

Content Availability and Bundling in Swarming Systems

Daniel S. Menasche
University of Massachusetts,
Amherst, MA, USA
sadoc@cs.umass.edu

Antonio A. A. Rocha
Federal University of Rio de
Janeiro, RJ, Brazil
arocha@land.ufrj.br

Bin Li
Tsinghua University,
Beijing, China
libinhere@gmail.com

Don Towsley
University of Massachusetts,
Amherst, MA, USA
towsley@cs.umass.edu

Arun Venkataramani
University of Massachusetts,
Amherst, MA, USA
arun@cs.umass.edu

ABSTRACT

BitTorrent, the immensely popular file swarming system, suffers a fundamental problem: content unavailability. Although swarming scales well to tolerate flash crowds for popular content, it is less useful for unpopular content as peers arriving after the initial rush find the content unavailable.

Our primary contribution is a model to quantify content availability in swarming systems. We use the model to analyze the availability and the performance implications of bundling, a strategy commonly adopted by many BitTorrent publishers today. We find that even a limited amount of bundling exponentially reduces content unavailability. Surprisingly, for swarms with highly unavailable publishers, the availability gain of bundling can result in a net improvement in download time, i.e., peers obtain more content in less time. We empirically confirm the model's conclusions through experiments on PlanetLab using the mainline BitTorrent client.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Modeling techniques

General Terms

Measurement; Performance; Reliability; Theory

1. INTRODUCTION

Despite the tremendous success of BitTorrent (estimated to account for 30–50% of all Internet traffic today), it suffers from a fundamental problem: availability. Although peer-to-peer swarming in BitTorrent scales impressively to tolerate massive flash crowds for popular content, swarming does little to disseminate unpopular content as their availability is limited by the presence of a seed or publisher. The

extent of publisher unavailability is severe, e.g., our measurement study shows that 40% of the swarms have no publishers available more than 50% of the time.

To appreciate the availability problem, consider a swarm for an episode of a popular TV show. When a publisher first posts the episode, a flash crowd of peers joins the swarm to download it. The original publisher goes offline at some point, but peers may continue to obtain the content from other peers while the swarm is active. If a peer arrives after the initial popularity wave, when the population of the swarm has dwindled down to near-zero, it finds the content unavailable and must wait until a publisher reappears.

Our primary contribution is a mathematical model to study content availability in swarming systems such as BitTorrent. We use an $M/G/\infty$ queue to model the self-scaling property of BitTorrent swarms, i.e., more peers bring in more capacity to the system. The key insight is to model uninterrupted intervals during which the content is available as *busy periods* of that queue. The busy period increases exponentially with the arrival rate of peers and publishers and the time spent by peers and publishers in the swarm.

Our model enables us to analyze the impact of *bundling*, a common strategy adopted by BitTorrent publishers wherein, instead of disseminating individual files via isolated swarms, a publisher packages a number of related files and disseminates it via a single larger swarm. To appreciate why bundling improves content availability, consider a bundle of K files. Assume that the popularity of the bundle is roughly K times the popularity of an individual file as a peer requesting any file requests the entire bundle. The size of the bundle is roughly K times the size of an individual file. Our model suggests that the busy period of the bundled swarm is a factor $e^{\Theta(K^2)}$ larger than that of an individual swarm. Indeed, if busy periods supported by peers alone last until a publisher reappears, the content will be available throughout.

Surprisingly, in some cases, the improved availability can reduce the download time experienced by peers, i.e., peers download more content in less time. The *download time* of peers in the system consists of the *waiting time* spent while content is unavailable and the *service time* spent in actively downloading content. If the reduction in waiting time due to bundling is greater than the corresponding increase in service time, the download time decreases. We validate this conclusion in Section 4 through large-scale controlled experiments using the mainline BitTorrent client over Planetlab.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT'09, December 1–4, 2009, Rome, Italy.

Copyright 2009 ACM 978-1-60558-636-6/09/12 ...\$10.00.

Our experiments also show that the conclusions of our model qualitatively hold even with realistic arrival patterns, peer upload capacities, and heterogeneous popularities.

In summary, we make the following contributions.

Measurement: We present a large-scale measurement study of real BitTorrent swarms that shows that (1) content availability is a serious problem due to publisher unavailability, (2) bundling of content is a widely prevalent, and (3) bundled content is more available than unbundled content.

