A Measurement Study on the Topologies of BitTorrent Networks

Majing Su, Hongli Zhang, Xiaojiang Du, Binxing Fang, and Mohsen Guizani, *Fellow, IEEE*

***Abstract*—BitTorrent (BT) is a widely-used peer-to-peer (P2P) application. Most of BT's characteristics (except the topology) have been studied extensively by measurement approaches. In this paper, we deploy a measurement system to examine some performance-related topology properties of BT. Our goal is to provide a measurement view of the real-world BT topologies and to verify the previous estimations via simulations and real-world experiments. We observe that at the steady stage, a BT topology has short distances and low clustering coefficients, and its degree-frequency exhibits a Gaussian-like distribution. These indicate that a BT network is very close to a random network rather than a scale-free network or a small world. The proportion of peers with large download percentages is very high at the steady stage, showing that the swarm is robust from the resource perspective. We also find out that most high-degree peers have a very fast download speed. However, the low Spearman's rank correlation coefficient indicates that there is no strong correlation between the peer connection degree and the download speed. Different from previous results, we find that the diameter of a BT network at the initial stage is small even when 95% of peers use the peer exchange extension.**

***Index Terms*—BitTorrent, topology, measurement, performance, peer exchange.**

I. INTRODUCTION

BITTORRENT (BT) is one of the most popular peer-to-peer (P2P) applications used to distribute large files to massive number of users efficiently. Researchers have been working to improve its performance by measurement, and many measurement studies have been performed on the swarm evolution [1], popularity [2], availability [3], piece population [4], attacks [5], etc. However, little is known about the topology characteristics of the real-world BT overlay, which is important to better understand the BT network, design a BT simulator and analyze the BT performance. A number of important BT properties have not been examined in real-world BT networks. For example, the resilience of a peer departure suggests that the BT overlay may be scale-free (a network whose degree distribution follows a power law distribution [33]), but simulations [7] and experiments [8] do not confirm this. The multiple peer-discovering mechanisms (Tracker [29], DHT [30] and PEX [12]) may lead to clustering or even a small-world of the BT swarms.

Previous works have studied BT topologies by theoretical analyses [6], simulations [7] and controlled-experiments [8]. Most of them provide detailed analysis on the BT topology properties. However, simulation (or controlled-experiment) results may be different from real world measurement results, due to the limitations of simulations (or controlled-environments). This is the main reason that we want to conduct measurements in real networks. To the best of our knowledge, the only work of measuring the topology of a real BT network is performed by Fauzie *et al.* [10], which analyzed the CDF of peer degree and clustering coefficients. In this paper, we analyze the BT swarm topology and its performance-related properties (e.g., peer distance, peer download speed) in different swarm scales and evolution stages by real-world measurement results. We also try to explain some measurement results and their effects on BT performance from the perspectives of protocol design and client implementation. Our goal is to provide a measurement view of the existing BT topology and verify the results (theoretical and simulation) in the literature [6], [7] and [8]. The measurement results would be helpful for designing better BT protocols and client software. Our main contributions are summarized as follows.

- We observe that in the steady stage, a BT swarm has a Gaussian-like peer-degree-frequency distribution, short peer distance and low clustering coefficient, especially for a large-scale swarm. This indicates that a BT network is more like a random network other than a scale-free network or a small world, which is different from the simulation results of Al-Hamra *et al.* [7]. The small amount of low-degree peers indicates that a BT network is more robust to the departure of large-degree peers than a scale-free network.
- Different from previous studies, we find that the diameter of a BT network in the initial stage is still very small (< 6), even if more than 95% of peers use the peer exchange extension. In the steady stage, the shortest path lengths between peers are small (but larger than those at the initial stage); the clustering coefficient of the network is very low (lower than initial stages). This shows that BT networks tend to be less clustered.
- There is no clear evidence that the number of connections of a peer is directly related to its download stages. Compared to the initial stage of a swarm, the proportion of peers with a large download percentage is higher in the steady stage, suggesting that a swarm is robust to peer

Manuscript received February 15, 2012; revised July 16, 2012.

M. Su, H. Zhang, and B. Fang are with the Harbin Institute of Technology, Harbin 150001, China (e-mail: {sumajing, bxfang}@pact518.hit.edu.cn; zhanghongli@hit.edu.cn).

X. Du is with the Department of Computer and Information Sciences, Temple University, Philadelphia, PA, 19122, USA (e-mail:dux@temple.edu).

M. Guizani is with Qatar University, Doha, Qatar (e-mail: mguizani@ieee.org).

This work was supported by the National High-Tech Development 863 Program of China under Grant Nos. 2010AA012504, 2012AA012506, and 2011AA010705; the National Natural Science Foundation of China under Grant No. 61173145; and by the U.S. National Science Foundation under Grants CNS-0963578, CNS-1022552, and CNS-1065444.

Digital Object Identifier 10.1109/JSAC.2013.SUP.0513030

departures from the resource perspective.

- Consistent with the common sense, peers with large degrees have high download speeds, especially in the initial swarm evolution stages. However, there are some low-degree peers whose download speed is close to or even higher than some large-degree peers. We compute the Spearman's rank correlation coefficient between the peer degree and the download speed. The low value of the coefficient (0.0174 $\sim$ 0.4754) suggests that there is no strong correlation between the two.

The rest of this paper is organized as follows. We introduce some related work in Section II and present our measurement methodology in Section III. In Section IV, we analyze the measurement results and study properties of the BT topology. We conclude this paper in Section V.

II. Related Work

There are many studies on characterizing the overlay topologies of P2P applications [19] -[21]. However, research work on BT topologies mainly focused on the locality distribution [13]-[16], [23]-[25]. Only a small number of papers studied the topology properties of BT networks extensively [6]-[8], [10]. Out of these, only [10] used measurement approaches.

Al-Hamra *et al.* [7] studied how the overlay properties impact BT efficiency by extensive simulations. They showed that the BT overlay is robust but not a random graph. They also found that the BT peer exchange (PEX) extension [12] generates a chain-like overlay with a large diameter in the initial swarm stage. On the other hand, Wu *et al.* [12] did not observe the chain-like topology in their PlanetLab [9] experiments. Dale *et al.* [8] studied the topology characteristics and swarm evolution by constructing experimental swarms with more than 400 nodes in PlanetLab. They deployed a modified BT client with a typical configuration to log the connection relationship. Their results showed that the network of peers unchoked by others is scale-free. They found out that there are no clear evidences of persistent clustering and presences of small word. However, in the Internet, due to Network Address Translation (NAT) and firewalls, the connectivity between peers is much worse than that of PlanetLab nodes. This may cause a different peer degree distribution in real world. Besides, the swarm size and diversity of BT clients may also affect peer connections. Zhong *et al.* [6] analyzed topological properties through a directed weighted graph model. They found that the node strength of a BT network follows a power-law distribution. When the proportion of seeders increases, a BT network tends to be less clustered. They also showed that the average shortest path length grows as the network size expands.

