Understanding Peer-level Performance in BitTorrent: A Measurement Study

Amir H. Rasti, Reza Rejaie
University of Oregon
{*amir,reza*}@*cs.uoregon.edu*

***Abstract*—The observed performance by individual peers in BitTorrent can be simply measured by their average download rate. While it is often stated that the observed peer-level performance by BitTorrent clients is high, it is difficult to accurately verify this claim due to the large scale, distributed and dynamic nature of this P2P system. To provide a "representative" characterization of peer-level performance in BitTorrent, the following two important questions should be addressed: *(i)* What is the distribution of observed performance among participating peers in a torrent? *(ii)* What are the primary peer- or group-level properties that determine observed performance by individual peers?**

In this paper, we conduct a measurement study to tackle these two questions. Toward this end, we derive observed performance for nearly all participating peers along with their main peer- and (peer-view of) group-level properties in three different torrents. Our results show that the probability of experiencing certain level of performance has a roughly uniform distribution across the entire range of observed values. Furthermore, while the performance of each peer has the highest correlation with its outgoing bandwidth, there is no dominant peer- and group-level property that primarily determines the observed performance by the majority of peers.

I. Introduction

During recent years, the Internet has witnessed a rapid increase in the popularity of BitTorrent and therefore its contribution in network traffic. BitTorrent is a peer-to-peer (P2P) content distribution mechanism that enables a single node to provide its static content to a large number of peers without requiring a large access link bandwidth. BitTorrent incorporates swarming content delivery to effectively utilize the outgoing bandwidth of participating peers and thus achieve scalability. The scalability of BitTorrent along with the ease of deployment has led to its increasing popularity over the Internet. This in turn has motivated researchers to examine the performance of BitTorrent using different techniques including modeling ([7]), simulations ([1], [7]), and in particular measurement ([3], [2], [4]).

One key aspect of performance in BitTorrent is the download rate that is achieved by participating peers. It is often stated (by users and developers) that BitTorrent provides a good performance to individual peers, *i.e.*, users can effectively utilize their available (or configured) incoming access link bandwidth. However, capturing a "representative" value of observed peer performance in practice is a non-trivial task. A typical measurement approach to study peer-level performance is to use one (or multiple) instrumented BitTorrent client(s) that participate in an existing torrent and download content [3], [5]. While this approach provides detailed information about the observed performance by a few instrumented peers, it is unclear whether its findings properly represent observed behavior by the entire population, *i.e.*, the results may not be *representative*. Intuitively, the observed performance by individual peers in a torrent could depend on their peer-level properties (*e.g.*, outgoing access link bandwidth) or group-level properties (*e.g.*, group population, content availability or churn). This implies that results of a measurement study using "instrumented client" could easily depend on time of measurement, location of instrumented clients or properties of the torrent and thus they are not representative. In essence, to investigate the peer-level performance in BitTorrent, the following two important and related questions should be addressed:

- What is the distribution of the observed performance by individual peers in a torrent?
- What peer- or group-level properties primarily determine the observed performance by individual peers in a torrent?

The first question reveals how similar (or dissimilar) are the observed performance by individual peers. This in turn determines whether a few peers can properly represent the entire population of participating peers and how they should be selected. The second question explores any potential dominant factor(s) that affect the observed performance by individual peers. To our knowledge, these questions have not been investigated by previous measurement studies on BitTorrent.

In this paper, we try to answer these two important questions by capturing the observed performance for almost all participating peers in several torrents with different groups of users. We present our methodology to derive peer- and group-level properties of participating peers in a torrent from its tracker log. We also describe various challenges in our approach including the difficulty to accurately estimate the observed performance by all participating peers. To tackle the first question, we examine the distribution of the derived peer- and group-level properties in our candidate torrents and illustrate that both the observed performance and the observed group level properties significantly vary among participating peers. To answer the second question, we investigate the correlation between the observed performance by each peer and both its peer-level and its observed group-level properties using several classic techniques. Our analyses demonstrate that while the performance of each peer has the highest correlation with its outgoing bandwidth, there is no dominant peer- or group-level property that primarily determines the observed performance by the majority of peers. This suggests that the dominant determining factors for the observed performance by individual peers is different. This paper makes the following contributions: *(i)* It presents set of techniques to accurately derive peer-level and group-level properties of all participating peers in a torrent, *(ii)* it illustrates that the commonly used approach of "instrumented clients" is inappropriate to characterize peer-level performance in BitTorrent and *(iii)* it provides several evidences that the relationship between the peer-level performance and other peer- or group-level properties is non-trivial, and there is no dominant factor that determine peer-level performance.

