Can identical BitTorrent peers experience different download times?

Fabrício Murai[1], Antonio Rocha[2], Daniel Figueiredo[1], Edmundo de Souza e Silva[1]

1 COPPE/Computer and System Engineering Dept., Federal University of Rio de Janeiro, Brazil

2 Mathematic Dept., Rural Federal University of Rio de Janeiro, Brazil

{fabricio,daniel,edmundo}@land.ufrj.br guto@ufrrj.br

1. Introduction

Understanding and characterizing the performance of BitTorrent (BT) through mathematical models has been an active topic of research mainly due to its high complexity and tremendous practical efficiency. One of the first BT models, presented in [3], predicts the download rates of homogeneous peers. More recently, BT models to capture heterogeneous peer population with respect to their upload and download capacities have also been proposed, as well as other relevant system characteristics [2, 1]. However, all models predict that identical peers (with respect to their upload capacities) will have identical performance (with respect to download times). We consider a BT swarm where all peers have identical upload capacities but unconstrained (or large) download capacities. In this context, we identify and characterize a phenomenon that has not been previously observed: homogeneous peers experience heterogeneous download rates. This is surprising as peers are identical and it has not been captured by any prior model. Moreover, this behavior has several negative impacts, including high variability of download times, unfairness with respect to peer arrival order, bursty departures and content synchronization among the peers. Although it is more prevalent in unpopular swarms (very few peers), these by far represent the most common type of swarm in BT.

2. The observed behavior

Consider a swarm formed by a single seed and 5 leechers, all of them having identical upload capacity (64 kBps), but unconstrained download capacities. Leechers download a file split in 1000 pieces (256MB) and exit the swarm after download completion. The seed never leaves the swarm. This system was evaluated using a detailed packet-level simulator of BT and also an instrumented client of BT running on PlanetLab.

Figures 1a and 1b show the evolution of the swarm size as a function of time for both simulation and experimental results for two different leecher arrival patterns. In Figure 1a, peers leave the swarm in the order they arrived (i.e., FIFO) and have a relatively similar download time. Thus, the download time is relatively indifferent to arrival order (with the exception of the first peer). Figure 1b shows that for a different arrival pattern (in fact, the inter-arrival times of peers are mostly preserved), an unexpected behavior can be observed in the system dynamic: despite the significant difference on arrival times, all five leechers completed their respective download nearly at the same time. The time inter departures is small comparing to the download time, which characterizes bursty departures. It means that peers that arrive later to the swarm have a smaller download time. In fact, the fifth peer completed the download in about half the time of the first leecher. Thus, the system is unfair with respect to the arrival order of leechers, with late arrivals being significantly favored. Why is BT exhibiting such dynamics?

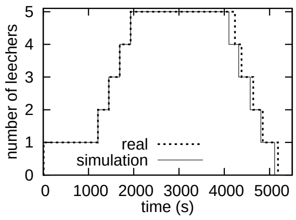

(a) Arrival intervals: 10sec, 10min, 4min, 4min, 4min.

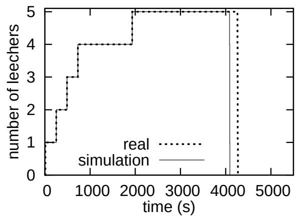

(b) Arrival intervals: 10sec, 4min, 4min, 4min, 10min.

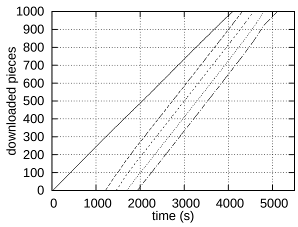

(c) Corresponding to Figure 1a.

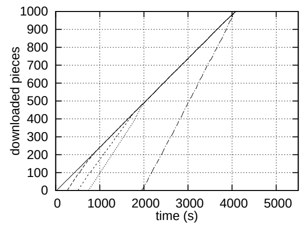

(d) Corresponding to Figure 1b.

Fig. 1: (a) and (b) Evolution of the swarm size; (c) and (d) Evolution of the number of downloaded pieces.

To understand the unexpected behavior exhibited by BT in Figure 1b, let us analyze the number of pieces each leecher has downloaded over time. Consider figures 1c and 1d where each curve indicates the number of pieces downloaded by a given peer for the corresponding scenario in figures 1a and 1b, respectively. One can note that the slope of each curve corresponds to respective leecher's download rate. In Figure 1c, which corresponds to the scenario considered in Figure 1b, we note that despite the slope of the first leecher is smaller than that of the remaining peers, the curves never meet. We also note that all other leechers have very similar slopes.