Fauzie's work [10] on BT topology properties is the closest one (of the recent research findings) to ours. Fauzie *et al.* [10] collected the information of peers and their connections by a modified BT client (rasterbar libtorrent [17]). They joined the top 35 TV series torrents from the PirateBay [31] and took snapshots with 3-minute durations for about 8 days. Fauzie *et al.* found that a power-law with exponential cutoff is a more plausible model and the clustering is not obvious. However, they did not describe their experimental deployment and sampling in details. We argue that the 3-minute duration

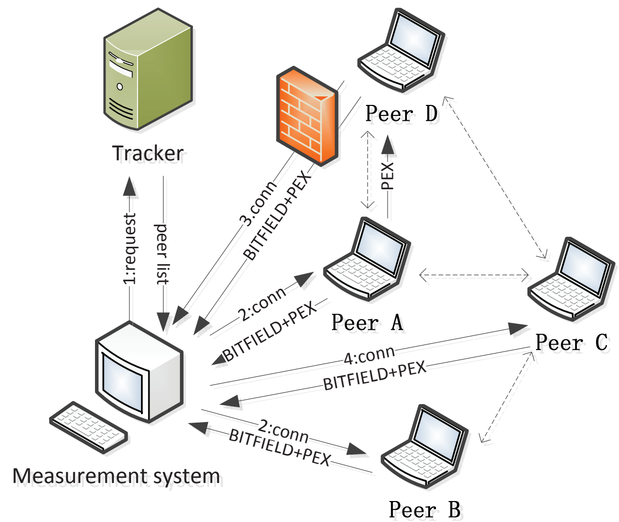

Fig. 1. Illustration of our measurement approach.

may not be sufficient to find all peers of hot torrents due to the discovering ability of BT clients. More importantly, Fauzie *et al.* did not provide sufficient explanations about the underlying reasons of their observations and the potential effects on BT networks.

Different from existing research work [6]-[8], in this paper, we study the topology characteristics of real BT networks by using measurements on real networks. Different from [10], which uses existing BT client software, we specifically design a measurement system to get comprehensive temporal topologies of BT networks in a short time. Furthermore, we provide detailed analyses on BT topology characteristics (e.g., peer degree distribution and clustering). In addition, we analyze the peer download speed and download percentage at different swarm stages, and then explain our results. To sum up, our analyses cover more aspects and are more detailed than [10].

III. BT Topology Measurement

We design a measurement system to capture snapshots of BT swarms. A *snapshot* is a temporal topology of a BT swarm observed at a particular moment, and it includes peers, the status of each peer (e.g., its download percentage and download speed), and connections among peers. In this section, we present the measurement system and our measurement results.

A. Design of Measurement System

A BT does not have a global topology discovering mechanism. We utilize the BT PEX extension [12], which allows peers to exchange the information of newly connected and disconnected peers periodically, to learn the connectivity among peers. Fig. 1 illustrates our measurement system.

We collect peer information from trackers and connection information from PEX messages, and then we construct BT network topologies with this information. We take snapshots in a different way as in [10]. In [10], a snapshot is taken every 3 minutes. However, it would take a longer time to capture snapshots of a hot torrent while shorter time for a cold torrent, and it would take more time at an initial stage than at steady stage. Therefore, in our approach, a snapshot is taken when the peer coverage is above 95%. To quickly capture BT snapshots, we use both active and passive approaches. The dynamic and large scales of BT networks are major challenges.

The active measurement of each snapshot consists of a *peer finding* stage and a *peer detecting* stage. In the peer finding stage, the measurement system requests peer lists from trackers extracted from the torrents, and then connects with these peers to find more peers from PEX messages. Repeat this process until 95% of the peers have been found. In the peer detecting stage, the system connects to all the peers (found in the first stage), exchanges BT messages, and records the BITFIELD (from BITFIELD messages) and neighbors (from PEX messages) of each connected peer. To save bandwidth, the system does not download any file data from peers.

We apply a passive approach to obtain information of a large portion of unreachable peers (behind firewalls or NATs). The system listens and responds to requests from these peers and logs their BITFIELD and neighbors. The system keeps connections with all connected peers to track their status. In our post-processing, the logs are split according to the start time and end time of each snapshot.

It may take a long time to capture a snapshot of a large swarm. To reduce the measurement time, we deploy several detection nodes and send requests to multiple trackers and peers at the same time. We use 10 parallel nodes to reduce the impact on BT swarms in deployment.

B. Tradeoff between Coverage Ratio and Measurement Time

The coverage ratio is important for evaluating the BT topology results. Due to the large scale and dynamics of BT swarms, it is nearly impossible to capture all the peers of a swarm in a short time. Hence, we have to make a tradeoff between coverage and measurement time.

1) Peer Coverage Ratio

The peer coverage ratio is the percentage of peers that a measurement system can detect, which represents the discovering ability of the system. We evaluate the peer coverage ratio in a similar way as in [11]. Assume a BT swarm has N peers (identified by the IP address and the listening port), and $T(m)$ is the total number of distinct peers discovered after m-th requests. The coverage ratio $P(m)$ is defined in (1):

$$P(m) = \frac{T(m)}{N} \tag{1}$$

According to [16], each tracker returns a subset of k peers randomly chosen from N peers. Therefore, at the $(m+1)$-th request, the number of newly discovered peers $\Delta(m+1)$ is

$$\begin{aligned} \Delta(m+1) &= T(m+1) - T(m) \\ &= \frac{N - T(m)}{N} \times k \\ &= [1 - P(m)] \times k \end{aligned} \tag{2}$$

$$\begin{aligned} P(m+1) &= \frac{P(m) \times N + \Delta(m+1)}{N} \\ &= (1 + \frac{\Delta(m+1)}{T(m)}) \times P(m) > P(m) \end{aligned} \tag{3}$$

We use equation (2) to calculate the stop condition of the measurement. If we expect a coverage of P_E, the system keeps detecting until $\Delta(m)$ is less than $\tau = (1 - P_E) \times k$ for C consecutive times, where the threshold C is used to avoid the error of a single measurement. In the implementation of our system, we set $P_E = 95\%$, $\tau = 3$ ($k = 60$) and $C = 20$. According to equation (3), the actual coverage ratio is more than P_E. We use 10 nodes to crawl peers from each tracker in parallel. Based on the results in [11], for a swarm of $N = 5000$ peers, we need at most $74 + 20 = 94$ times of requests, which is about 10 requests from each node. A request is sent to the trackers every 5s. Hence, it takes less than a minute $(10 \times 5s = 50s)$ to crawl a 5000-peer swarm.

2) Connection Coverage Ratio

The connection coverage ratio is the percentage of *definite connection relationship* (connected or unconnected) between peers that a system can observe. Let $Q(r)$ be the total number of definite connections discovered at snapshot r, the connection coverage ratio $C(r)$ is defined in (4):

$$C(r) = \frac{Q(r) \times 2}{N \times (N-1)} \tag{4}$$

There are three types of peers in a snapshot: those directly connected to us (referred to as $\boldsymbol{P_c}$), those not connected to us, but known to us because they are directly connected to peers in $\boldsymbol{P_c}$ (referred to as $\boldsymbol{P_k}$), and those we observed from trackers but not connected to peers in $\boldsymbol{P_c}$ (referred to as $\boldsymbol{P_n}$). We discover that most of popular BT clients (such as μTorrent [26], Vuze [27] and BitTorrent [28]) return all their connected peers in the first PEX message. By exchanging one PEX message, we can get all of the peer connections in $\boldsymbol{P_c}$, and we can obtain the connections between $\boldsymbol{P_c}$ and $\boldsymbol{P_k}$. It takes less than 1 minute to exchange one PEX message. Consequently, the connection coverage ratio can be calculated by (5), where $N = |\boldsymbol{P_c}| + |\boldsymbol{P_k}| + |\boldsymbol{P_n}|$.

$$\begin{aligned} C(r) &= \frac{Q(r) \times 2}{N \times (N-1)} \\ &= 2 \times \frac{\frac{|\boldsymbol{P_c}| \times (|\boldsymbol{P_c}|-1)}{2} + |\boldsymbol{P_c}| \times (|\boldsymbol{P_k}| + |\boldsymbol{P_n}|)}{N \times (N-1)} \\ &= \frac{|\boldsymbol{P_c}| \times (|\boldsymbol{P_c}| - 1) + |\boldsymbol{P_c}| \times (|\boldsymbol{P_k}| + |\boldsymbol{P_n}|) \times 2}{N \times (N-1)} \end{aligned} \tag{5}$$

The peer connections may change during the measurement. Therefore, we keep the connections with peers during the detecting stages and every minute peers send PEX messages containing the newly connected and disconnected peers. The connection set of a peer at the end of a snapshot is its initial connections plus its newly connected peers then minus its disconnected peers. In the implementation of our measurement system, the peer detecting stage lasts less than 130s.