The rest of this paper is organized as follows: In Section II, we present a brief overview of BitTorrent to provide the required background for our study. We will briefly present and categorize the previous work on BitTorrent in Section III. Our measurement methodology, our dataset and our data processing are described in Section IV. We present the peer-level and group-level properties in Section V. In Section VI, we examine the correlation between peer performance and various properties. Finally, Section VII concludes the paper and sketches our future plans.

II. BitTorrent: An Overview

To provide the proper context for our study, we present a brief overview of those aspects of BitTorrent that are relevant to our measurement and characterization. In BitTorrent, all participating peers that join the system to download the same file are referred to as a "torrent". All peers in a torrent form a random mesh and incorporate swarming content by pulling their missing segments from connected peers. Peers are generally divided into two groups, seeds and leechers, that have the entire or part of the entire file, respectively. All peers provide content to their neighbors but only leechers need to download content.

BitTorrent features a peer-level incentive mechanism among connected peers called tit-for-tat. This mechanism tends to connect together peers with similar ability to provide content. Therefore, the tit-for-tat mechanism can affect achieved download rate by individual peers. In general, the observed download and upload rates by each peer is limited by its incoming and outgoing access link bandwidth, respectively. However, these rates could be further limited by the user and by cross traffic. Peers can join and leave a torrent in an arbitrary fashion. These dynamics of peer participations (or churn) could also affect observed performance by individual peers.

For each torrent, there is a well known node called *tracker*. The tracker keeps track of all the participating peers in a torrent as well as their download progress. Each peer contacts the tracker when it joins or leaves a torrent, or requires more neighbors. Each peer also periodically (every 30 minutes) reports its total amount of uploaded and downloaded bytes, among other information, to the tracker. The tracker records all these interactions (including peer arrival, departure and periodic updates) in a log file. In the next section, we describe how the tracker log of a torrent can be leveraged to characterize its peer- and group-level properties.

III. Related Work

Due to its growing popularity, BitTorrent has been the subject of numerous studies in the last few years. Previous studies on BitTorrent can be broadly divided into three groups. *(i)* One group of studies use modeling and statistical analysis to represent some (often group level) properties of BitTorrent system (*e.g.*, [7]). These studies often do not consider various dynamics in the system nor incorporate subtle aspects of the protocol. More importantly, validating a model requires a good understanding of "representative" behavior of the system through measurement which is difficult as we discuss in this paper. *(ii)* Another group of studies use simulations to investigate some aspects of BitTorrent (*e.g.*, [1]). Using simulation allows one to control various details and capture any desired aspects of the protocol behavior. However, the key challenge is to properly simulate different aspects of peer behavior including dynamic properties of user participation during and after completion of download (*i.e.*, lingering time), as well as bandwidth heterogeneity and asymmetry. A good understanding of these properties and dynamics is required in order to incorporate them in a simulation. Again, such an understanding should be obtained through a representative measurement study. *(iii)* The last group of studies conduct measurement on a real torrent to capture some peer-level or group-level properties of BitTorrent (*e.g.*, [4], [5], [3]). The key question in a measurement study is "whether the captured population of peers provides a representative view of the desired property?". To our knowledge, none of the previous studies have addressed this issue. They either captured their desired property at a few instrumented clients (*e.g.*, [5]) without providing any evidence that these peers properly represent the entire group or use tracker logs to derive evolution (or distribution) of some group level properties (*e.g.*, [3], [6]). In summary, while conducting a representative measurement study on BitTorrent is necessary to characterize its behavior, previous studies did not explicitly illustrate that they have captured a representative view of their desired property.

