Availability in Swarming Systems

Daniel Sadoc Menasche
University of Massachusetts

Antonio A. Aragao Rocha
Federal Univ. of Rio de Janeiro

Bin Li
Tsinghua University

Don Towsley
University of Massachusetts

Arun Venkataramani
University of Massachusetts

1. INTRODUCTION

Unstructured peer-to-peer systems have a number of properties which explain their tremendous popularity, e.g. the capacity of the system scales with the number of peers and peers can come and go in a random fashion. Nevertheless, when a peer leaves the system it may carry the last copy of a block of a file and for this reason it is hard to predict the availability of the files and consequently the performance of the system as a whole. In this article we model peer-to-peer systems, à la BitTorrent, where content may become unavailable.

BitTorrent uses P2P file swarming to disseminate content. A swarm is formed by a set of users interested in downloading a given file (or set of files, packaged in a bundle). The file is chopped into blocks that peers download from the server and from each other. A tracker is used to coordinate the interaction among the peers. Periodically, a peer queries the tracker and obtains a random subset of other peers in the swarm. A leecher is a peer that hasn't finished its download yet. After concluding the download, the leecher becomes a seed. Peers that have incentives to make the content available are referred to as publishers. The very first peer that makes available a given content is the first publisher.

This paper presents novel peer-to-peer models that serve as the foundation of a thesis in the more general context of simple mechanisms (e.g., swarming coupled with bundling) robust to faults and workload changes. The key contributions of this work are the following: (1) we propose a model to capture the effect of content unavailability on performance; (2) we show the applicability of our model to the problem of deciding the optimal bundling strategy for a publisher.

Assumptions. We consider peers arriving to each swarm according to a Poisson process with rate λ peers/s. Peers download chunks from each other and from the publishers, and leave the system as soon as they finish their downloads. The distribution of chunks across peers is assumed to be uniform. The peer selection mechanism is also uniform. The download rate of peers is equal to μ download completions/s if the content is available. Note that the capacity of swarming systems scales with the number of peers, hence if the content is available μ does not depend on the current state of the system.

Related work. [4] point out that peers who want a file may need to wait for a long period of time for some blocks that are not available. In this paper we propose bundling to mitigate this availability problem.

[2, 3] proposed models to analyze the performance of p2p systems using Markov Chains or differential equations. However, none of these models has taken content availability into account. [1] analyze file availability but their performance model does not take into account file unavailability due to churn.

2. ONE PUBLISHER THAT LEAVES AND NEVER RETURNS

We begin with the case of a single publisher that leaves and never returns. We introduce a simplifying assumption regarding the conditions for content availability, namely, that content becomes unavailable if and only if the number of peers in the system goes to zero (this assumption is eliminated in Section 4). Furthermore, for simplicity, we assume that the publisher and peers each remain in the system for an exponentially distributed period of time equal to the one of a typical peer. Under these conditions we can model this system as an $M/M/\infty$ queue. The period of time during which the content is available corresponds to the busy period of the queue. Then, the busy period is given by $(e^{\lambda/\mu} - 1)/\lambda$.

3. IMPATIENT PEERS

We now consider a system where publishers of the same content arrive according to a Poisson process with parameter γ. Each publisher remains for an average time $1/\alpha$. Peers that arrive to the system when the content is not available leave immediately. We model the system as a Markov Chain with state (l, p) where l and p stand for the number of leechers and publishers in the system, respectively. Starting from state (l, p), the transitions are given as follows:

1) $l = 0$, $p = 0$: $(0, 0) \xrightarrow{\gamma} (0, 1)$

2) otherwise: $(l, p) \xrightarrow{\lambda} (l + 1, p)$, $(l, p) \xrightarrow{l\mu} (l - 1, p)$ $(l > 0)$, $(l, p) \xrightarrow{\gamma} (l, p + 1)$, $(l, p) \xrightarrow{p\alpha} (l, p - 1)$ $(p > 0)$

(Note that if $\lambda = \gamma$ and $\mu = \alpha$ we have an $M/M/\infty$ queue, as in the last section.) Using this model we can compute, for instance, the fraction of leechers that arrive to find the content unavailable, which is given by the probability of state $(0, 0)$.

4. PATIENT PEERS

In this section we consider the case where a single publisher for the content goes *on* and *off*. The state (l, p, a) of the system is characterized by l, the number of leechers,

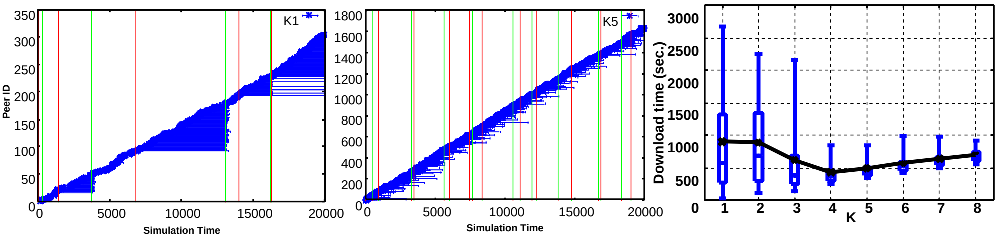

Figure 1: (a) Simulation results for K=1; (b) Simulation results for K=5; (c) PlanetLab results.

p, the state of the publisher [on (1) or off (0)] and a, the content availability [available (1) or not (0)]. Starting from state (l, p, a), the transitions go as follows:

1) $l > 2$: (a) $p = 0$, $a = 1$: $(l, 0, 1) \xrightarrow{l\mu} (l-1, 0, 1)$, $(l, 0, 1) \xrightarrow{\lambda} (l+1, 0, 1)$, $(l, 0, 1) \xrightarrow{\gamma} (l, 1, 1)$; (b) $p = 0$, $a = 0$: $(l, 0, 0) \xrightarrow{\lambda} (l+1, 0, 0)$, $(l, 0, 0) \xrightarrow{\gamma} (l, 1, 1)$; (c) $p = 1$, $a = 1$: $(l, 1, 1) \xrightarrow{\lambda} (l+1, 1, 1)$, $(l, 1, 1) \xrightarrow{l\mu} (l-1, 1, 1)$, $(l, 1, 1) \xrightarrow{\alpha} (l, 0, 1)$.

2) $l = 2$: (a) $p = 0$, $a = 1$: $(2, 0, 1) \xrightarrow{2\mu} (1, 0, 0)$, $(2, 0, 1) \xrightarrow{\lambda} (3, 0, 1)$, $(2, 0, 1) \xrightarrow{\gamma} (2, 1, 1)$ (b) otherwise: b) and c) for case $l > 2$.

3) $l = 1$: (a) $p = 0, a = 0$: $(1, 0, 0) \xrightarrow{\lambda} (2, 0, 0)$, $(1, 0, 0) \xrightarrow{\gamma} (1, 1, 1)$; (b) $p = 1, a = 1$: $(1, 1, 1) \xrightarrow{\lambda} (2, 1, 1)$, $(1, 1, 1) \xrightarrow{\alpha} (1, 0, 0)$, $(1, 1, 1) \xrightarrow{\mu} (0, 1, 1)$.

4) $l = 0$: (a) $p = 1, a = 1$: $(0, 1, 1) \xrightarrow{\lambda} (1, 1, 1)$, $(0, 1, 1) \xrightarrow{\alpha} (0, 0, 0)$; (b) $p = 0, a = 0$: $(0, 0, 0) \xrightarrow{\lambda} (1, 0, 0)$, $(0, 0, 0) \xrightarrow{\gamma} (0, 1, 1)$.

Premature Content Unavailability: We now eliminate the assumption that content becomes unavailable only when the number of peers in the system goes to 0. For that purpose, we define a *coverage threshold*, T. We modify the Markov Chain presented in the last section to account for the fact that if the number of peers in the systems is less than T there is a probability $q > 0$ that content becomes unavailable.

1) $T > l > 2, p = 1, a = 1$: $(l, 1, 1) \xrightarrow{\alpha q} (l, 0, 1)$, $(l, 1, 1) \xrightarrow{\alpha(1-q)} (l, 0, 0)$

2) $T > l > 2, p = 0$, $a = 1$: $(l, 0, 1) \xrightarrow{l\mu q} (l-1, 0, 1)$, $(l, 0, 1) \xrightarrow{l\mu(1-q)} (l-1, 0, 0)$

3) other states: as in the previous model

5. BUNDLING

To show the applicability of our model, we consider the problem of deciding the optimal level of bundling in a peer-to-peer system. Consider a server publishing F files using swarming. The server has the flexibility to decide how to combine the files. Assume a homogeneous population of files each one with popularity λ and size s. If K files are bundled, the arrival rate to this swarm is $K\lambda$ and the size of the bundled file is Ks.

We want to find the value of K that minimizes the net download 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 simulate the model presented in Section 4 with $q = 1$, $T = 11$, $\lambda = (1/60)K$, $\mu = (1/80)/K$, $\gamma = (1/300)$, $\alpha = (1/900)$ using Tangram II (`http://www.land.ufrj.br`). In Figures 1(a) and (b), each line segment starts at the instant the peer arrives and ends when the peer departs. Green and red vertical lines represent arrivals and departures of the publisher. When $K = 1$ (Figure 1(a)) many peers get blocked due to the unavailability of the content. When $K = 5$ (Figure 1(b)), on the other hand, the swarm is *self sustaining*: even when the publisher is not available the peers still smoothly complete their downloads. For K varying between 1 and 8, we obtained the following values for the expected download time of the peers: (1586, 1563, 1336, 522, 412, 487, 568, 650).

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

6. CONCLUSION AND FUTURE WORK

In this work we model content unavailability in swarming systems and show that bundling can be used to mitigate it. In terms of future work, in the models presented in this paper we do not account for ISP locality/friendliness. Incorporating such aspects may lead, from an economic perspective, to interesting extensions of the model.

Acknowledgment: This work was supported in part by the NSF under award numbers CNS-0519922 and CNS-0721779. Research of DSM also funded in part by a scholarship from CAPES/Fulbright (Brazil).

7. REFERENCES

[1] Fan, B., Chiu, D.-M., and Lui, J. Stochastic analysis and file availability enhancement for bt like file sharing systems. In *IWQoS* (2006).

[2] Ge, Z., Figueiredo, D., Jaiswal, S., Kurose, J., and Towsley, D. Modeling peer-peer file sharing systems. In *Infocom* (2003).

[3] Qiu, D., and Srikant, R. Modeling and performance analysis of BT-like p2p nets. In *SIGCOMM* (2004).

[4] Ramachandran, A., Sarma, A., and Feamster, N. Bitstore: An incentive compatible solution for blocked downloads in BT. In *NetEcon* (2007).