C. Deployment and Measurement Results

We conducted the measurement from 13 December 2011 to 31 January 2012. We randomly chose 106 torrents whose population is larger than 300 peers from a well-known BT torrent search engine TorrentZ [18], and deployed our system on 2 servers in Harbin, China, which connect to CERNET [32]. Each server has 10 concurrent measurement nodes. The

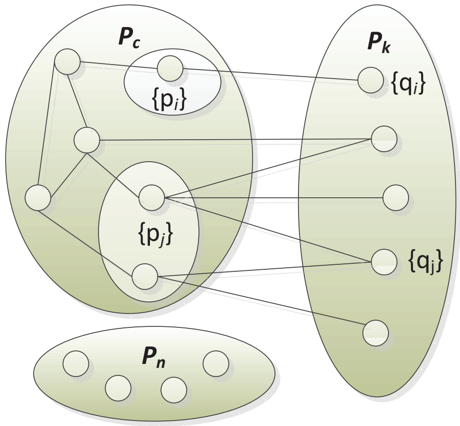

Fig. 2. Illustration of the connectivity graph.

system continuously captured the snapshots of each swarm, and each snapshot took about 3 minutes on average. We started to measure about 30 torrents only a few hours after they were published to the website, and most of their swarm evolutions were caught.

In each snapshot, the percentage of peers that our system could connect to (in $\boldsymbol{P_c}$) is about 30% $\sim$ 50%, and the percentage of peers that we observed from PEX message (in $\boldsymbol{P_c} \cup \boldsymbol{P_k}$) is about 85% $\sim$ 90%. There are about 10% $\sim$ 15% peers that we can observe from trackers but we could not connect them or observe them from PEX messages (in $\boldsymbol{P_n}$). According to equation (5), the connection coverage ratio is about 51% $\sim$ 75%.

D. Factors Influencing the Coverage

There are some factors that affect the coverage ratio of the measurement results. First, we cannot get the peer lists of the seeders behind firewalls and NATs. Second, some BT clients (e.g., μTorrent) do not send PEX messages to peers if they connect less than 3 peers. Third, we cannot establish connections with peers whose number of connections has reached the limit. The status of these peers and their connections cannot be discovered and therefore are missing from the results.

IV. Topology Analysis

In this section, we analyze the properties of BT topologies by using the measurement results. We select 5 torrents of different scales to demonstrate our analysis results, as shown in Table I. Other torrents have similar characteristics.

The topology of each snapshot can be represented by an undirected graph $\boldsymbol{G} = (\boldsymbol{V}, \boldsymbol{E})$, in which vertices $\boldsymbol{V} = \boldsymbol{P_c} \cup \boldsymbol{P_k} \cup \boldsymbol{P_n}$ denote the peers and edges $\boldsymbol{E}$ denote the connections, i.e., $(u, v) \in \boldsymbol{E}$ if and only if peers u and v are connected to each other. Fig. 2 gives an illustration of such graph.

A. Peer Degree Distribution

The degree of a peer v is the number of peers to which it connects in each snapshot, denoted as $deg(v)$. The distribution of peer degree reflects the robustness of a BT network. A previous experimental study by Dale *et al.* [8] showed that the BT peer degree does not follow a power-law distribution. To examine the peer degree distribution, we study the rank-degree distribution and degree-frequency distribution based on our measurement results.

1) Rank-degree Distribution

Fig. 3 plots the steady-stage snapshots (the 4th day after the measurement starts) of the rank-degree distributions of reachable peers. The x-axis is the peer rank in descending order of degrees. Contrary to previous expectations, it can be seen that there are too few low-degree nodes to form a pure power-law network. In each snapshot, only a few peers have very large (or very small) number of connections; most peers have a degree close to a threshold γ (≤ 50). Even if some high-degree peers quit the system, a peer can still download files from other connected peers, suggesting the robustness of the BT swarms. We fit the measurement results with equation (6) and (7), and use the correlation coefficient R to evaluate the fitting results.

$$f(x) = k(C - x)^{\alpha} \tag{6}$$

$$f(x) = kx^{\alpha} \tag{7}$$

Fig. 3 shows that the rank-degree distributions have a two-segment power-law distribution, especially for large scale swarms. If peer degrees are descending ranked, the peer degrees larger than γ follow a power-law distribution. On the other hand, if peer degrees are ascending ranked, peer degrees less than γ follow a power-law distribution. This result is similar to the ultra-leaf degree distribution in the modern Gnutella networks [19].

Potential reasons of the two-segment power-law distribution may include: 1) Due to the limitations of BT clients and network connectivity, a peer cannot establish connections with all peers it knows. Hence, the degree cannot be too large. 2) Trackers, DHT and PEX increase the opportunity of peers finding each other and BT clients may try to setup more connections to increase the download speed. 3) Some clients (e.g., μTorrent) do not return PEX if the number of connected peers is less than 3. Hence, the actual number of low-degree peers is more than the number measured by our system.

2) Degree-frequency Distribution

The corresponding degree-frequency distributions of the snapshots are plotted in Fig. 4. It can be seen that the percentage of peers with high degree ($>$100) is very low (less than 0.002). The percentage of peers with degrees close to γ is about 20 times more than that of peers with degree larger than 100. The results are consistent with the results of the rank-degree distribution. We use Gaussian distribution (8) to fit the degree-frequency of each snapshot.

$$f(x) = \frac{A}{\sqrt{2\pi}\sigma} \exp\{-\frac{(x-\mu)^2}{2\sigma^2}\} \tag{8}$$

Results show that in steady stages, the degree-frequency distribution of a large-scale swarm follows a Gaussian distribution with a correlation coefficient R larger than 0.9 (e.g., Fig. 4(a)$\sim$4(d)). R is around 0.5 for small-scale swarms (e.g., Fig. 4(e)). This indicates that the BT network is very close to a random network, especially for a large-scale swarm. This

TABLE I
INFORMATION OF 5 TORRENTS

Description	Torrent 1	Torrent 2	Torrent 3	Torrent 4	Torrent 5
Created Time	2011-12-24 15:46:35	2011-11-09 8:47:45	2011-12-24 3:52:54	2011-12-20 22:58:57	2011-12-22 19:16:42
Start time	2011-12-25 11:30:27	2011-12-13 19:38:05	2011-12-24 09:41:52	2011-12-21 11:40:32	2011-12-23 22:56:00
File type	TV show	Game	Movie	Music	Movie
File Size	801MB	12.8GB	1.77GB	118MB	1.12GB
Swarm/Connected[1]	35004/9272	27483/6745	4980/1630	1379/675	392/128
Swarm/Connected[2]	31593/8832	24954/5083	4150/1259	1096/415	378/102

[1] Average swarm population and connected peers observed during the 1st 50 snapshots at the initial measurement.
[2] Average swarm population and connected peers observed during the 1st 50 snapshots at 4th day of measurement.