IV. Measurement Methodology

A common approach to study BitTorrent is to run multiple instrumented clients and capture their observed performance. This approach provides detailed information (*e.g.*, access link bandwidth, variations of download rate over short timescales) about observed performance by several peers. However, this approach has two important limitations: *(i)* Since the distribution of observed performance among participating peers is unknown, the observed performance may not provide a representative view of the entire population. *(ii)* This approach does not provide any group-level information (*e.g.*, average content availability, group population) that might have a significant impact on the peer-level performance.

To address this problem, we leverage BitTorrent tracker log to estimate both peer-level properties (namely download and upload rate) for *nearly all* participating peers and key group-level properties (*i.e.*, churn rate, content availability, and group population) that are observed by individual peers. This information allows us to answer our two key questions. It is worth noting that this approach has its own limitations as follows: First, as we explained earlier, each peer sends an update of its download (and upload) progress once every 30 minutes. Therefore, we can only estimate "average" peer-level properties over 30 minute timescale. This implies that variations on download and upload rates over shorter timescale can not be captured by this approach. Second, the tracker log does not contain any information about the connectivity between participating peers (*i.e.*, shape of the overlay topology). Therefore, we are not able to examine the potential effect of content availability among neighbors of a given peer on its performance. Third, the tracker log does not provide any explicit information about the maximum download or upload rate that each peer is able (willing or configured) to achieve. This could affect the accuracy of estimated performance by each peer. We further elaborate on this issue and explain our approach to address this problem in subsection IV-C. In the next two subsections, we describe how peer- and group-level properties are derived from tracker logs.

A. Deriving Peer-level Properties

We only focus on two peer-level properties: download and upload rates. Toward this end, we define a *session* as a collection of events associated with a single appearance of a particular peer in a torrent. A complete session starts with a *sign-in* event in the tracker log, continues with several periodic updates, and finally ends with a *sign-out* event. Note that the tracker log is missing "sing-in" (or "sign-out") events of a session if these events occur outside our logging window. A session may also include a *download completion* event which implies that a peer has become a seed. In our study, we only focus on the observed performance by leechers until they complete their download.

The average download (or upload) rate for a particular peer between two consecutive updates is estimated by dividing the increase in the amount of downloaded (or uploaded) bytes during this interval updates by its duration. This leads to several rather short term average upload and download rates (one per update) for each peer. The average download and upload rate during the entire session can be similarly estimated by comparing the first and the last (or the download completion if it occurs during the session) reports. Figure 1(a) depicts the evolution of downloaded and uploaded bytes over time for a single peer. The slope between two consecutive points represents the average download/upload rate for that interval whereas the slope of the line that connects the first and last points represents the average rate across the entire session. This figure clearly shows that: *(i)* This peer completes its download shortly before 5pm but remains in the system as a seed, and *(ii)*its average upload rate is higher than its average download rate.

B. Deriving Peer-View of Group-Level Properties

Our goal is to derive the average value of key group-level properties, namely population, churn rate, and content availability, that are viewed by individual peers during their session. To achieve this goal, we first derive evolution of these properties over time. We sample the value of these group-level properties at evenly spaced points in time. Figure 1(b) demonstrates this approach by showing the arrival time of all received updates from each peer (with a circle) on a horizontal line. At each sampling point, we only consider the last report before and the first report after the sampling point for each active session (shown with a filled circle). Given the available content at each peer in these two reports, we can estimate the available content at that peer at the sampling point. Then, we can average the available content across all active peers to estimate average content availability at that sampling point. Counting the number of active sessions at a sampling point provides an estimate for the population of peers at that point of time. Comparing the identity of peers at a sampling point with the last sampling point reveals the number of departed or arrived peers since the last sampling point. Using this information, we can estimate the evolution of these group-level properties during the appearance of individual peers (*i.e.*, a session). Then, the peer-view of these group-level properties can be derived for each peer by averaging their values during that peer's presence as a leecher. More specifically, we focus on the average value of group population, churn rate and content availability during the downloading time of each peer to derive the observed value of these group properties for that particular peer. In essence, peer view of these properties represents the state of the group during the appearance of a peer as a leecher.