The results illustrated in Figure 1d which corresponds to the scenario considered in Figure 1b shows a very different behavior. The slope of the first peer is practically constant, remaining unchanged by the arrival of other peers. The slope of all other peers is larger than that of the first peer, meaning the curves may eventually meet. When two curves meet, the corresponding leechers have the same number of pieces. The figure also shows that younger peers do not pass ahead the first peer, but instead follow it in the number of pieces. Finally, their slope are rather similar, except the fifth whose slope is slightly larger.

2.1 Consequences of heterogeneity in homogeneous swarms

The observations above imply essentially that the download time of peers are quite different, despite their homogeneous upload capacity. In summary, the consequences are: (i) **Variability in download times.** Since peers can experience different download rates, download times can also differ; (ii) **Unfairness with respect to peer arrival order.** Since peers download rates, and thus download times, may depend on their arrival order, the system is inherently unfair; (iii) **Content synchronization.** As the first leecher is downloading the file at the same rate at which the seed pushes new pieces into the swarm, whenever a leecher reaches the same number of pieces than it, they have exactly the same content; and (iv) **Bursty departures.** A direct consequence of content synchronization is bursty departures. Peers tend to leave the swarm within a small time interval despite arriving at relatively far apart instants.

Of course, the prevalence of the phenomenon and its consequences depend directly on the parameters of the swarm, such as the arrival pattern of peers, upload capacity of seed and leechers and file size. Intuitively, a file with a larger number of pieces or a seed with a lower upload capacity increase the probability that the consequences above occur.

2.2 Heterogeneity under Poisson arrivals

The above behavior does not require deterministic arrivals or any crafted leecher arrival pattern. We investigate the swarm behavior under Poisson arrivals with a large amount of evaluations, using both simulations and prototype-based experiments. Because of space limitation, we present here only a single result obtained from an experimental evaluation deployed on Planetlab. For the real experiments, we used an instrumented BT client and created private torrents swarms. In those swarms, a single seed with capacity $c_s = 50$ KBps was always available uploading content, all leechers with identical capacity $c_l = 50$ KBps arrived according to a Poisson process with rate $\lambda = 1/100$ leechers/sec and left after concluding the download of a $S = 20$ MB file long.

Figure 2a illustrates the evolution of the swarm size in the experiment. We observe many occurrences of bursty departures. This is consistent with our findings in Figure 1b, despite the fact that the arrival pattern is now Poisson. Figure 2b shows the complementary cumulative distribution (CCDF) of download times. The curve exhibits several sharp drops caused by the significant amount of peers with similar download times. Considering the minimum and maximum values shown in the distribution (approximately 260 and 500 seconds, respectively), we can also verify the high variability in download times (in which the mean value is 356 seconds).

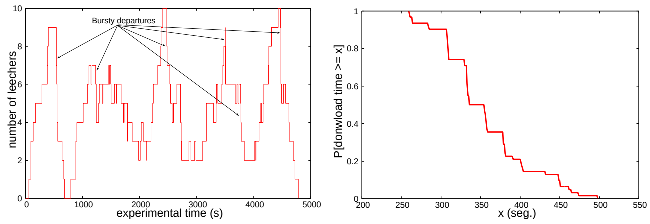

(a) Evolution of swarm size. (b) CCDF of download time.

Fig. 2: Results from experiments with Poisson arrivals.

3. Ongoing work

We are now developing a mathematical model attaining to understand the origin of the heterogeneous download times in a homogeneous swarm. From a detailed analysis of our simulation and experimental results, we observe that a young leecher is less likely to have pieces interesting to older ones than the opposite. Due to this fact, the uplink capacity which would otherwise be dedicated to older leechers is either redistributed or remains idle until it receives a piece interesting to them, i.e. a piece that is not owned by an older leecher. We conclude that the relationship of interests plays an important role on the download rates between peers. Moreover, this is the root cause for the heterogeneous download rates in unpopular swarms. Thus, our goal is to embed this relationship in a flow model to estimate the rate at which a leecher can upload to its neighbors. This model would enable us to analyze the impact of system parameters on the emergence of the referred phenomenum and to predict its consequences.

References

[1] A.L.H. Chow, L. Golubchik, and V. Misra. Bittorrent: An extensible heterogeneous model. In *IEEE INFOCOM*, pages 585 –593, apr. 2009.

[2] W.C. Liao, F. Papadopoulos, and K. Psounis. Performance analysis of bittorrent-like systems with heterogeneous users. *Performance Evaluation*, 64(9-12):876 – 891, 2007.

[3] D. Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *ACM SIGCOMM*, pages 367–378, 2004.