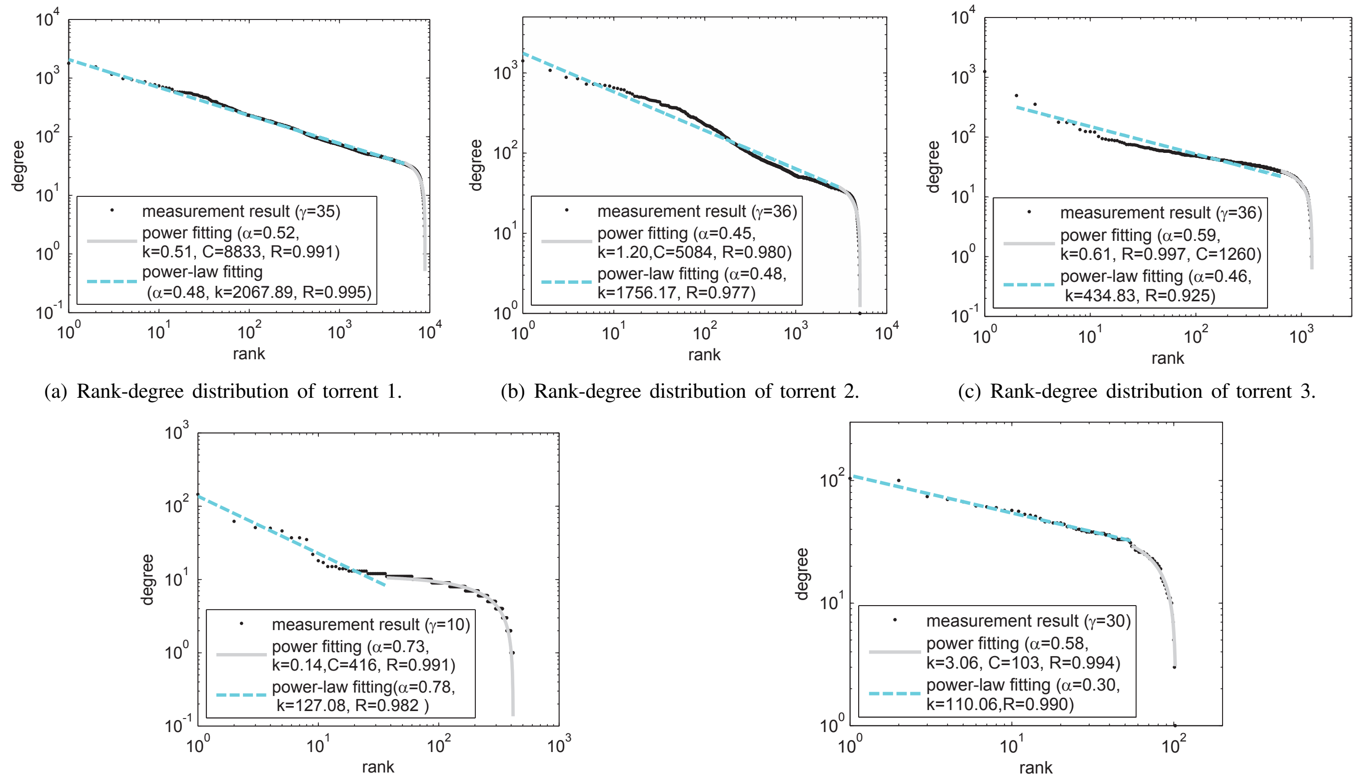

(a) Rank-degree distribution of torrent 1.
(b) Rank-degree distribution of torrent 2.
(c) Rank-degree distribution of torrent 3.
(d) Rank-degree distribution of torrent 4.
(e) Rank-degree distribution of torrent 5.

Fig. 3. Rank-degree distribution of 5 torrents with different scales.

result is different from the simulation results in [7], which shows that the BT network is not random. Our result is also different from [10], which suggests a power-law distribution with exponential cutoff. When a peer leaves a BT, this may change the connection status of peers connected to it. But these peers could continue downloading from other peers. The random characteristics of a BT network may be attributed to the random peer selection policy used by the trackers and peers.

We notice that the analysis by Al-Hamra *et al.* [7] was based on the network connectivity at the 10th minute, which is at an initial stage. We also examine the peer degree distribution at the initial stage, one result is plotted in Fig. 4(f) (torrent 3, a medium scale swarm). We find that although the degree-frequency distribution exhibits a Gaussian-like distribution, the R is small ($R = 0.759$ for torrent 3), which is similar to [7]. At the initial stage, the swarm has not reached a steady-state and a lot of peers joining the network may cause a sudden connection change.

In Fig. 4, we notice that for peer-degrees between 45 and 50, the percentage of peers is high, and there is a spike around the degree of 50. Then the percentage drops dramatically at the degree of 51. We believe that the most important reason of this phenomenon is the connection limitation of BT clients. For most BT clients, by default the maximum number of connected peers per torrent is set to 50, which limits the number of connections with peers that have high bandwidth. On the other hand, peers reaching their connections' limits may actively free some idle or slow connections and try to connect to new peers. Moreover, some clients will reserve some connections for the incoming peers. We confirm these speculations on μTorrent.

In Fig. 5, we plot the CDF of the number of peers with respect to their degrees. It shows that there are more than 70%

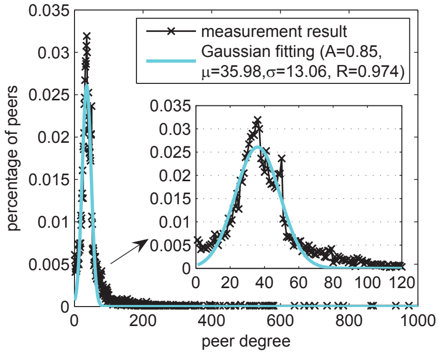

(a) Degree-frequency distribution of torrent 1.

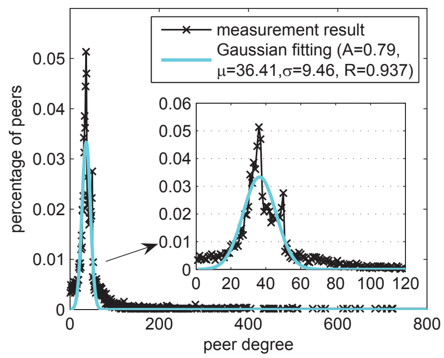

(b) Degree-frequency distribution of torrent 2.

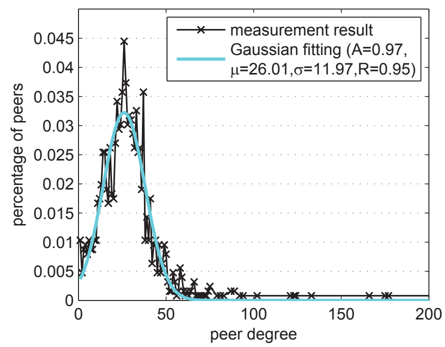

(c) Degree-frequency distribution of torrent 3 (The 117th snapshot on 2011-12-28.)