C. Performance Metrics

The main goal of individual peers in a torrent is to maximize their download rate. To determine observed performance by individual peers, we should measure their ability to utilize their access link bandwidth, *i.e.*, the ratio of average download rate to the maximum rate that a peer is able and willing to receive content (*i.e.*, its physical, available or configured incoming bandwidth). However, the tracker log does not provide any

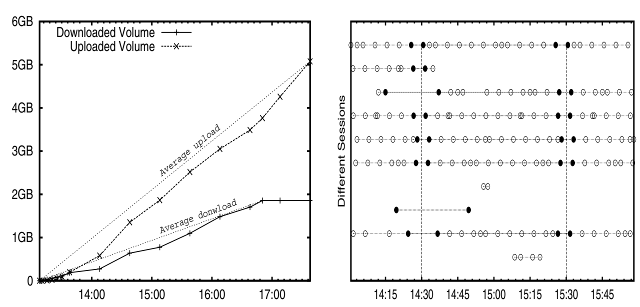

(a) Calculating average download and upload rate (b) Sampling method for group properties

Fig. 1. Capturing peer- and group-level properties

explicit information about incoming access link bandwidth of individual peers. Therefore, we use the maximum value of per-interval download rate for each peer as an estimate for its access link bandwidth. Since the measured download rates are averaged over 30 minute intervals, they provide a "loose" lower bound for maximum download rate. Using this estimate of incoming access link bandwidth, we define the following two performance metrics for each peer:

- *Access Link Utilization*: The ratio of average download rate to its maximum value during a session estimates the utilization (or relative performance) of a peer.
- *Variability of Download Rate*: The ratio of standard deviation of per-interval download rate to its average value during a session represents the normalized variations or the variability of observed performance by each peer.

The first metric accurately captures performance of each peer but it is sensitive to the estimated access link bandwidth for individual peers. The second metric does not depend on access link bandwidth but it is a rather indirect measure of performance.

D. Data Set

We have examined the tracker logs for more than 4185 torrents and selected three torrents from different user communities, namely RedHat (RH) and Debian (DE) Linux distributions, and a 3D Game software(GA). Table I summarizes the characteristics of these three torrents. This table shows that the tracker logs have been collected at different points of time that are at least a few weeks long, and have different population and different number of sessions. The diversity across these tracker logs enables us to determine whether our findings are rather common or specific to a particular torrent. We have conducted several sanity checks on tracker logs to identify any potential error in our dataset. For example we examined whether reported amount of download/upload data by all peers always monotonically grows. We discovered that a small fraction of (potentially buggy) clients do not pass this condition. We also noticed that there are some gaps in some tracker logs (*i.e.*, no event is recorded for a couple of hours). This could occur when tracker becomes unreachable for any reason. We have removed information about any misbehaving session from our logs and only focus on the portion of logs that does not contain any gap to avoid any significant error in our analysis. We also remove all the short-lived sessions (with uptime less than 30 minutes) since their performance could be significantly affected by their short stay in the system.

V. Distribution of Observed Properties

In this section, we examine the stability of peer-level and group-level properties in our three candidate torrents. Toward this end, we try to answer the first question that we raised earlier as follows: *What is the distribution of the observed peer-level and group-level properties among participating peers in a torrent?* Note that participating peers in a torrent may appear at different points of time during our long measurement period.

TABLE I
Characteristics of three selected torrents

Community	#Start Time	End Time	#Sessions	Max. Pop.
Red Hat	3/2003	8/2003	170814	3684
Debian	2/2005	3/2005	139736	91
3D Games	10/2004	12/2004	195660	1530

We explore this issue for peer-level and group-level properties in the following subsections:

A. Peer-Level Properties

Figures 2(a) and 2(b) present the distribution (CDF) of two peer-level performance metrics, the average utilization of incoming access link and the normalized standard deviation of download rate, among participating peers in all three torrents where each torrent is labeled with its corresponding community. These distributions reveal several interesting points as follows: First, despite the differences among these torrents, the distribution of each performance metric has an interestingly similar shape for all three torrents. Second, around 10% of participating peers in Figure 2(b) exhibit significant variations in their download rates, *i.e.*, experience poor performance. If we exclude these low-performing outliers from Figure 2(b), both figures depict a pretty smooth distribution without any dominant mode. This implies that the probability of experiencing a certain level of performance (between 0 and 1) is rather similar. *In a nutshell, our results from all three torrents illustrate that participating peers in a torrent experience a rather diverse performance with a roughly uniform distribution.*

