

Unraveling BitTorrent’s File Unavailability: Measurements and Analysis

Sebastian Kaune
Technische Universität Darmstadt
kaune@kom.tu-darmstadt.de

Rubén Cuevas Rumín
Universidad Carlos III de Madrid
rcuevas@it.uc3m.es

Gareth Tyson
Lancaster University
g.tyson@comp.lancs.ac.uk

Andreas Mauthe
Lancaster University
mauthe@comp.lancs.ac.uk

Carmen Guerrero
Universidad Carlos III de Madrid
guerrero@it.uc3m.es

Ralf Steinmetz
Technische Universität Darmstadt
steinmetz@kom.tu-darmstadt.de

Abstract—BitTorrent suffers from one fundamental problem: the long-term availability of content. This occurs on a massive-scale with 38% of torrents becoming unavailable within the first month. In this paper we explore this problem by performing two large-scale measurement studies including 46K torrents and 29M users. The studies go significantly beyond any previous work by combining per-node, per-torrent and system-wide observations to ascertain the causes, characteristics and repercussions of file unavailability. The study confirms the conclusion from previous works that seeders have a significant impact on both performance and availability. However, we also present some crucial new findings: (i) the presence of seeders is *not* the sole factor involved in file availability, (ii) 23.5% of nodes that operate in seedless torrents *can* finish their downloads, and (iii) BitTorrent availability is discontinuous, operating in cycles of temporary unavailability.

I. INTRODUCTION

BitTorrent [1] has become a de-facto standard for scalable content distribution over the Internet. The reason for its success is its ability to efficiently leverage the uplink capacity of nodes whilst achieving high scalability during peak demands [2], [3]. This efficiency is largely attributable to BitTorrent’s tit-for-tat mechanism, which encourages users to share their resources whilst downloading files.

Despite the success of BitTorrent, it still suffers from a significant problem: the *long term availability of content*. More specifically, content that is distributed using BitTorrent often becomes unavailable after a relatively short period of time. For example, [4] found that the available lifespan of most torrents is between 30-300 hours whilst 10% of all users fail to successfully download their desired content.

A file can be considered unavailable if one or more of its data pieces are inaccessible to users wishing to download it. The most intuitive reason for this occurrence is that previously successful users in possession of the entire file (seeders) have left the system leaving only users that possess a subset of the file (leechers). Subsequently, unavailability occurs when this subset cannot collectively rebuild the complete file with their remaining pieces. Previous research (such as [5][4][6]) has promoted the importance of seeders in regard to availability and concluded that a seedless torrent is unable to reconstruct

the file. However, this conclusion is challenged by the observation that some torrents continue to effectively serve files despite lacking any seeders.

In this paper, we devote our attention to understanding and characterizing BitTorrent’s file unavailability problem. Specifically, we strive to discover the scale, causes and repercussions of the problem. To achieve this we have performed two large-scale measurement studies; the first investigates BitTorrent on a macroscopic level by periodically probing over 46K torrents to ascertain their high level characteristics, such as swarm size and seeder/leecher ratio. Whilst, the second study investigates BitTorrent on a microscopic level by contacting over 700,000 individual peers in 832 torrents to discover relevant properties such as their download rates and piece availability. To the best of the authors’ knowledge, this is the largest dataset in terms of size and collected information used to investigate file availability in BitTorrent. This allow us to extend previous works to obtain far more accurate results; through this we make a number of interesting findings,

- In 86% of cases, leechers are unable to reconstruct files in the absence of seeders. However, in 14% of cases, leechers can reconstruct the file without any seeders present. We therefore discover that seeders are not the sole factor involved in BitTorrent’s unavailability problem. Such torrents achieve this through the possession of large and stable populations as well as high aggregate download rates that enable leechers to quickly replicate rare chunks.
- In 64% of torrents, unavailability is *not* immutable and, instead, occurs in cyclic periods followed by reoccurring availability. This is due to old seeders returning to swarms where they previously participated in.
- The combination of the two previous observations results in 23.5% of users affected by a lack of seeders actually being able to complete their downloads.
- Users often become frustrated with unavailable torrents that exhibit poor download rates. We observe a chain reaction in which such users abort their downloads thereby exacerbating unavailability, resulting in further abortions.

- User arrival rates follow an exponentially increasing pattern; consequently, to ensure 99% availability, it would be necessary to increase average seeding times by a factor of 10.

The rest of the paper is structured as follows; Section 2 provides related work. Section 3 then details the problem and our measurement methodology. Following this, Section 4 characterises the causes and impact of unavailability. We discover a primary cause is a lack of seeders and therefore Section 5 investigates seedless states in BitTorrent. Section 6 then utilises the measurements to briefly analyse the potential solution-space to offer some guidelines for future work in addressing this problem. Finally, we conclude the paper in Section 7.

