Ross. 2011. Unraveling the BitTorrent ecosystem. *IEEE Transactions on Parallel and Distributed Systems*, 22, 7, 1164-1177.