To explore the behavior of other peer-level properties, Figure 2(c) shows the distribution of normalized standard deviation of upload rate for all three torrents. These distributions are very similar to those for normalized download rate (in Figure 2(b)) which suggest that the contribution of participating peers into their torrent is rather diverse. Furthermore, the contribution of participating peers in the RedHat torrent is higher than the Debian torrent, and in the Debian torrent is higher than the Gaming torrent. To explain this we note that participating users in the RedHat and Debian torrents are usually tech-savvy clients that have nodes with higher bandwidth connectivity and processing capabilities. The similarity between the distribution of normalized download and upload rates suggest that they might be correlated. We will further examine this issue in the next section.

B. Peer-view of Group-Level Properties

We now turn our attention to the observed group-level properties and examine their variability among participating peers in a torrent. Figure 3(a), 3(b), and 3(c) present the distribution of peer-view of three key group-level properties among participating peers in our three candidate torrents. The distribution of average group population (in 3(a)) is clearly different across three torrents. The RedHat torrent contains the initial flash crowd where the population of peers varies between 200 to 3500 peers, and around 55% of peers complete their download during this initial phase. The Debian and Gaming tracker logs do not contain the flash crowd phase.

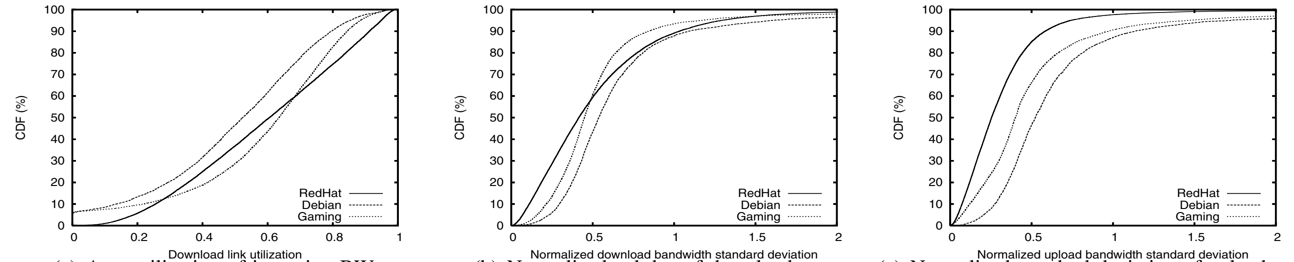

(a) Avg. utilization of incoming BW (b) Normalized std-dev of download rate (c) Normalized standard deviation of upload rate

Fig. 2. Distribution of performance metrics across all peers for three torrents

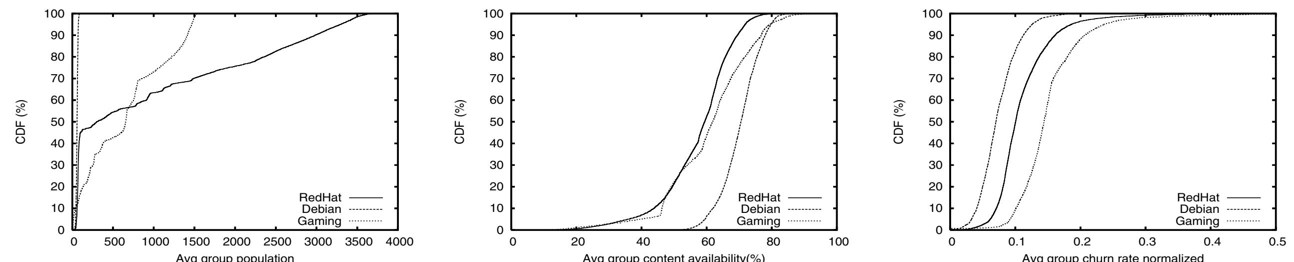

(a) Dist. of Group Population (b) Dist. of Avg. Content Availability (c) Dist. of Avg Churn Rate

Fig. 3. Distribution of peer-view of group-level properties for three torrents

The observed group population by peers in the Debian torrent changes between 50 to 1500 whereas the population of the Gaming torrent remains rather stable around 50 peers. In short, the observed average group population among participating peers in these three torrents exhibit significantly different characteristics. Despite this difference in group population, the distribution of observed content availability and churn among participating peers in each torrent is rather similar (*i.e.*, distribution is not too skewed). More specifically, peers in each torrent have experienced around 50-70% average content availability among their coexisting peers in the system (not necessarily among their neighbors), and the average observed churn rate by peers is different across these torrents but is rather similar among peers in each torrent. In summary, our results show that the participating peers in a torrent do not experience a similar performance. Except for group population, other peer- and group-level properties exhibit similar overall trends across different torrents.