II. RELATED WORK

BitTorrent Measurements: BitTorrent measurement studies can be classified into two different groups. The first type uses log traces from trackers [2], [5], [4], [7] whereas the second type relies on crawling techniques to retrieve the information from the system [3], [8], [9], [6], [10]. The first type of measurements is less intrusive since they do not actively interfere with the system. However, they are often problematic to obtain since they require the agreement from content providers. The crawling techniques, on the other hand, can be divided into two categories. In its simplest form, a crawler exploits the BitTorrent protocol to periodically request the IP addresses of the clients participating in the torrent from the tracker [9]. This makes it possible to study the demographics and dynamics of the torrents under analysis. This is what we name *macroscopic crawling*. More sophisticated crawlers also contact the clients and retrieve detailed information such as the client ID and their piece bitmap. We name this *microscopic crawling*. Although the microscopic crawling gives more detailed information, it is noticeably less scalable and only allows a few thousand torrents to be studied in parallel [3], [8]. Each approach is effective for addressing particular needs; however, these have not yet been combined to investigate BitTorrent in a holistic way.

BitTorrent’s File Availability Analysis: There are only a few works investigating availability issues in BitTorrent systems [4], [11], [6]. Neglia et al. mainly study the tracker/DHT availability of 22,000 torrents obtained from two torrent indexing sites [11]. Guo et al. [4] extended this, to model the lifespan of torrents by analyzing a limited number of tracker traces from [11]; it was found that most torrents are short-lived because of an exponentially decreasing peer arrival rate. This model starts from the basis that content is unavailable when there are no seeders present in the swarm. This is, so far, an unverified hypothesis that is important to investigate. Similarly, Menasche et al. [6] also use this hypothesis to investigate the availability of seeders in 45,000 torrents obtained from the Mininova website, finding that 40% of swarms lack seeders for more than 15 days in the first month after the torrent’s birth. Finally, Qui et al. [12] analyses the effectiveness and stability of BitTorrent systems through a fluid model. However, their

results consider that peers arrive in a constant Poisson process which is a strong assumption given the measurement results presented in [3], [4] and also this paper. Thus, this model is only predictive for BitTorrent swarms that are in steady state. As shown in [4], this stable period is typically very short in reality.

III. PROBLEM BACKGROUND AND METHODOLOGY

A. Defining File Availability

To study and understand the availability of files in BitTorrent we first present a simple model. Let’s assume that we have a torrent T , formed by N nodes, managing the download of a file composed by P pieces. Thus, we can define the vector $V_i = [V_{i1}, V_{i2}, \dots, V_{iP}]$ that contains the information about the pieces stored by peer i : $V_{ij} = 1$ if node i has the piece j ; $V_{ij} = 0$ if node i does not have piece j . V_i is typically known as the *bitfield* of node i .

We define the *Percentage of Available Pieces* of torrent T at a time instant t as

$$U(T) = \frac{\sum_{j=1}^P OR(V_{ij})}{P}. \quad (1)$$

Where $OR(V_{ij})$ represents the logical *OR*-operation over the piece j across all the nodes in the torrent T .

B. The Circumstances of Unavailability

It is important to understand in which circumstances a file becomes unavailable, based on our definition. A file is considered unavailable if at least one of its pieces is not accessible within a swarm. This situation arises if there are no peers in the swarm that possess a given piece or, alternatively, if the peer(s) that possess the piece are inaccessible (e.g. due to firewalls, NAT or overlay graph disconnection). It is intuitive to consider the former as a far more likely circumstance (e.g. most BitTorrent clients implement techniques such as NAT traversal [13]. Moreover, they include neighbors discovery techniques such as the Peer Exchange Protocol (PEX) [14] and periodical tracker polling that prevents graph disconnection). Therefore, given this assumption, a file can be considered available if (i) there is at least one seeder or (ii) there is no seeder but the bitfields of the leechers collectively fit the condition $U(T) = 1$. Without detailed analysis, we can therefore currently state that:

- With an accessible seeder, a file *is* available,
- Without an accessible seeder, a file *may* be available.

This paper uses these two observations as a starting point to investigate unavailability in BitTorrent. In the following sections, we denote time periods in a torrent’s lifecycle in which no seeder is online as a *seedless state*. To this end, the file is *unavailable* if torrent T is in seedless state and $U(T) < 1$.

C. Measurement Methodology

To study the unavailability problem and specifically the seeders’ role in it, we have performed two large-scale measurement studies using microscopic and macroscopic crawling.

To the best of our knowledge, this paper is the first to combine both microscopic and macroscopic crawling techniques to better understand BitTorrent (specifically BitTorrent’s file availability).

Microscopic Crawling: To truly understand unavailability in BitTorrent, it is necessary to be able to view the microscopic characteristics of any given swarm, e.g. piece distribution or nodes’ download rates. Without this, one can only get a rough estimation of availability using metrics such as the number of seeders. The information regarding the behaviour of individual peers provides the necessary data to make new, more accurate findings. To gain this information we developed and deployed a distributed BitTorrent crawler that can investigate swarms on a microscopic level, using 20 nodes in the Emulab testbed [15].

