

Modeling Content Availability in Peer-to-Peer Swarming Systems

Daniel Sadoc Menasche
University of Massachusetts
sadoc@cs.umass.edu

Antonio A. Aragao Rocha
Federal Univ. of Rio de Janeiro
arocha@land.ufrj.br

Bin Li
Tsinghua University
libinhere@gmail.com

Don Towsley
University of Massachusetts
towsley@cs.umass.edu

Arun Venkataramani
University of Massachusetts
arun@cs.umass.edu

1. INTRODUCTION

BitTorrent, the immensely popular file swarming system, has a fundamental problem: availability. Swarming systems are peer-to-peer systems where users simultaneously interested in the same content cooperate with each other. Although swarming in BitTorrent scales impressively to tolerate massive flash crowds for popular content, it does little to disseminate unpopular content as they are limited by the availability 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.

In this paper we develop a model to quantify content availability in BitTorrent. An M/G/ ∞ queueing system is used to capture BitTorrent's scale-free property, i.e., more peers bring in more capacity to the system. The *busy period* of the queue captures an uninterrupted duration for which the content is available. The busy period increases exponentially with the arrival rate of customers as well as the time spent by each customer in the system. Armed with this model, we ask how content availability can be improved.

Our model suggests an intriguing answer, namely, *bundling builds availability in BitTorrent*. Bundling is a growing trend today among BitTorrent publishers; 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. 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, so each peer stays K times as long in the system. 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. If the busy period lasts until the publisher reappears, the content will be available throughout.

In related work, [2] modeled the *performance* of peer-to-peer systems. In contrast, we present a model for *availability*. [3] points out that peers who want a file may have to wait indefinitely to obtain certain chunks that are not available, a problem referred to as the blocked leecher problem (BLP). Here we show that bundling is one solution to mitigate this availability problem.

2. THE MODEL

Figure 1(a) illustrates how content availability in BitTorrent depends upon the arrival and departure of publish-

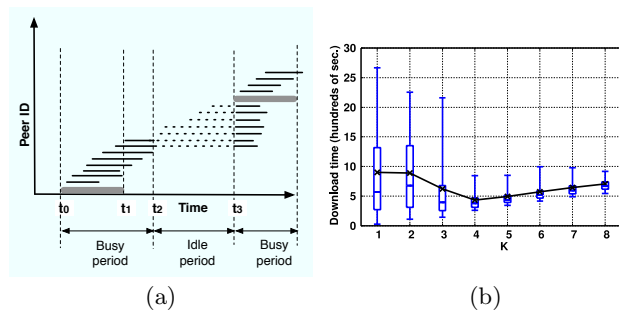


Figure 1: (a) Illustration of busy and idle periods. (b) Search for the optimal bundle size (PlanetLab).

ers and peers. Each horizontal line segment represents the length of time for which a peer (represented using thin lines) or a publisher (represented using thick lines) stays online. A swarm is initiated by the arrival of a publisher, which also marks the start of the first *busy period*. The swarm's lifetime is divided into alternating *busy* and *idle periods*. Content is available during busy periods and unavailable during idle periods. If a publisher is always online, the first busy period lasts forever and content remains always available.

A busy period ends when the following two conditions are satisfied: 1) there are no publishers online, and 2) the *coverage*, i.e., the number of currently online peers, drops below a fixed small threshold. For example, Figure 1(a) shows that after the publisher leaves at time t_1 , the busy period continues until time t_2 when the number of peers drops below the coverage threshold of 3. Peers arriving between t_1 and t_2 are likely to complete downloading the content despite the absence of a publisher. At t_2 , the swarm enters an idle period that lasts until a publisher reappears at time t_3 . Extant peers at the end of the busy period as well as peers arriving during the idle period find the content unavailable (represented by dotted lines) until the idle period ends at t_3 . Because of idle waiting, these peers experience a longer *wait-time*, i.e., the time since a peer arrives until it completes the download (and departs).

In what follows, we assume that peer and publisher arrivals are Poisson, and their staying times are exponentially distributed. For each swarm k , the publisher arrival rate, r_k , is small and in particular is much smaller than the peer arrival rate, λ_k . The mean stay time u_k of a publisher is assumed to be the same as the mean time s_k/μ that a peer spends in the system fully downloading the file in a single

busy period.

We initially consider the case where the busy period is defined w.r.t. a coverage threshold of one, i.e., a peer arriving during a busy period always finishes the download in that busy period and the last peer to finish ends the busy period. A peer arriving during an idle period finds the file unavailable and immediately leaves (*impatient peers*).

The average number of peers served in a busy period for swarm k is $E[N_k] = e^{\lambda_k s_k / \mu} - 1$. The average length of a busy period is simply $E[B_k] = (e^{\lambda_k s_k / \mu} - 1) / \lambda_k$. The probability P_k that a peer arrives to find the content unavailable is $P_k = \frac{1/r_k}{E[B_k] + 1/r_k}$.

To exemplify what happens when the coverage threshold differs from one, let us now consider the scenario where, as soon as the population reaches a critical size n , the first publisher leaves. After the publisher leaves, peers complete the download only if they arrive before the coverage drops below a threshold $m < n$. A peer arriving during an idle period finds the file unavailable and waits until a publisher returns (*patient peers*).

Let $B[n, m]$ be the expected length of a *residual busy period* which initiates with n leechers and ends as soon as the population size reaches m .