VI. Identifying Underlying Factors

In this section, we tackle the second question that we raised earlier as follows: *What are the peer- or group-level properties that primarily determine the observed performance by individual peers in a torrent?*. To answer this question, we derive the observed performance by each peer (based on both performance metrics) along with its peer- and group-level properties. Given this information for all peers in a torrent, we leverage several classic techniques to identify (either qualitatively or quantitatively) any correlation between each performance metric and the following key properties that we discussed in the previous section: upload rate, group size, content availability and churn. Simple techniques such as scatter-plots did not reveal any clearly visible correlation between the observed performance and peer- or group-level properties. Therefore, we focus on more elaborate techniques in this section.

A. Linear Regression

We perform linear regression (using Splus) as a classic statistical technique to establish a linear relationship between observed performance by each peer and its main peer- and group-level properties(*i.e.*, median upload rate, average group population, average content availability and average churn rate). The derived model also quantifies the impact of each property on the overall performance. To minimize the effect of outliers, we remove any session whose properties are within the top or bottom 10% of observed range of values.

Each row of Table II presents the coefficients for different properties that represent the derived linear model by this technique for RedHat torrent. The results for other torrents are similar. For each torrent, we examined each performance metric in the following three scenarios in a progressive fashion: *(i)* The base model that relates a performance metric with all properties, *(ii)* the model that relates a performance metric to the log value of some properties (population and upload rate), and *(iii)* same as step *(ii)* but we use the "step" function in Splus to simplify the model by removing least important factors when possible. The goal in examining log value of properties is to reduce the range of values for properties which in turn could reveal any non-linear relationship that might exist as well. Each row also includes "R-Squared" value which estimates the percentage of sessions that can be properly predicted by the derived model.

As Table II indicates, all R-square values are smaller than 0.1 (*i.e.*, models can predict performance in less than 10% of sessions). In essence, this table provides a clear evidence that there is no simple linear (or non-linear) model that properly captures the relationship between the observed performance and other examined properties for individual peers. Therefore,

TABLE II
LINEAR REGRESSION RESULTS FOR REDHAT TORRENT. (COEFFICIENT, P-VALUE)

Model	R-square	outbw.50p	avg.grp.pop	avg.grp.cont.avail	avg.grp.churn
util	0.0651	0.0091, 0	-0.1206, 0	0.3493, 0	0.0015, 0
util-log	0.0603	0.0965, 0	-0.0311, 0	0.4367, 0	0, 0
util-step	0.0603	0.0965, 0	-0.0309, 0	0.4358, 0	removed
sdev	0.0709	-0.0142, 0	0.2245, 0	-0.3344, 0	-0.0029, 0
sdev-log	0.0741	-0.1585, 0	0.0778, 0	-0.6486, 0	-0.0005, 0.0095
sdev-step	0.0741	-0.1585, 0	0.0778, 0	-0.6486, 0	-0.0005, 0.0095

instead of deriving a model that incorporates all the properties, we explore the pairwise correlation between the observed performance and each property which is easier to observe in the next subsection.

B. Spearman's Rank Correlation

We use Spearman's rank correlation test as powerful technique to quantify the degree of correlation between each performance metric and peer- and group-level properties. Spearman's rank correlation coefficient is a non-parametric measure of correlation that assesses how well an arbitrary monotonic function could describe the relationship between two variables, without making any assumptions about the frequency distribution of the variables. Table III presents the Spearman's rank correlation coefficient between our two performance metrics and the following properties for all three candidate torrents in Table I: normalized standard deviation of upload rate (dev.upload), average group population (pop), average content availability (cont), and average churn rate (churn).