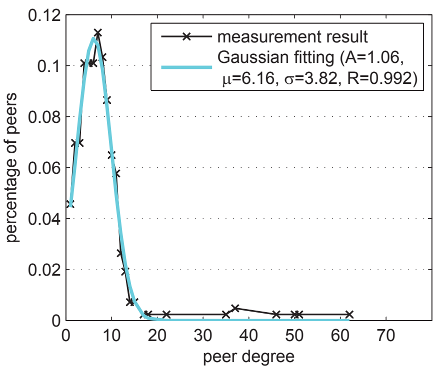

(d) Degree-frequency distribution of torrent 4.

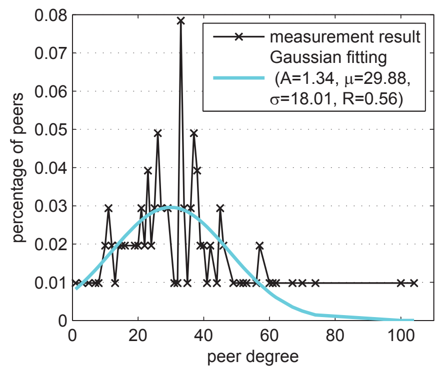

(e) Degree-frequency distribution of torrent 5.

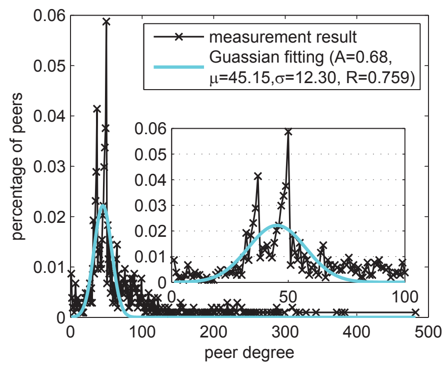

(f) Degree-frequency distribution of torrent 3 (The 41st snapshot on 2011-12-24).

Fig. 4. Degree-frequency distribution of 5 torrents with different scales.

of peers whose degrees are less than 50; 90% of peers have a degree less than 60, and less than 5% of peers have degree larger than 100. The results suggest that for most peers, the dominating factor of peer degree is not the connection limit of BT clients.

B. Small World

Many existing P2P networks (e.g., Gnutella) have been showed to be a small-world [19], in which the mean distance between nodes is small and nodes are highly clustered. Legout *et al.* [22] observed node clustering in the early stages of swarms based on PlanetLab experiments. On the other hand, Dale *et al.* [8] found no clear evidence of clustering in their experiments. Al-Hamra *et al.* [7] showed that a short diameter is essential to provide fast distributions of pieces, and PEX may generate a chain-like overlay with a large diameter. We believe that PEX helps peers know each other within a short distance. Hence, we study the peer distance and clustering coefficient to check whether the topologies of real BT networks exhibit the small-world characteristics or not.

1) Distance between Peers

In BT topologies, a path represents a peer sequence such that from each peer in the sequence there is a connection to the next peer in the sequence. The distance between each pair of peers is the shortest path length between them. If there are two peers with no path between them, the network is not fully connected. The diameter of a network is the maximum length of the shortest path between peers. Although we cannot know the complete topology of a BT swarm due to the limitation of measurements, we can use our measurement results and obtain an approximate value of the shortest path length between peers u and v, denoted as $ASL(u, v)$. Let's take the graph in Fig. 2 as an example. First, for peers $p_i, p_j \in \boldsymbol{P_c}$, we simply use the known portion of $\boldsymbol{P_c}$'s topology ($\boldsymbol{G_c}$, a sub-graph of $\boldsymbol{G}$, consisting of all peers in $\boldsymbol{P_c}$ and the connections among them), and directly compute the shortest path length between p_i, p_j as the $ASL(p_i, p_j)$. Second, $\boldsymbol{P_k}$'s topology is unknown, but we can approximate the shortest path length between peers $q_i, q_j \in \boldsymbol{P_k}$ according to their connectivity with peers in $\boldsymbol{P_c}$. For example, if q_i and q_j are connected to peers in $\{p_i\} \subseteq \boldsymbol{P_c}$ and $\{p_j\} \subseteq \boldsymbol{P_c}$ respectively, we have:

$$ASL(q_i, q_j) \le \min_{p_i \in \{p_i\}, p_j \in \{p_j\}} ASL(p_i, p_j) + 2 \qquad (9)$$

Third, for peers $p_i \in \boldsymbol{P_c}$ and $q_j \in \boldsymbol{P_k}$, if q_j is connected to peers in $\{p_j\} \subseteq \boldsymbol{P_c}$, likewise we have:

$$ASL(p_i, q_j) \le \min_{p_j \in \{p_j\}} ASL(p_i, p_j) + 1 \qquad (10)$$

We cannot obtain the shortest path length between peers in $\boldsymbol{P_n}$ because we know nothing about their topologies. However, according to our measurement result, $\boldsymbol{P_c}$ and $\boldsymbol{P_k}$ cover about 90% of the swarm population, and 51% - 75% of the connections. Hence, our measurement can reflect the connectivity of the network to a good extent.

According to the above analyses, we only need to calculate the ASL of each pair of peers in $\boldsymbol{G_c}$, in which the connections are definite. The ASL of peer pairs in $\boldsymbol{P_c}$ may be smaller than

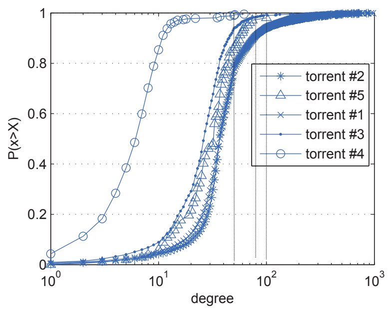

Fig. 5. CDF of the number of peers with respect to their degrees.

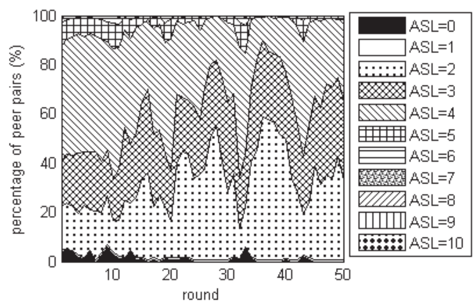

(a) Initial stage (50 snapshots on 2011-12-24).

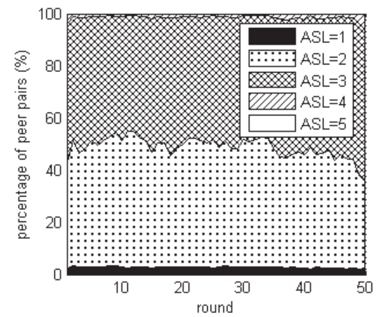

(b) Steady stage (50 snapshots on 2011-12-28).

Fig. 6. ASLs and their percentage of peer pairs in $\boldsymbol{P_c}$ of torrent 3.

our calculation. In Fig. 6(a), we plot 50 steady-stage snapshots of torrent 3 as an example. We find that $\boldsymbol{G_c}$ of most snapshots are unconnected, with a few pairs of peers ($<$10%) where there is no path between them (ASL=0 in Fig. 6(a)). The ASLs of other peer pairs are less than 10. Particularly, the ASLs of more than 80% of the pairs are less than 4. Our measurement result is consistent with the simulation result of [7]. By examining the download percentage of peers in $\boldsymbol{P_c}$, we find that most of the peers are seeders in the steady stage (discussed in subsection C), so there is no direct path between them as they do not connect with each other to download pieces, and this may be the reason why the graphs are split. The short distances indicate that a leecher is able to find out any other peers (especially seeders) by exchanging peer lists for several times, suggesting that PEX is efficient of peer discovering. In addition, the partition of the network topology suggests that trackers and DHT are necessary because they provide different peer discovery approaches that help peers in different partitions find each other.

