Early Stage Denial of Service Attacks in BitTorrent: An Experimental Study

Soufiane Rouibia
TMG, France
Email: rouibia@tmg.eu

Jonathan Vayn, Olivier Beauvais, Guillaume Urvoy-Keller
Institut Eurecom, France
Email: {vayn, beauvais, urvoy}@eurecom.fr

***Abstract*— BitTorrent is a popular p2p file replication application, which aims at replicating a given content as fast as possible on a set of peers. The algorithms of BitTorrent used to elect remote peers with whom a peer collaborates and also which pieces of the content it offers, have proved to be highly efficient. This means that a high level of parallelism is achieved among the peers as a given peer always has a high chance to find another peer that holds content it is currently missing. Still, at the beginning of a BitTorrent session, pieces of the content have to be obtained from only a few peers (in general a single one called the initial seed) that hold a full copy of the file to be replicated. In this work, we aim at evaluating the ability of a BitTorrent session to survive to a denial of service attack that would disconnect the initial seed from the network. We address this issue through experimentation. Our main conclusion is that BitTorrent is highly resilient to this attack as neither the ability to obtain a full copy of the content nor the actual replication speed are affected by the disconnection of the initial seed if the attack is not carried out at the very early stage of the session.**

I. Motivation

A BitTorrent session, a.k.a. a swarm, gathers a set of peers that want to quickly replicate a given content. In the BitTorrent terminology, one distinguishes between leechers, i.e., peers that have not completed yet the download of the content and seeds that have completed the download but remain in the swarm to service other peers.

A host of works have focused on the performance of BitTorrent, e.g., [4] studies the availability of content at a large scale. Fewer studies focus on possible attacks against BitTorrent. In [1], the authors demonstrate through passive and active measurements that torrents corresponding to top box-office movies are attacked. Attackers are modified BitTorrent clients that were either sending fake control messages or fake content. In [2], attacks using the signaling plane of BitTorrent are evaluated through simulations.

The main objective of this work is to assess the feasibility of an attack against the seeds of a swarm. An attack is considered successful if the leechers cannot reach the seed state or, alternatively if the download rate of the leechers becomes too low after the departure of the seeds, as we can expect here that users will become impatient and disconnect from the swarm.

II. Experimental Set-Up

Our work is based on empirical experimentation. We consider a swarm created to download an mp3 file of 28 Mbytes. The swarm consists of 45 leechers. While modest, such a value is in line with typical swarm sizes observed in the wild that range between a few tens and a few hundreds peers in general - see [5]. To implement and test our attack, we used OpenVZ, a virtual network environment distributed under a public license (http://openvz.org/). Constraints on the number of virtual machines that could be run simultaneously on our test machine further dictated this choice of swarm size.

As each leecher gets a list of 40 peers from the tracker, it ends up being connected to almost all the peers in the swarm, thus maximizing the efficiency of the protocol.

We proceeded as follows to test the impact of the seed disconnection on the application efficiency. We start the seed after the leechers in order to be sure that all leechers are ready to perform requests to the seed. At time T, we shut down the seed. Then, we wait enough to be sure that the swarm reaches its steady state where all leechers have downloaded all the pieces available in the swarm at time T^+, i.e. just after the initial seed as been shut down. We performed more than 200 experiments for different T values between 60 seconds and the state of the peers at the end of the experiment.

III. Results

A. Likelihood of download completion

Figure 1(a) represents the final download ratio as a function of the average download ratio when the seed is shut down. Each cross represents one experiment.

For an average completion ratio greater than 60%, we observe that disconnection of the seed has no effect. In this situation, leechers are able to complete their download in every case. For an average download ratio lower than 20%, the attack is almost always a success. In between 20% and 60%, the outcome of the attack seems difficult to predict. To better understand what is happening in this gray zone, between 20 and 60%, we computed the probability of success of the attack as a function of the average completion ratio when the seed is shut down. We obtain the graph in Figure 1(b). Probability of download completion ramps up very quickly between 40% and 60%. We can thus conclude that it is useless to perform an attack once leechers have reached a ratio above 40%.