THEOREM 2.1. *For all n ,*

$$B[n, 0] = \left[\sum_{i=1}^n \frac{s_k}{i\mu} + \frac{s_k}{\mu} \sum_{i=1}^{\infty} \left(\frac{s_k \lambda_k}{\mu} \right)^i \frac{(n+i)! - n!i!}{i!(n+i)!i} \right] \quad (1)$$

For $l < n$, $B[n, l] = B[n, 0] - B[l, 0]$.

We heuristically use Theorem 2.1 to estimate the expected wait-time of peers:¹

$$E[T] = \frac{s_k}{\mu} + \frac{1/r_k}{1/r_k + B[\lambda_k s_k / \mu, m] + u_k} 1/r_k \quad (2)$$

Equation (2) models a system which cycles through three regimes: (1) a publisher is available; (2) the publisher is off line but the content is still available; (3) the content is not available. A peer always incurs a waiting time of at least s_k / μ . If it arrives during the third regime, its wait-time is extended by $1/r_k$, corresponding to the mean time for the arrival of a publisher. Note that we have assumed that no publishers arrive in the second regime.

3. BUNDLING

In this section we show how the proposed model can be used to provide insights on the benefits of bundling. For that purpose, we consider the special case when the peer arrival rates are the same for all files, i.e., $\lambda_k = \lambda$, and the file sizes are the same, i.e., $s_k = s$ for all K files. Note that if K files are bundled, the arrival rate to this swarm is $K\lambda$ and the size of the bundled file is Ks . Omitting the subscripts of the variables described above when referring to their bundle counterparts, we have: $E[N] = e^{K^2 \lambda s / \mu} - 1$, $E[B] = (e^{K^2 \lambda s / \mu} - 1) / K\lambda$, $P = \frac{1}{1 + R(e^{K^2 \lambda s / \mu} - 1) / (K\lambda)}$. Note that $E[N]$ and $E[B]$ are a factor $e^{\Theta(K^2)}$ larger than the corresponding values for an individual swarm.

THEOREM 3.1. *The average busy period length is given by $E[B] = e^{\Theta(K^2)}$.*

¹In [1] we present exact closed form expressions for the wait-time of patient peers when the coverage threshold is 1.

PROOF. All the proofs are available in [1]. \square

It can also be verified that $-\log P_k = \Theta(1)$ and $-\log P = \Theta(K^2)$. Thus, bundling reduces the probability of not finding the content by an $e^{-\Theta(K^2)}$ factor.

THEOREM 3.2. *Let $R = e^{-cK^2}$ be the bundled arrival rate of publishers ($c > 0$). Then, $P_k = e^{-\Theta(1)}$ and $P = e^{-\Theta(K^2)}$.*

To illustrate in quantitative terms the applicability of our model we consider the problem of deciding the optimal level of bundling, i.e., we want to find the value of K that minimizes the wait-time. Note that small values of K lead to an increase in download time without a significant improvement in availability, while large values of K lead to large bundles, which necessarily will take longer to download.

We numerically evaluate eq. (2) with $s_k / \mu = K/80$ s, $\lambda = K/60$ 1/s, $1/r_k = 1/300$ s and $u_k = 900$ s. We assume that whenever a seed departs it leaves behind it $n = K^2 \lambda_k s / \mu$ peers. Fixing a coverage threshold of $m = 4$, for K varying between 1 and 8 we obtained the following values for the expected download time of the peers, in hundreds of seconds: (7.55, 7.76, 3.37, 3.20, 4.00, 4.80, 5.60, 6.40). When $K = 1$ many peers get blocked due to the unavailability of the content. When $K \geq 5$, on the other hand, the swarm is *self sustaining*: even when the publisher is not available the peers still smoothly complete their downloads.

To evaluate the model against real experiments, we ran BitTorrent clients in PlanetLab under conditions similar to those described in the previous paragraph. We varied the bundling factor K from 1 up to 8. In Figure 2(b) we show the expected download time observed in the PlanetLab experiments as a function of K . The optimal bundling factor predicted by the model, $K^* = 4$, matched the one observed in the experiments and the linear increase in the download time for $K > 5$ was correctly captured by the model.

4. CONCLUSION

Content availability is a serious problem in peer-to-peer swarming systems today, e.g., almost half of the BitTorrent swarms are unavailable half of the time, which does not bode well for the increasing commercial interest in integrating swarming with server-based content dissemination. Our work is a first step towards developing a foundational understanding of content availability in swarming systems. Our model suggests that bundling of content, a common practice among swarm publishers today, improves availability.

Acknowledgement: This work was supported in part by the NSF under award CNS-0519922 and CNS-0721779. Research of DSM also funded by CAPES/Fullbright (Brazil).

5. REFERENCES

- [1] MENASCHE, D., ROCHA, A., LI, B., TOWSLEY, D., AND VENKATARAMANI, A. Bundling builds availability in BitTorrent. *UMass TR UM-CS-2009-010* (2009).
- [2] QIU, D., AND SRIKANT, R. Modeling and performance analysis of BT-like P2P Nets. In *SIGCOMM* (2004).
- [3] RAMACHANDRAN, A., DAS SARMA, A., AND FEAMSTER, N. Bitstore: An incentive-compatible solution for blocked downloads. In *NetEcon* (2007).