The crawler operated from July 18, 2009 to July 29, 2009 (*micros-1*) and then again from August 19, 2009 to September 5, 2009 (*micros-2*). To discover all the online users in a torrent it periodically contacted the torrent’s tracker as well as using the Peer Exchange Protocol (PEX). From every peer, every 10 minutes it requested their piece bitmap to discover the real-time distribution of pieces. For the *micros-1* study, the crawler followed 255 torrents appearing on Mininova¹ after the first measurement hour; in these torrents, we observed 246,750 users. The *micros-2* dataset contains information from 577 torrents and 531,089 users.

Macroscopic Crawling: The microscopic measurements provide detailed insight into the distribution of pieces and download rates within the swarm, as well as between different peers. However, due to scalability issues it is difficult to perform such detailed measurements on a very large-scale (e.g. several thousand torrents). To complement these results, we therefore also implemented a higher level crawler that followed every torrent published on the Mininova website after December 09, 2008 for a period of 38 days. This crawler periodically requested, from multiple sites in Europe, tracker information regarding each torrent’s number of seeders and leechers alongside the members’ ip addresses (we were able to systematically collect 98% of all the ip addresses from within the swarms). This study allowed us to gain an extremely large number of measurements regarding details such as peer arrival patterns, seeder/leecher ratios and torrent sizes. This information can subsequently be correlated with our smaller-scale microscopic measurements to derive such things as the scale of seedless states and the causes for seedless states occurring. Our final macroscopic dataset consisted of reports from 46,227 torrents and 29,066,139 users.

IV. CHARACTERISING UNAVAILABILITY: CAUSES AND IMPACT

In this section, we first investigate the role that seeders play in file unavailability. Following this we study the exceptions and variations we discovered. Lastly, we then investigate the

¹At this time, Mininova was the largest BitTorrent Community based on the Alexa Ranking.




Fig. 1. Piece availability in torrents affected by seedless states.

real-time impact that a lack of seeders has on client performance and their subsequent reactions that can be observed.

A. Investigating the Role of Seeders in File Unavailability

It is intuitive to think that $U(T) < 1$ in a torrent without any seeder (that is, leechers are unable to reconstruct the file). However, this is, so far, an unverified assumption that must be investigated (and quantified). To ascertain this, we inspect the (i) nodes’ bitfield and (ii) nodes’ download rates in all the torrents of our microscopic traces affected by seedless states.

1) *Bitfield Analysis:* We have collected every nodes’ bitfields for all the torrents in our microscopic measurements as they have evolved over time. For each torrent we have computed $U(T)$ periodically every 10 minutes during any period a torrent is without any seeders (i.e. it is in a seedless state). This allows us to ascertain whether a full copy of the file exists in the torrent at any given time. Fig. 1 shows the cumulative distribution function (CDF) of $max(U(T))$ observed in the seedless state for each torrent that we studied. From this data, we can extract two pieces of information; first, in the majority of cases (86%) our hypothesis is confirmed and the contacted leechers are unable to collectively reconstruct the file once a seeder has left (i.e. $max(U(T)) < 1$). Clearly, this means that seeders *do* have a significant impact on the availability of files in BitTorrent. Importantly, however, we also find that a notable proportion of torrents (14%) actually remain available even without a seeder. Collectively, this makes up 24% of all leechers that operate in seedless swarms. This is a crucial finding that has not been observed before; it is therefore in contrast with previous measurement studies [6], [5] that consider all seedless torrents to be unavailable.

2) *Download Rate Analysis:* A limitation of the bitfield analysis is that not all nodes are accessible due to NATs. To address this, we also inspect the aggregate torrent download rates. Through this, we can infer that a file is unavailable when the download rate of all the peers participating in a specific torrent drops close to 0 KBps. From this we can derive that the node cannot find any new pieces to download.

To highlight our findings, we first inspect a representative

torrent from our microscopic trace², shown in Fig. 2. The figure shows the median instant download rate of the online leechers over time, sampled every 10 minutes. It also plots the number of seeders and leechers, as well as the number of copies of the least replicated piece. Note that when the number of seeders becomes 0, the torrent enters a seedless state.

The torrent can be observed to enter a seedless state after the middle of day 3, remaining in this state for roughly two days. When the final seed departs the download rate of the leechers drops to approximately 0-3 KBps after only a few minutes. This also coincides with the number of least replicated pieces dropping to zero. It can therefore be confidently inferred that the file is, indeed, unavailable during this period due to the departure of the last seeder.

Interestingly, it can also be seen that the torrent becomes available again during day 5. As the seeders return, the download rate increases and the file becomes available again. In contrast to past assumptions, it is therefore evident that unavailability is not continuous. This important phenomenon will be investigated further in Section V-C.




Fig. 2. Snapshot from a torrent in our microscopic trace.

The above analysis has inspected a representative torrent. To validate its widespread applicability we also look at the download rate degradation in all torrents. To achieve this, we have taken all the users that have been affected by a seedless

²We have observed the same behaviour in most of the torrents affected by seedless states.