B. Download rates before and after the attack

Our second criterion to determine if an attack is a success is the download rate of the leechers in the system. Indeed, the shutdown of the initial seed could be considered as a

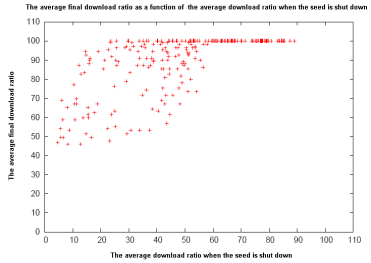

(a) Final download ratio w.r.t. the average download ratio at the attack time

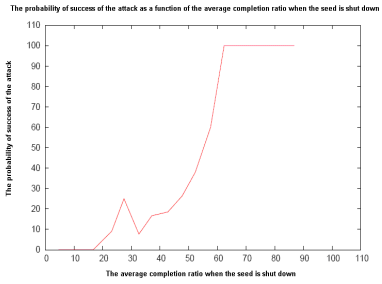

(b) Probability of success

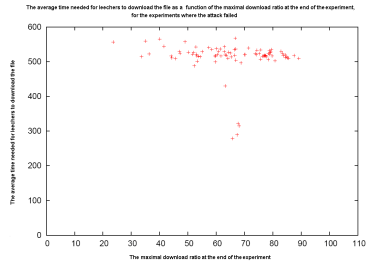

(c) Average time to completion of leechers

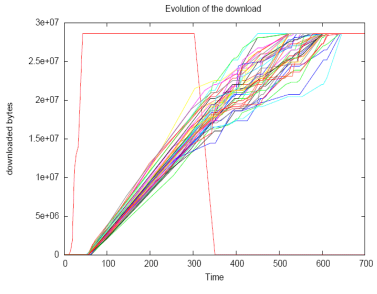

(d) Download rates over time of leechers

Fig. 1.

success if it slows down enough the download rates of the leechers. We computed - see Figure 1(c)- the average time needed for leechers to download the file as a function of the maximal download ratio[1] at the end of the experiment, for the experiments where the attack failed. We observe that the time for leechers to complete the download is almost independent from the state at time T^+. This means that the departure of the seed has no real effect on the download rate of the remaining leechers. Due to BitTorrent algorithms, data transfers between leechers are the main source of traffic in the swarm. Those results are in line with the ones in [3] where the authors highlight the efficiency of the chunk and peer selection strategies of BitTorrent. We can conclude that the shut down of the seed can only affect the ability of peers to gather all chunks of the file but does not influence their download rates. Figure 1(d) further confirms this finding. This figure depicts the cumulative downloaded bytes per peer for an experiment where the attack fails and the seed is shut down at $T = 300$ seconds. The lonesome curve on the left side corresponds to the initial seed. We do observe no major qualitative difference between the evolution of each peer before and after the departure of the initial seed.

[1]We checked empirically that considering the average or the maximum completion ratio among the peers leads to the same conclusions. This is because peers are homogeneous in our swarm and thus progress similarly during the download phase.

IV. Discussion

In this work, we aimed at quantifying the effect of an attack that disconnects the initial seed from its swarm. We observed that this attack is a success only if the average completion ratio of leechers is less than 40%. It should even be less than 20% if we want a high probability of success. With our work, we have only scratched the surface of the problem. For instance, it is highly likely that the threshold (here 40% of average download ratio) to be used needs to be a function of the swarm size: the larger the swarm, the smaller the threshold. Still, our methodology and experimental set-up have the potential to bring useful information about the resilience of BitTorrent against denial of service and other attacks in general, and we intend to continue our efforts using a faster testbed to be able to scale to larger swarm size and evaluate other attacks against the BitTorrent protocol.

References

[1] P. Dhungel, D. Wu, B. Schonhorst, and K. W. Ross, "A Measurement Study of Attacks on BitTorrent Leechers", In *IPTPS*, 2008.

[2] M. A. Konrath, M. P. Barcellos, and R. B. Mansilha, "Attacking a Swarm with a Band of Liars: evaluating the impact of attacks on BitTorrent", In *IEEE P2P 2007*, pp. 37–44, 2007.

[3] A. Legout, G. Urvoy-Keller, and P. Michiardi, "Rarest First and Choke Algorithms Are Enough", In *ACM SIGCOMM/USENIX IMC'2006*, October 2006.

[4] G. Neglia, G. Reina, H. Zhang, D. Towsley, A. Venkataramani, and J. Danaher, "Availability in BitTorrent Systems", In *INFOCOM 2007*, pp. 2216–2224, 2007.

[5] G. Urvoy-Keller and P. Michiardi, "Impact of Inner Parameters and Overlay Structure on the Performance of BitTorrent", In *9th IEEE Global Internet Symposium 2006 in conjunction with IEEE Infocom 2006*, Barcelona, Spain, April 2006.