Availability model: We present a novel queuing-theoretic model to analyze content availability and download times in BitTorrent-like swarming systems. To our knowledge, this is the first model that relates content availability to arrivals and departures of peers as well as publishers.

Implications of bundling: We use the model to analyze the implications of bundling, a widely prevalent yet little studied phenomenon, and show that (1) bundling improves availability, and (2) bundling can reduce download times for unpopular content when publishers are highly unavailable.

Experimental validation: We validate the model using large-scale controlled experiments with the mainline BitTorrent client on PlanetLab showing that the model accurately predicts download times in swarms with intermittently available publishers for both bundled and individual content.

2. MEASURING CONTENT AVAILABILITY AND BUNDLING IN BITTORRENT

In this section, we present a large-scale measurement study of BitTorrent that shows that 1) content unavailability is a serious problem in BitTorrent today, and 2) bundling of content is widely prevalent and bundled contents shows greater availability. We begin with a brief overview of how swarming in BitTorrent works and why content becomes unavailable.

2.1 Why unavailability?

A swarm consists of a set of peers concurrently sharing (downloading or uploading) content (a file or a bundle of files) of common interest with the help of a coordinating tracker. Content is divided into blocks and peers obtain meta-data about constituent blocks as well as identities of other peers in the swarm from the tracker. A peer exchanges blocks with other peers using a tit-for-tat incentive strategy until it completes its download. Peers that have not yet completed their downloads are called *leechers* while peers that possess all blocks in the content are called *seeds*.

Content is available if either at least one seed is present or sufficiently many active leechers are present so as to collectively make all constituent blocks available. Seeds may become unavailable in practice due to several reasons. Publishing sites serving a large number of files may take down seeds after the initial popularity wave subsides in order to reduce bandwidth costs. A seed may also be an average user publishing home-generated content that can not afford to stay online all the time. Seeds illegally uploading copyrighted material often disappear quickly for obvious reasons. Even for legitimate content, maintaining highly available seeds entails administrative effort and cost, which runs counter to the goals of content publishers that value BitTorrent as a cheap alternative to a client-server approach.

Throughout this section, we measure content availability by equating it with seed availability. In the next section, we

model content availability resulting both from seeds as well as from leechers alone. In the rest of this paper, we use the terms *publishers* and *peers* interchangeably with *seeds* and *leechers* respectively.

2.2 Measuring unavailability

How available is content in BitTorrent swarms today? To answer this question, we conducted a seven-month long measurement study of BitTorrent swarms as follows. We developed and deployed BitTorrent monitoring agents at 300 nodes on Planetlab from August 3, 2008 to March 6, 2009. Once every hour, a host at the University of Massachusetts Amherst receives an RSS feed advertised by GoogleReader of recently created torrent URLs from Mininova (a large torrent hosting site), and sends each URL to a subset of the monitoring agents on Planetlab. The agents fetch the torrent metadata by joining the swarm and begin to monitor its peers. Our agents leverage the Peer Exchange (PEX) protocol extension, that enables it to discover new neighbors from other peers in addition to the tracker. To avoid copyright issues, our agents collect information only about the control plane without actually uploading or downloading content, which suffices for our purposes as we equate content availability with seed availability.

To distinguish seeds from leechers, our agents record the bitmaps received from connected peers. The bitmaps are part of the BitTorrent protocol and a peer uses them to convey the blocks it possesses to its neighbors. Each entry in the trace collected by the agents consists of a swarm identifier, a peer identifier (IP address and port number) and its bitmap recorded roughly periodically for each discovered peer in the swarm. Our traces consist of more than 14 million distinct IP addresses and 66K distinct swarms.

Figure 1 shows the distribution of seed availability for the monitored swarms. The solid curve shows the availability in the first month after the creation of the swarm, when we expect the content to be more popular. The extent of publisher unavailability is severe: less than 35% of the swarms had at least one seed available all the time. The availability of swarms over the entire duration of the measurement is even lower as shown by the dotted curve: almost 80% of the swarms are unavailable 80% of the time.




Figure 1: CDF of seed availability in 45,693 swarms each monitored for at least one month.