Fig. 3. Interval download rates for nodes affected by the lack of seeders.

state and separated their downloading time into two periods: (i) periods in which they have suffered from a seedless state and (ii) periods in which they have not. Fig. 3 presents the download rate distribution for both periods. First, we can observe that the download rate in a non-seedless state is much higher than in a seedless state. 80-85% of the nodes experience an average download rate lower than 1 KBps when in a seedless torrent, indicating that the peers cannot locate any required pieces and the file is, indeed, unavailable. Second, however, we also observe that 15-20% of users, in fact, maintain a reasonable level of performance even without any seeders. This can be attributed to two reasons: (i) the aforementioned 14% of torrents are capable of reconstructing their file without a seeder at an average rate of 21.3 KBps; and (ii) newly joined peers can download the subset of available pieces at an effective rate. This can be observed in the representative torrent (cf. Fig. 2): between days 4 and 5 there is a peak in the number of leechers which results in a short peak in the download rate as newcomers download the available pieces.

B. Investigating the Causes of Swarm Resilience

The previous section has identified a notable percentage (14%) of torrents that can maintain availability even without any seeders; this represents 24% of all leechers that encounter seedless states.

To investigate this, we separate torrents into those that survive in the absence of seeders (*resilient torrents*) and those that do not (*susceptible torrents*). We then investigate quantitative properties of these two groups to ascertain how they differ at various points in their lifecycles. All of the identified metrics have been calculated for each group (across all member torrents) every 10 minutes using information from the microscopic traces and the tracker reports. These values have then been averaged together over each time period investigated.

Table I gives an overview of all metrics used in this analysis. We calculate these over two time periods: the beginning of the torrents' lifecycle and just before the last seeder goes offline. Although not included in the table, we also investigated the effects of file size and content type without ascertaining any correlation. Most metrics are straight-forward, however, two

require some explanation: Distribution Entropy ($E(T)$) and the Churn Factor (CF).

The $E(T)$ investigates the distribution of pieces within the swarm; this is to investigate whether torrents that can survive achieve a superior distribution of pieces. We therefore characterize the distribution entropy in a torrent T at a given time t by introducing the following *Entropy Index*:

$$E(T) = \frac{\left(\sum_{j=1}^P \sum_{i=1}^N V_{ij}\right)^2}{P \cdot \sum_{j=1}^P \left(\sum_{i=1}^N V_{ij}\right)^2} \quad (2)$$

Recall that N defines the number of nodes in the swarm, P is the number of pieces a file is composed of and V_i is the bitfield of node i . This index is similar to Jain’s Fairness Index [16] and achieves a value of 1 if all pieces are equally distributed among the peers.

The Churn Factor CF investigates whether torrents that can survive have more stable populations. This factor is defined by N_{disc}/N_{all} where N_{disc} is the number of users that have left the swarm during a given time period (t) and N_{all} is the total number of users observed during this same period. A factor of 0 indicates that no user disconnected within t ; by default $t = 10$ mins.

From the data in Table I, we can make the following important observations,

- *Torrent Popularity*: From the beginning, resilient torrents exhibit higher leecher population sizes. Larger torrents possess an increased probability of replicating rare pieces before the loss of seeders.
- *Low Churn Factor*: High churn in small torrents creates a greater risk of losing vital pieces; if this coincides with the loss of a seeder then it becomes impossible to recover these pieces again until a seeder returns. Resilient torrents have significantly lower churn factors than susceptible torrents.
- *Seeder/Leecher Ratio*: Resilient torrents exhibit a higher seeder/leecher ratio and, as a derivative of this, experience download rates that over twice as high as susceptible torrents. This superior performance is highly beneficial for the survival of piece replicas as it allows the quick duplication of rare pieces. Before seedless state occurring, resilient torrents therefore have many more replicas of the rarest piece when compared to susceptible torrents.

In summary, these results show that swarm resilience is a product of large, stable populations that can achieve higher download rates due to beneficial seeder/leecher ratios. The combination of these factors results in rarest piece replication rates that are over 5 times greater than their susceptible counterparts. This makes such swarms highly resilient to the loss of any seeders. Importantly, it also can be concluded that unavailability cannot be addressed by modifying any of BitTorrent’s algorithms (e.g. piece selection) but, instead, must be solved by incentivising users to modify their behaviour. This is exemplified by the lack of any correlation between resilience and distribution entropy.

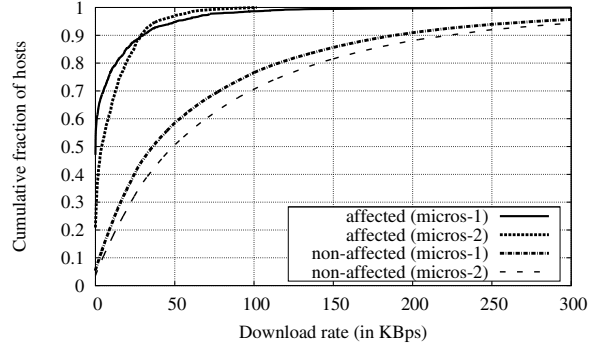


Fig. 4. Comparison of download performance between peers affected and not affected by seedless states.

C. Effects and Trends of Seed Departure

Despite a notable percentage of torrents surviving without seeders, it is evident that the loss of all seeders can often result in unavailability. This section now investigates the effects that this has on both individual users and wider system performance. Three stages can be identified which we now discuss.