Both [12] and our measurement find out that more than 95% BT peers use the PEX protocol. Therefore, we plot the ASL of torrent 3 in its early stage to examine whether the PEX generates the large-diameter chain-like topologies in real world as in Al-Hamra *et al.*'s [7] simulation results. The result (shown as Fig. 6(b)) actually differs from [7] because the network diameter of the initial stage is very short (less than 5+2=7). In addition, the result shows that the topology in its initial stage is a connected graph and, especially, 90% of peer pairs have a shortest path length less than 3, meaning that the connectivity in the early stage is better than that in the steady stage. Together with the results in subsection C, we confirm Zhong *et al.*'s [6] result that the increase of seeders proportion causes the increase of ASL. The results also indicate that: in the early swarm stages, a tracker can reduce its load by reducing the frequency of client requests.

2) Clustering Coefficient

Clustering coefficient is a measure of degree to which nodes in a graph tend to cluster together, and it is defined in equation (11), where T_e is the number of connections between neighbors of peer v. The clustering coefficient is usually large in a scale-free graph while small in a random graph. The clustering coefficient for the entire network is calculated by equation (12).

$$C_v = \frac{2 \times T_e}{deg(v) \times (deg(v) - 1)} \tag{11}$$

$$C = \frac{1}{N} \sum_{i=1}^{N} C_i \tag{12}$$

In Fig. 7, we plot the network clustering coefficient of 50 snapshots of torrent 3 in its early stage and steady stage of its swarm evolution. The results show that the network clustering coefficient is small (< 0.15) in the steady stage. This observation is consistent with Dale *et al.* [8] and Fauzie *et al.* [10], and further proves that BT topologies are more likely random graphs. In addition, similar to Dale *et al.*, we also notice that in the real network, the clustering coefficient of the initial stage is larger than that in the steady stage and exhibits a declining thread as the swarm evolves. This may also be caused by increased proportion of seeders who can provide more uploading connections than the leechers. Our further analysis in subsection C shows that there are more seeders in steady stage than in the initial stage. Our measurement result is also consistent with Zhong *et al.* [6] which show that the BT networks tend to be less clustered when the seeder proportion increases. A high clustering coefficient in the initial stage implies that peers collaborate in sharing the files while a low clustering coefficient in the steady stage suggest that seeders may contribute more.

Summing up 1) and 2), in the steady stage, BT networks are not fully connected and the clustering coefficient is much low, indicating that the BT networks are not small-world but are more likely random graphs.

C. Peer Download Percentage

The peer degree may be different in different download stages. For example, peers in their early download stages will try to connect as many peers as possible to increase download speed. The degrees of seeders may be also very large, which can distribute files faster to more peers. In this subsection, we examine the relationship between a peer degree and the download percentage.

We extract the BITFIELD from the logs. Each bit b_i of the BITFIELD represents whether the piece i has been downloaded (b_i=0: not downloaded; b_i=1: downloaded). The

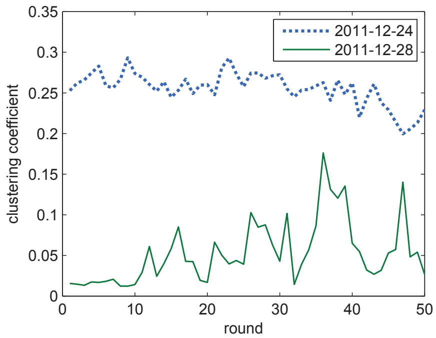

Fig. 7. Comparison of clustering coefficient of torrent 3, 50 snapshots in initial stages (2011-12-24) and the steady stages (2011-12-28).

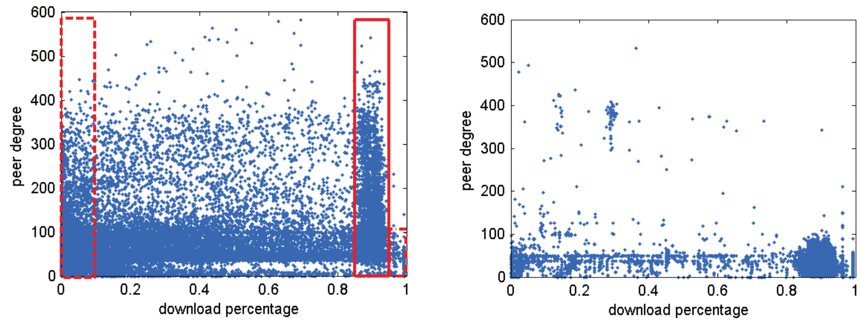

(a) Initial stage (50 snapshots on 2011-12-24, 6 hours after the torrent birth).

(b) Steady stage (50 snapshots on 2011-12-28, 102 hours after the torrent birth).

Fig. 8. Download percentage of peers and their corresponding degrees of torrent 3.

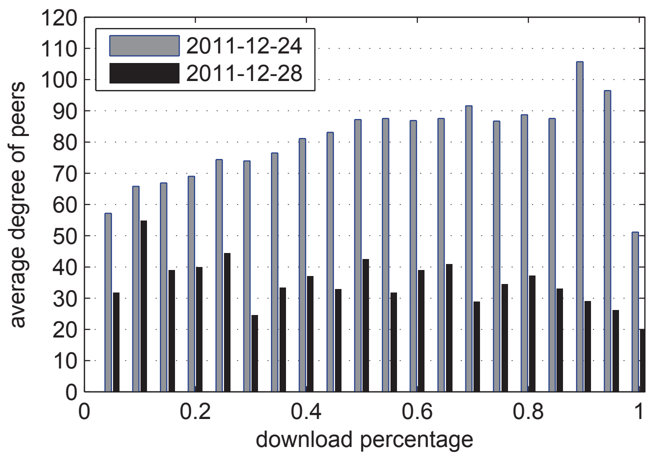

Fig. 9. Average degree of peers in download percentage ranges of torrent 3.

download percentage of a peer at time t is calculated using equation (13), where PN is the piece number of the file.

$$FR^t = \sum_{i=1}^{PN} \frac{b_i^t}{PN} \tag{13}$$

Fig. 8 shows the peers' degrees and their download percentages of torrent 3 in the initial and steady stages, respectively. To reduce the error of one snapshot, we plot 50 continuous snapshots. We also plot the average degree of peers (in Fig. 9) and the percentage of peers (in Fig. 10) within each download percentage range (with 0.05 widths). It can be seen from the figures that the degrees of peers in their early downloading stages and seeding stages are not as large as previously expected. There is no direct and obvious relationship between the peer degree and the peer download stages. We observe that in the initial stages of most torrents, as for the leechers whose degree is larger than 100, the number of peers whose $FR > 0.85$ (marked by the rectangle in Fig. 8(a)) is larger than that of peers whose FR is in other ranges. The degrees of peers with FR between 0.5 and 0.9 are usually larger than 50. The potential reason may be that in an initial-stage swarm, the number of seeders is small and peers may connect to the peers with higher complete rate because they are more likely to have the pieces demanded by other peers.