This table illustrates several interesting points: First, despite difference between three torrents, both performance metrics appears to have the highest correlation with the upload rate. This suggests that the variability of upload rate has the highest effect on the observed performance which in turn implies that tit-for-tat mechanism has the most noticeable impact on performance. Furthermore, in all three torrents, the correlation coefficient between our two performance metrics and outgoing bandwidth have a close absolute value with opposite signs. Second, aside from the upload rate, the effect of other parameters on observed performance by individual peers seems to vary across different torrents and in some cases between different performance metrics. In the RedHat torrent, both performance metrics have a relatively stronger correlation with group population and churn. This could be due to the fact that the log for this torrent contains the initial flash crowd phase where more than half of the captured sessions lie. Average content availability has a relatively larger coefficient for both metrics in the Debian torrent. This is most likely due to the small population in this torrent. Finally, in the Gaming torrent, the first performance metric (access link utilization) has a small and comparable coefficient for all three properties while the second performance metric (normalized standard deviation of download rate) has larger coefficients for all three properties. Clearly, the potential error in estimating the incoming access link bandwidth could affect the derived coefficients in our analysis. However, since the coefficients for upload rates and both metrics are similar, we believe that the impact of error on the largest coefficients is rather small. In summary, our results suggest that upload rate by individual peers (*i.e.*, its contributed upload rate to the system) has the primary effect on their observed performance. This finding is based on both performance metrics and across all three torrents.

TABLE III
SPEARMAN'S RANK CORRELATION COEFFICIENT (3 TORRENTS)

Torrent	**Perf.**	dev.upload	Pop	Cont	Churn
RH	inbw.util	-0.46	-0.13	0.05	-0.12
RH	inbw.nsdev	0.49	0.20	-0.03	0.19
DE	inbw.util	-0.42	-0.02	0.10	-0.02
DE	inbw.nsdev	0.47	0.03	-0.10	0.00
GA	inbw.util	-0.36	-0.05	0.04	-0.05
GA	inbw.nsdev	0.47	0.14	-0.11	0.14

VII. CONCLUSION

In this paper, we examined a repeated claim that BitTorrent can provide high performance (*i.e.*, download rate) to participating peers in a torrent. We derived peer-level performance along with observed peer- and group-level properties among all peers in three different torrents. We showed that the distribution of performance among participating peers in the a torrent is roughly uniform. We also investigated the impact of various properties on the observed performance by individual peers and illustrated two points: *(i)* There is no clear relationship between peer-level performance and main peer- and group-level properties, *i.e.*, the relationship could significantly vary among peers, and *(ii)* average upload rate (*i.e.*, contribution) of individual peers has the highest correlation with its observed performance. This suggests that the tit-for-tat mechanism in BitTorrent is the primary factor that affects peer-level performance. These findings reveal that a common approach of using a few instrumented clients does not provide a representative view of BitTorrent behavior. Instead, a more global view must be considered in order to derive a reliable and general conclusion.

REFERENCES

[1] A. Bharambe, C. Herley, and V. Padmanabhan. Analyzing and Improving a BitTorrent Network's Performance Mechanisms. In *INFOCOM*, 2006.
[2] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. Measurements, Analysis, and Modeling of BitTorrent-like Systems. In *IMC*, 2005.
[3] M. Izal, G. Urvoy-Keller, E. W. Biersack, *et al.* Dissecting BitTorrent: Five Months in a Torrent's Lifetime. In *PAM*, 2004.
[4] T. Karagiannis, P. Rodriguez, and K. Papagiannaki. Should Internet Service Providers Fear Peer-Assisted Content Distribution? In *IMC*, 2005.
[5] A. Legout, G. Urvoy-Keller, and P. Michiardi. Rarest First and Choke Algorithms Are Enough. In *IMC*, 2006.
[6] J. Pouwelse, P. Garbacki, D. Epema, and H. Sips. The Bittorrent P2P File-sharing System: Measurements and Analysis. In *IPTPS*, 2005.
[7] D. Qiu and R. Srikant. Modeling and Performance Analysis of Bit Torrent-Like Peer-to-Peer Networks. In *SIGCOMM*, 2004.