The first repercussion of the loss of seeders in susceptible torrents is a significant and rapid drop in download rates. To extend the earlier analysis, we now compare the average download rate of users that suffer from a seedless state at some point during their download against the average download rate of users that always find content available. We first categorise users into two groups: affected vs. non-affected. The first group of users consists of leechers that are (at some point) affected by a lack of seeders. The non-affected users, on the other hand, have at least one seeder available during their entire download. Fig. 4 gives the download rate distribution of both user groups as obtained from the two microscopic crawlings. Whereas the median download rate for the non-affected users is 36 KBps in *micros-1* and 48 KBps in *micros-2*, the performance for peers attempting to download unavailable content is only 0.06 KBps and 3.8 KBps, respectively.

The second observable stage is a direct derivative of the decrease in download performance. Specifically, we observe a large increase in download abortions. To study this we examine the session times in our microscopic traces. We observe that 89% of users affected by file unavailability (i.e. participating in *susceptible* torrents) abort their downloads due to the bad performance. Sadly, this is an unnecessary action as we have found that seeders often return, making files available again. In contrast to these results, users operating in resilient torrents only have an abortion rate of 34.47%. Although this seems initially high, we also find that many users operating in other torrents that do not suffer from unavailability also abort their downloads. On closer inspection, these ‘unnecessary’ abortions occur in torrents that have particularly low download rates that are under a third of the average.

The third stage in this process is the worrying emergence of a chain reaction. We find that as the number of abortions

Metric	Time before seedless state				Time after torrent's birth			
	1 hour		6 hours		6 hours		24 hours	
	Resilient	Susceptible	Resilient	Susceptible	Resilient	Susceptible	Resilient	Susceptible
Swarm speed (in KBps)	58.83	23.88	68.58	24.99	95.72	53.50	62.99	44.82
Seeder/Leecher ratio	0.15	0.03	0.14	0.04	0.19	0.19	0.32	0.43
Firewalled/NATed peers (in %)	70.86	61.09	62.30	60.67	51.30	56.05	54.94	58.44
Distribution Entropy $E(T)$	0.93	0.94	0.92	0.93	0.91	0.90	0.92	0.91
Least repl. piece (# of copies)	9.21	1.61	8.38	2.82	15.34	14.21	21.33	23.55
Churn factor CF	0.03	0.21	0.08	0.15	0.07	0.11	0.08	0.07
Online leechers	281.34	111.61	250.48	101.55	134.05	82.71	163.12	89.55
Online seeders	5.28	1.51	6.11	1.80	12.05	11.25	18.17	21.34

TABLE I
CHARACTERISTICS OF RESILIENT TORRENTS (THOSE THAT MAINTAIN AVAILABILITY IN SEEDLESS STATE) AND SUSCEPTIBLE TORRENTS (THOSE THAT CANNOT RECONSTRUCT THE FILE).

increase, the number of available chunks decrease. This results in an exacerbation of the torrent's unavailability and a further drop in download rates for those trying to access the remaining chunks. As other users witness this trend, they too abort their downloads. This process results in fewer users becoming seeders and therefore greater unavailability and more abortions. Frequently, the above two repercussions of unavailability and the creation of this chain reaction often spells the end for a torrent.

From these findings we derive that users are highly sensitive to their perceived instant quality of service and therefore any solutions must maintain an acceptable download rate whilst also improving file availability.

V. CHARACTERISING SEEDLESS STATES

The previous section has validated and quantified the importance of seeders in regard to file availability in BitTorrent and discussed under which circumstances a file of a seedless torrent becomes unavailable. It has been found that in the majority of torrents (86%), the loss of all seeders results in unavailability. In this section we therefore investigate the behaviour of seeders and characterise the nature of seedless states using our large scale dataset. In Section 5.1 we look at the frequency at which seedless states occur in BitTorrent and in Section 5.2 we investigate what causes the system to enter in such states. In Section 5.3 we address why torrents revive after extended periods of unavailability.

A. How Prevalent are Seedless States?

To quantify how prevalent seedless states are in BitTorrent, we ask the following question: how *many torrents* and to *what extent* are torrents affected by seedless states? To answer this, we use the logs from our macroscopic trace that give us a large-scale view on the system comprising of 46k torrents.

The measurements show that more than 38% of torrents (17,568 out of 46,227) lose their seeders within the first month, out of which 72% lack seeders after only 5 days. Similarly, we find that more than 45% of the torrents suffer from a lack of seeders for half of their monitoring time. To exemplify the scale of this, in 50% of the torrents observed for periods longer than 30 days, no seeder was available for more than 16 days.




Fig. 5. Illustration of a seedless state.

Finally, in our study, more than 9.68 million users (33% of all users seen) participated in torrents with highly unavailable seeders suggesting that this is not only a long tail problem. Out of these users, more than 1.59 million were directly affected by seedless states.