2.3 Content bundling

Bundling of content is a common practice in BitTorrent today. In this section, we study the extent of bundling and its impact on availability. The trace used in this section is a snapshot of BitTorrent swarms taken on May 6, 2009. For each of the 1,087,933 swarms in this snapshot, we record its content category (e.g., movies, TV, books etc.), names and sizes of constituent files, creation date, and instantaneous number of seeds and leechers. Note that we could afford to monitor many more swarms in this dataset than the previous one as we did not have to measure details of peer arrivals and departures inside each swarm.

2.3.1 Extent of bundling

We analyze the extent of bundling in three of the nine categories present in Mininova, namely, music, TV shows and books. These three categories together account for 45.98% of the swarms and 31.93% of the peers in the system. We chose these three categories because it is easier to automatically detect bundling by checking for the presence of multiple files with known extensions (e.g., `.mp3` for songs, `.mpg` for TV shows and `.pdf` for books). Detecting bundling is non-trivial in some categories, e.g., a DVD for a single movie is often organized as a collection of video files that are never distributed individually, making it difficult to check for the presence of multiple movies without manual inspection.

Among music swarms, albums are common. We classify a music swarm as a bundle if it has two or more files with common audio file extensions such as `.mp3`, `.mid` and `.wav`, which results in 193,491 of the 267,117 monitored swarms being classified as bundles.

Among TV show swarms, many bundles consist of sets of episodes in a season. We classify swarms that have two or more files with common video file extensions such as `.mpg` and `.avi` as bundles, which results in 25,990 of the 164,930 monitored swarms being classified as bundles.

Among book swarms, we observe that collections, i.e., torrents containing the keyword “collection” in their titles, usually consist of a bundle of contents connected by a broad theme, e.g., the “Ultimate Math Collection (1)” of size 5.81 GB has 642 books. We classified 841 of the 66,387 monitored swarms as collections. Classifying swarms that contain 2 or more files with common file extensions such as `.pdf` and `.djvu` as bundles results in an additional 6,270 bundles.

2.3.2 Bundled content is more available

In this section, we present evidence suggesting that bundling is correlated with higher availability. We first consider book swarms. We find that 62% of all book swarms had no seed available on May 6, 2009, whereas that number drops to 36% if we consider only collections. Furthermore, the average number of downloads for a typical book swarm is 2,578, whereas for collections it is 4,216.

One reason for higher seed availability may be that content publishers are intrinsically more willing to support seeds for bundled content. The higher number of downloads for bundled content may be either because of higher demand for bundled content (as any peer seeking any of the constituent files may opt to download the bundle), or because of higher availability, or both. Higher seed availability in turn may in part be because of the increased number of downloads as some peers may choose to altruistically disseminate the content further. Although it is difficult to discern cause and

Figure 2: Illustration of busy and idle periods.

effect in our measurement data, our analytic model in the next section quantifies how the higher demand and higher seed availability for bundled content *produce* improved content availability.

We next analyze our traces more closely for content that is available both in isolation and as part of a larger bundle. We observe that among the unavailable collections, some of them were subsets of bigger collections, e.g., the 23 swarms consisting of collections of Garfield comics from 1978 to 2000 had no seeds. However, each of these collections can be found in a single super-collection aggregating all Garfield comics. The super-collection had seven seeds. After a manual inspection of all 841 book collections, we concluded that 210 had no seeds and were not subsets of other collections, which results in $210/841 = 25\%$ unavailability for content disseminated through collections (compared to 62% above for a typical swarm).

As another example, we consider swarms for the popular TV show “Friends”. There were a total of 52 swarms associated with this show. Among them, 23 had one or more seeds available, and the remaining 29 had no seeds. The 23 available swarms consisted of 21 bundles (and 2 single episodes), whereas the 29 unavailable swarms consisted of only 7 bundles. These observations suggest a strong correlation between bundling and higher availability. The next section presents an analytic model that quantifies the causal relationship between the two.

3. MODEL

In this section, we develop a model for content availability in BitTorrent. The key insight underlying the model is to view BitTorrent as a *coverage process* or equivalently an $M/G/\infty$ queuing system. The model shows that 1) bundling improves availability, and 2) for swarms with highly unavailable publishers, the availability benefit of bundling more than offsets the increased time to actively download more content, resulting in a net decrease in user-perceived download times.

3.1 Model overview

Figure 2 illustrates how content availability in BitTorrent depends upon the arrivals and departures of publishers and peers. Each horizontal line segment represents the time interval during which a peer (represented using thin lines) or