Comparing Fig. 8(a) and Fig. 8(b), we can also find that in the initial stage of the swarm evolution, the proportion of peers with small download percentage (typically, $FR < 0.2$) is higher than that of peers with large download percentage (typically, $FR > 0.8$). Fig. 10 shows that 42.46% of peers with $FR < 0.2$ and 23.98% of peers with $FR > 0.8$. While in the steady stage, the proportion of peers with large FR is much more than that of peers with small FR, e.g., 2.31% of peers with $FR < 0.2$ vs. 95.41% of peers with $FR > 0.8$, as shown in Fig. 10. This result is similar to the observation by our BT clients. The result confirms Guo's [1] modeling of swarm evolution. It also indicates that if a group of peers (even some seeders) leave the system, pieces from the remaining peers can make up a complete copy of files with a very high probability, suggesting that the swarm is robust in the steady stage from the resource perspective.

We want to point out that the actual percentage of seeders is larger than that shown in Fig. 10 because some seeding peers send BITFIELD messages with some bits being set to 0 rather than 1, and then send HAVE messages to remote peers to announce these "absent" pieces.

D. Peer Download Speed

Usually, it is considered that the more peers a node connects, the higher download speed it can have. To verify this in real networks, we compute the average download speed of each peer and its corresponding degree. The average download speed can be calculated by the incremental download amounts between two successive measurements. That is

$$downSpeed^t = \frac{(FR^{t'} - FR^t) \times fileSize}{t' - t} \tag{14}$$

where t' is the next time after t the system acquires the peer's BITFIELD, and $fileSize$ is the length of the distributed file(s). In our calculation, if we did not get the BITFIELD of the peer after t', we set the download speed at time t' the same as that at time t.

Fig. 11 plots the degree and average download speed of peers of torrent 3, for the same snapshots in Fig. 8. The figure does not include seeders with zero download speed. Intuitively, Fig. 11 shows that peers with large degrees usually have high download speed, especially in the initial swarm stage, which is consistent with common sense. However, the download speed

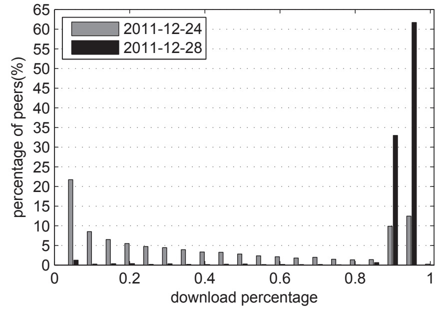

Fig. 10. Proportion of peers in download percentage ranges of torrent 3.

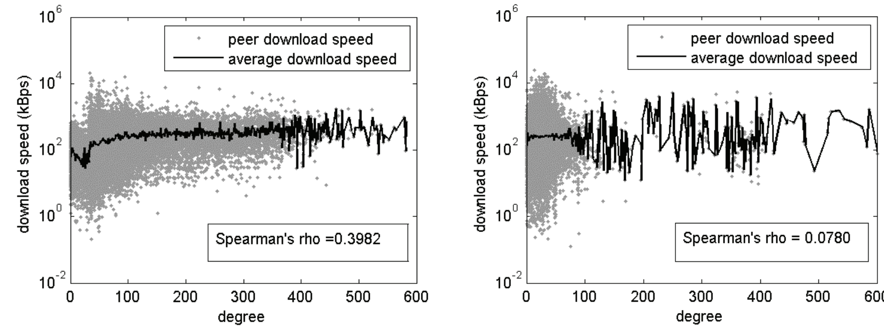

(a) Initial stage (torrent 3, 50 snapshots on 2011-12-24).

(b) Steady stage (torrent 3, 50 snapshots on 2011-12-28).

Fig. 11. Relationship between peer degree and download speed.

is not in direct proportion to the peer degree, due to the limit of the link bandwidth and other peers' upload speeds. In addition, there are a number of low-degree peers that can get a download speed close to even higher than the large-degree peers (e.g., 5.7% peers have larger than 200kBps download speed but their connection degree are less than 30). To measure the dependence between the peer connection degree and the download speed, we compute the Spearman's rank correlation coefficient (Spearman's rho) of the 100 snapshots, the same snapshots used to plot Fig. 11. We find that the Spearman's rhos are low (around 0.0174 to 0.4754), which shows that the correlation between the peer degree and the download speed is not significantly positive. This suggests that increasing peer connections may not necessarily improve the download speed.

The possible reason for the high download speed but low degrees is that: the low-degree peer has high bandwidth and is downloading from high-uploading-rate peers. However, the number of connections is limited due to the following possible reasons: a) it knows (or can only connect to) these limited peers, or b) it may be downloading multiple torrents, but the total connection number of other torrents has reached the global maximum number of connections. Therefore, it will not connect to more peers for this torrent.

V. Conclusion

In this paper, we presented a measurement study on the topology characteristics of a real BT network. We designed a hybrid BT measurement system with a combination of active and passive approaches and explored the BT peer exchange extension (PEX) to collect connection information among peers. Our measurement results cover more than 95% of peers and 51% $\sim$ 75% connections of each snapshot.

The main results of this paper are summarized below: (1) We discovered that a steady-stage BT network has a Gaussian-like peer-degree-frequency distribution, especially for a large scale swarm. We believe that the steady stage of a BT network is closer to a random graph rather than a scale-free graph, which is different from the previous results. The small proportion of low-degree peers indicates the robustness of BT networks. (2) We showed that the BT is not a small world due to the short peer distance and low clustering coefficient. Especially, the topologies are not fully connected in the steady stage, which implies that trackers and DHT are required to help peers in different partitions find each other. Different from previous work, we did not observe the large-diameter chain-like topologies although more than 95% of peers use the PEX extension. (3) Neither the newly joined peers nor seeders have the very high degrees that we expected. We found that the download percentages of peers in the steady stage are larger than that in the initial stage, showing that the BT swarm is robust to peer departures. We also found that the large proportion of seeders increases the distances of peers and even makes the topologies not fully connected. (4) We observed that peers with large degrees can get relatively high download speeds while some low-degree peers can download at fast speeds. The low Spearman's rank correlation coefficient indicates that there is no strong correlation between the peer connection degree and the download speed. This suggests that increasing peer connections may not necessarily improve the download speed.

References

[1] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "A performance study of bittorrent-like peer-to-peer systems," *IEEE J. Sel. Areas Commun.*, vol. 25, pp. 155–169, Jan. 2007.

[2] J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips, "The bittorrent P2P file-sharing system: Measurements and analysis," in *Proc. 4th Int. Workshop Peer-to-peer Syst.*, 2005.

[3] G. Neglia, G. Reina, and H. G. Zhang, "Availability in bittorrent systems," in *Proc. 26th IEEE Int. Conf. Comput. Commun.*, 2007, pp. 2216–2224.

[4] C. Dale and J. C. Liu, "A measurement study of piece population in bittorrent," in *Proc. IEEE Global Telecommun. Conf.*, 2007, pp. 405-410.

[5] P. Dhungel, X. Hei, D. Wu, and K. W. Ross, "A measurement study of attacks on bittorrent seeds," in *Proc. IEEE Int. Conf. Commun.*, 2011.

[6] L. Zhong, X. F. Wang, and M. Kihl, "Topological model and analysis of the P2P bittorrent protocol," in *Proc. 9th World Congress Intelligent Control Automation*, 2011, pp. 753–758.

[7] A. Al-Hamra, A. Legout, and C. Barakat, "Understanding the properties of the bittorrent overlay," INRIA, Sophia Antipolis, France, Tech. Rep., July 2007.