B. Why do Seedless States Occur?

Since seedless states are highly prevalent in real swarms, an intuitive question is: *why do they occur in the wild?* In this section, we first identify and then further investigate the influencing factors responsible for triggering seedless states.

1) *Identifying Influencing Factors:* There are two main factors that directly influence the existence of seedless states: (i) the session time of seeders and (ii) the inter-arrival rate of the users. To illustrate the influencing factors, we use a simple example shown in Fig. 5. In this figure, each horizontal line represents the lifetime of a user; these users can either be in a leecher state (thin lines) or a seeding state (thick lines).

It seems straightforward that the longer a seeder serves content, the more leechers are able to finish their downloads. Unfortunately, (as demonstrated later on) the seeding time is typically quite short contributing significantly to the frequency and length of seedless states.

Let's now assume that user n is the last available seeder in our example torrent and none of the previous seeders return to the torrent. In this case, a seedless state occurs when the time

required for leechers to download the file exceeds the online time of the last seeder. For example, Fig. 5 shows that after the last available seeder leaves the swarm at time t_3 , none of the remaining leechers were able to finish the download. If we focus on the n -th node and its subsequent successor in the torrent ($n+1$ -th), the inter-arrival time between both users is given by $\tau_{n+1}(=t_2-t_1)$ whereas the seeding time of node n is given by μ_n . Assume that both users n and $n+1$ download a file of size F_s with rate D_n and D_{n+1} respectively. Thus, the swarm enters a seedless state when Eq. 3 is fulfilled.

$$\frac{F_s}{D_n} + \mu_n < \tau_{n+1} + \frac{F_s}{D_{n+1}} \quad (3)$$

To simplify the analysis, we assume that $D_n = D_{n+1}$ ³. In this case, the seedless state is reached if the inter-arrival time is larger than the seeding time.

To summarise, seeding times as well as inter-arrival times play an important role in the generation of seedless states and subsequently in the long-term availability of content. Since both parameters are not directly correlated, we individually analyse both of them in the following.

2) *Arrival Behaviour of Users*: The first behavioural characteristic that is paramount to seedless state generation is the inter-arrival times of users. In this regard, intuitive questions are: (i) what inter-arrival times do we expect in reality and (ii) how do inter-arrival times evolve over time?

By analysing a few hundred torrents in a small community, previous work [4] has shown that user inter-arrival times are exponentially increasing. Our goal is to generalize this finding for 'open' communities such as Mininova.org that are orders of magnitude larger. For our analysis, we use similar techniques as applied in [4]. We consider all torrents in our macroscopic trace. We use linear regression to fit the logarithm of the complementary⁴ of the number of node arrivals of each torrent along time. Let X_t denote the complementary number of node arrivals at time epoch t and Y_t be the fitting result. We define the relative deviation of the actual node arrivals over an ideally exponentially increasing function by $\frac{\log X_t - \log Y_t}{\log X_t}$. Thus, a relative deviation of 0% indicates that both curves overlap. Fig. 6 shows the deviation for each torrent of our macroscopic trace. The x-axis depicts the torrents ordered by ascending population size while the y-axis shows the relative deviation. For most of the torrents, the relative deviation is less than 10% whereas the deviation tends to decrease with increasing torrent popularity. Altogether, the average relative deviation of all torrents is 4.8%. Therefore, we conclude that the inter-arrival time of the nodes exponentially increases with time.

Notably, we observed especially high inter-arrival times in torrents affected by seedless states; this is in line with our analysis in the previous section. For instance, Fig. 7 plots the

³Our microscopic measurements show that the download rate of users that finish downloads (D_n in the example) is higher than the download rate of those that do not (D_{n+1}) validating our assumption.

⁴We use the complementary number of node arrivals to avoid domains in which the logarithm is undefined, e.g., epochs with no peer arrivals.




Fig. 6. Deviation from linear regression.

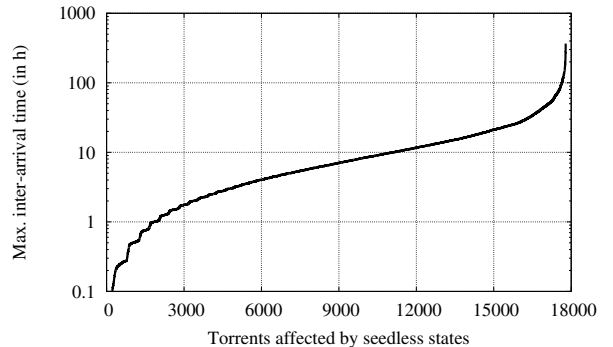


Fig. 7. Maximum inter-arrival times of torrents with highly unavailable seeders.

maximum inter-arrival time observed in these torrents with unavailable seeders. More than 45% of the torrents exhibit inter-arrival times far beyond 10 hours.

3) *Seeding Times of Users*: The second behavioural characteristic that is paramount to the creation of seedless states is the seeding time of a node, i.e. how long seeds stay online for. As already shown in our example torrent (cf. Fig. 5), to maintain file availability it is necessary for seeders to remain online for long enough for new seeds to be generated. Fig. 8 shows the cumulative distribution of the seeding times of the nodes obtained from the two microscopic measurements. It can be seen that seeding times are generally short-lasting with 75% of the seeders staying online for less than 4 hours. When this data is compared to the inter-arrival time of users it can be identified that the current seeding times in BitTorrent are not sufficient to avoid seedless states, thus preventing to achieve long-term file availability in BitTorrent.

C. How long are Seedless States?

The representative snapshot presented in Fig. 2 has highlighted that torrents can become available again after a extended periods of unavailability. In this section we validate (using both the macroscopic and microscopic datasets) that file unavailability is, in fact, discontinuous with reoccurring periods of temporary availability. Through our measurement studies we can state that this occurs because seeders often return to swarms that they have previously participated in. This allows the 11% of users that choose to remain online




Fig. 8. Seeding time distribution.




Fig. 9. Length of cycles of temporary availability and unavailability periods.

during periods of unavailability (i.e. in *susceptible* torrents) to eventually complete their downloads. Alongside the existence of *resilient* torrents, this means that 23.5% of all leechers affected by seedless states can actually still gain access to the file.

The reoccurrence of seeders happens in over 64% of torrents that suffer from seedless states in our macroscopic study. To investigate this, Fig. 9 shows the CDFs of both the duration of seedless states as well as the duration of the subsequent periods in which content becomes available again, computed over all torrents exhibiting this phenomenon. Note that the x-axis is in log scale. It can be observed that seedless periods are typically long-lasting with an average of 43.19 hours whereas the subsequent availability periods only last 12.56 hours on average.

The primary reason for the (seemingly) altruistic return of seeders is likely to be the default settings of many BitTorrent clients (e.g. Vuze, μ Torrent) that automatically rejoin torrents at their start-up even after a user has completed a file download. Unfortunately, BitTorrent users do not have permanent identifiers and thus we cannot make quantitative statements on exactly how many unique seeders rejoin a swarm and over what time period. However, the length of the seedless periods as depicted in Fig. 9 offers a lower bound for the inter-seeding time distribution of such users. This is obviously very coarse grained and therefore we expect the inter-seeding times actually to be higher.

The reoccurrence of seeders is obviously in contrast to previous work that has assumed unavailability is continuous

and immutable. To investigate the impact that this finding has on previous work, we briefly look at the relative deviation that the assumption has when compared against our dataset. We define the relative deviation as,

$$\text{relative deviation} = \frac{\text{measured avail. time} - \text{assumed avail. time}}{\text{assumed avail. time}} \quad (4)$$

We find that in approximately 35 % of the torrents in our macroscopic dataset, the assumption works well. In these torrents, we did not observe any temporary availability period after the torrent first enters a seedless state. However, in 50% of the torrents, the content is actually available for at least twice as that assumed when considering immutable unavailability.

VI. DISCUSSION ON IMPROVING FILE AVAILABILITY

The previous sections have investigated file unavailability in BitTorrent; this section now briefly explores some existing and intuitive solutions to gain an understanding of the best future approach to take. Specifically, we show that bilateral incentives are insufficient when considering current peer-to-peer file sharing workloads. This is because they provide no incentive for those users that have already completed their download.

The most recent work [6] tries to tackle the file unavailability problem with file bundling, which involves forcibly lengthening the online times of users by making them download multiple files to gain access to one. Although this can work in the short-term, the measurement results show that it only circumvents the problem - it does not solve it. This is because it is only possible to lengthen online times by a relatively limited period (e.g. 10 minutes, 1 hour, etc.); our studies have shown that inter-arrival times between two consecutive peers can quickly exceed 24 hours (c.f. Section V-B2), leaving file bundling totally ineffective. In fact, this is the same for any approach that attempts to lengthen online times: the measurement data indicates that, on average, seeding times would have to be increased by a factor of 10 to gain 99% availability! The cause for previous confidence in this technique is the incorrect assumption of a constant Poisson arrival rate; although this produces positive results, our measurements show that it is entirely unrepresentative of the real-world (as confirmed by [3], [5]).

These conclusions lead us to believe that a single-torrent approach is not appropriate for addressing file unavailability. The alternative to this is therefore following a *cross-torrent* approach which distributes incentives over multiple torrents. The feasibility of this is confirmed by our traces which show that 36 hours after a torrent's birth, there is usually at least one user online (in another torrent) that could also act as a replica for content it has previously downloaded. These users are therefore untapped resources that have the potential to greatly increase both availability and performance on a per-torrent basis. A key line of future work in this area is therefore understanding how these users can be incentivised to



Fig. 10. Cross-torrent bartering (left) and cross-torrent tit-for-tat (right).

contribute their resources to their past torrents. In the current tit-for-tat model this is actively discouraged because the need to divert upload resources from a user's current torrent would disincentivise cross-torrent collaboration. This is because it would decrease the probability of a user being unchoked for its own content download. Briefly, there exists two intuitive cross-torrent solutions that could be investigated:

- *Cross-torrent bartering*: Assume that user A has previously download torrent X (fully or partially) and is a leecher in torrent Y. User B, on the other hand, has obtained torrent Y earlier and is now a leecher in torrent X. Both A and B could mutually exchange chunks while still conforming to BitTorrent's tit-for-tat strategy (cf. Fig. 10).
- *Cross-torrent tit-for-tat*: Traditional tit-for-tat and cross-torrent bartering are real-time incentives. That is, contributions and rewards are instant. An alternative is that each peer remembers all contributions and reciprocates across time and swarms by using persistent identities. For instance, as depicted in Fig. 10, peer A contributes to B in torrent X because in a future swarm Y, B will recognize A and reciprocate.