[8] C. Dale, J. C. Liu, J. Peters, and B. Li, "Evolution and enhancement of bittorrent network topologies," in *Proc. IEEE 16th Int. Workshop Quality Service*, 2008.

[9] *PlanetLab* [Online]. Available: http://www.planet-lab.org/

[10] M. D. Fauzie, A. H. Thamrin, R. V. Meter, and J. Murai, "A temporal view of the topology of dynamic bittorrent swarms," in *Proc. IEEE Conf. Comput. Commun. Workshops*, 2011, pp. 894–899.

[11] L. Ye, H. L. Zhang, F. Li, and M. J. Su, "A measurement study on bittorrent system," *Int. J. Commun., Netw. Syst. Sci.*, vol. 3, no. 12, Dec. 2010.

[12] D. Wu, P. Dhungel, X. J Hei, C. Zhang, and K. W. Ross, "Understanding peer exchange in bittorrent systems," in *Proc. 10th IEEE Int. Conf. P2P Comput.*, 2010.
[13] H. Y. Wang, J. C. Liu, and K. Xu, "On the locality of bittorrent-based video file swarming," in *Proc. 8th Int. Conf. P2P Syst.*, 2009.
[14] J. C. Liu, H. Y. Wang, and K. Xu, "Understanding peer distribution in the global Internet," *IEEE Netw.*, vol. 24, pp. 40–41, July 2010.
[15] L. D. Yu and M. Chen, "Zipf-law: A measurement study of geographical distribution of peers in bittorrent swarms," in *Proc. 6th Int. Conf. Wireless Commun. Netw. Mobile Comput.*, 2010.
[16] T. Hossfeld, F. Lehrieder, D. Hock, S. Oechsner, Z. Despotovic, W. Kellerer, *et al.*, "Characterization of bittorrent swarms and their distribution in the Internet," *Comput. Netw.*, vol. 55, no. 5, pp. 1197–1215, Apr. 2011.
[17] *Libtorrent* [Online]. Avaliable: http://www.rasterbar.com/products/libtorrent/
[18] *TorrentZ* [Online]. Avaliable: http://torrentz.eu/
[19] D. Stutzbach, R. Rejaie, and S. Sen, "Characterizing unstructured overlay topologies in modern P2P file-sharing systems," *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 267–280, Apr. 2008.
[20] S. Saroiu, K. P. Gummadi, and S. D. Gribble, "Measuring and analyzing the characteristics of Napster and Gnutella hosts," *Multimedia Syst.*, vol. 9, no. 2, pp. 170–184, Aug. 2003.
[21] L. Vu, I. Gupta, K. Nahrstedt, and J. Liang, "Understanding overlay characteristics of a large-scale peer-to-peer IPTV systems," *ACM Trans. Multimedia, Comput., Commun. Appl.*, vol. 6, no. 4, pp. 1-24, Nov. 2010.
[22] A. Legout, N. Liogkas, E. Kohler, and L. Zhang, "Clustering and sharing incentives in bittorrent systems," in *Proc. ACM Int. Conf. Meas. Modeling Comput. Syst.*, 2007, pp. 301–312.
[23] Y. H. Liu, X. M. Liu, L. Xiao, L. M. Ni, and X. D. Zhang, "Location-aware topology matching in P2P systems," in *Proc. IEEE Conf. Comput. Commun.*, 2004, pp. 2220-2230.
[24] B. Liu, Y. Cui, Y. S. Lu, and Y. Xue, "Locality-awareness in bittorrent-like P2P applications," *IEEE Trans. Multimedia*, vol. 11, pp. 361–371, Apr. 2009.
[25] S. S. Ren, E. H. Tan, T Luo, S. Q. Chen, L. Guo, and X. D. Zhang, "TopBT: A topology-aware and infrastructure-independent bittorrent client," in *Proc. IEEE Conf. Comput. Commun.*, 2010, pp. 1–9.
[26] *µTorrent* [Online]. Avaliable: http://www.utorrent.com/
[27] *Vuze* [Online]. Avaliable: http://www.vuze.com/
[28] *BitTorrent Client* [Online]. Avaliable: http://www.bittorrent.com/
[29] *BitTorrent Tracker* [Online]. Avaliable: http://en.wikipedia.org/wiki/BitTorrent_tracker
[30] *BitTorrent DHT Protocol* [Online]. Avaliable: http://bittorrent.org/beps/bep_0005.html
[31] *The PirateBay* [Online]. Avaliable: http://thepiratebay.se/
[32] *CERNET* [Online]. Avalible: http://www.edu.cn/
[33] A. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, Oct. 1999.

Majing Su received the B.S. and M.S. degrees in computer science from Harbin Institute of Technology (HIT), Harbin, China, from 2003 to 2009. Since 2009, she has been working as a Ph.D. candidate in the Department of Computer Science and Technology at HIT. Her research interests include P2P measurement and security.

Hongli Zhang received her B.S. degree in computer science from Sichuan University, Chengdu, China, in 1994, and her Ph.D. degree in computer science from Harbin Institute of Technology (HIT), Harbin, China, in 1999. She is a professor in the Computer Science and Technology Department of HIT and is vice director of the National Computer Information Content Security Key Laboratory. Her research interests include network and information security, network measurement and modeling, and parallel processing.

Xiaojiang (James) Du is currently an associate professor in the Department of Computer and Information Sciences at Temple University. Dr. Du received his B.E. degree from Tsinghua University, China, in 1996, and his M.S. and Ph.D. degrees from the University of Maryland, College Park, in 2002 and 2003, respectively, all in electrical engineering. His research interests are security, systems, wireless networks, and computer networks. Dr. Du has published over 100 journal and conference papers in these areas, and has been awarded more than $3M in research grants from the U.S. National Science Foundation and Army Research Office. He serves on the editorial board of four international journals.

Binxing Fang received his B.S. degree in computer science from Harbin Institute of Technology of China in 1981. He received his M.S. and Ph.D degrees in computer science from Tsinghua University and Harbin Institute of Technology of China in 1984 and 1989, respectively. He is currently a member of the Chinese Academy of Engineering. His current research interests include information security, information retrieval, and distributed systems.

Mohsen Guizani (S'85-M'89-SM'99-F'09) is currently a Professor and the Associate Vice President for Graduate Studies at Qatar University, Doha, Qatar. He was Chair of the Computer Science Department at Western Michigan University from 2002 to 2006 and Chair of the Computer Science Department at the University of West Florida from 1999 to 2002. He has also served in academic positions at the University of Missouri-Kansas City, University of Colorado-Boulder, Syracuse University, and Kuwait University. He received his B.S. (with distinction) and M.S. degrees in electrical engineering; and M.S. and Ph.D. degrees in computer engineering in 1984, 1986, 1987, and 1990, respectively, from Syracuse University, Syracuse, New York. His research interests include computer networks, wireless communications, mobile computing, and optical networking. He currently serves on the editorial board of six technical journals and is the Founder and EIC of the *Wireless Communications and Mobile Computing* journal published by John Wiley (http://www.interscience.wiley.com/jpages/1530-8669/). He is the author of eight books and more than 300 publications in refereed journals and conferences. He guest-edited a number of special issues in IEEE journals and magazines. He has also served as member, Chair, and General Chair of a number of conferences. Dr. Guizani served as Chair of the IEEE Communications Society Wireless Technical Committee (WTC) and Chair of the TAOS Technical Committee. He was an IEEE Computer Society Distinguished Lecturer from 2003 to 2005. Dr. Guizani is an IEEE Fellow and a Senior member of ACM.