To provide a rough idea of the potential of these two approaches, our macroscopic measurement data can be inspected. The key requirement for cross-torrent bartering to success is for peers to frequently encounter other users with synchronous interests. Unfortunately, however, we find that this assumption holds for only 0.1% of cases. In contrast, cross-torrent tit-for-tat requires peers to often have repeat interactions over time. Our measurement study similarly disproves this requirement, however, as we observed that 81% of peers that interact do so only once.

These results suggest that the use of cross-torrent bilateral incentives are just as ineffective as single-torrent bilateral incentives at solving file unavailability. This has been found to be true even when incentives are detached from time. Consequently, we believe future work in this field should focus on cross-torrent incentives that are not only detached from time but also from individual peers. In such a system, peers would be able to claim back contributions at any time from any peer, i.e. reciprocation would be indirect. Early work in this field includes [5], [10], however, a number of challenges still remain. These include managing overhead and complexity, addressing potential security threats, handling persistent identities and realistically evaluating its potential.

VII. CONCLUSIONS

This paper has investigated BitTorrent's unavailability problem in the wild and explored the feasibility of the potential

solution-space. To achieve this, two large-scale measurements studies were performed to ascertain the characteristics, causes and repercussions of file unavailability in BitTorrent. Based on this, we made a number of interesting findings that offer the most accurate study of file availability in BitTorrent so far. Most notably, it was found that (i) a lack of seeders often results in unavailability but *not* always, (ii) the churn level, the fast replication of rare chunks and the population size largely defines a swarm's ability to survive without a seeder, (iii) unavailability usually occurs in cyclic periods with intermittent availability, and (iv) unavailability often results in a chain effect that leads to future download failures. Lastly, these results have been used to briefly explore avenues of future work to reveal that the most promising line of investigation is the use of cross-torrent indirection reciprocation schemes.

ACKNOWLEDGMENT

This work has partially been supported by the German Federal Ministry of Education and Research (BMBF) in the project Premium Services (O1IA08003A).

REFERENCES

- [1] B. Cohen, "Incentives build robustness in bittorrent," in *1st Workshop on Economics of Peer-to-Peer Systems*, 2003.
- [2] M. Izal, G. Uroy-Keller, E. Biersack, P. A. Felber, A. A. Hamra, and L. Garces-Erice, "Dissecting bittorrent: Five months in torrent's lifetime," in *PAM*, 2004.
- [3] J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips, "The bittorrent p2p file-sharing system: Measurements and analysis," in *IPADS*, 2005.
- [4] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurements, analysis, and modeling of bittorrent-like systems," in *IMC*, 2005.
- [5] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurements, analysis, and modeling of bittorrent-like systems," *IEEE Journal on Selected Areas in Communications*, 2007.
- [6] D. S. Menasche, A. A. A. Rocha, B. Li, D. Towsley, and A. Venkataramani, "Content availability and bundling in swarming systems," in *CoNEXT*, 2009.
- [7] A. Bellissimo, B. N. Levine, and P. Shenoy, "Exploring the use of bittorrent as the basis for a large trace repository," Tech. Rep. 04-41, UMASS Amherst, Tech. Rep., 2004.
- [8] G. Sigamos, J. Pujol, and P. Rodriguez, "Monitoring the bittorrent monitors: A bird's eye view," in *PAM*, 2009.
- [9] C. Zhang, P. Dughel, D. Wu, and K. W. Ross, "Unraveling the bittorrent ecosystem," in *Unpublished*, 2009.
- [10] M. Piatek, T. Isdal, A. Krishnamurthy, and T. Anderson, "One hop reputations for peer to peer file sharing workloads," in *NSDI*, 2008.
- [11] G. Neglia, G. Reina, H. Zhang, D. Towsley, A. Venkataramani, and J. Danaher, "Availability in bittorrent systems," in *INFOCOM*, 2007.
- [12] D. Qiu and R. Srikant, "Modeling and performance analysis of bittorrent-like peer-to-peer networks," in *SIGCOMM*, 2004.
- [13] J. J.-D. Mol, J. A. Pouwelse, D. H. J. Epema, and H. J. Sips, "Free-riding, fairness, and firewalls in p2p file-sharing," in *P2P*, 2008.
- [14] Azureus, <http://sourceforge.net/projects/azureus>.
- [15] Emulab – Network Emulation Testbed, <https://www.emulab.net>.
- [16] R. K. Jain, D.-M. W. Chiu, and W. R. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," TR-301, Digital Equipment Corporation, Tech. Rep., 1984.