Who is going to be the next BitTorrent Peer Idol?

Florian Adamsky[1,3], Syed Ali Khayam[2], Rudolf Jäger[3] and Muttukrishnan Rajarajan[1]

[1]City University London, London, United Kingdom
[2]PLUMgrid, Incorporation, Sunnyvale, CA, United States of America
[3]Technische Hochschule Mittelhessen (THM), Friedberg, Germany

Abstract

Active measurement studies show that the P2P file sharing protocol BitTorrent is highly under attack. Moreover, malicious peers can easily exploit the original seeding algorithm and therefore reduce the efficiency of this protocol. In this research work, we propose a novel seeding algorithm that requests peers to vote for their best sharing peers. Our results show that this incentive mechanism makes BitTorrent harder to exploit without losing performance. In some situations our algorithm even outperform other seeding algorithms. The peer exchange—that comes as a side effect—reduces the dependency on a centralized tracker and increases the robustness and the efficiency. We studied the effectiveness of our approach in a real testbed comprising 32 peer.

Introduction

- If a P2P network allows peers to be selfish (only downloading but not uploading), then this degrads the system performances and make the network weak. This problem is called free-riding.
- Malicious peers can exploit this problem to attack the network.
- BitTorrent uses an incentive mechanism called choking algorithm that favors peers who upload.

Problem Statement

- The choking algorithm is different when the peer has the complete file (seeder).
- A seeder uses its upload rate rather than its download rate to decide which peer to unchoke..
- This does can be easily exploited by malicious peers.

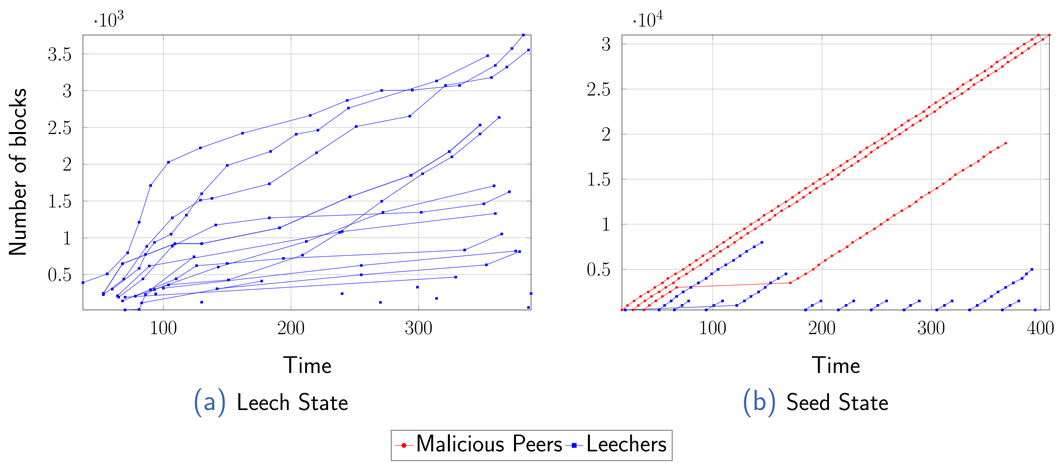

(a) Leech State (b) Seed State

Figure: Upload piece distribution of the choking algorithms in leech state and in seed state. Every line represents another peer. The data was produced with 1 seeder that has 5 Mbit/s upload limit, 29 leechers with 1 Mbit/s and 3 fast peers with no limit.

Seeding Algorithm: Peer Idol

Leecher Side

- If a peer wants to acquire pieces from a seeder, it has to send every unchoke round a new BitTorrent message to the seeder, that we call a vote.
- This vote contains n peers that have shared the most.

Seeder Side

- The seeder who gets the vote awards each voted peer with points: Peer A gets 3 points, peer B gets 2 points and peer C gets 1 point.
- Every unchoke round the seeder sorts all candidates by their score of each peer and unchokes the peers with the highest score.
- If $N = \{1, 2, \ldots, n\}$ is the peer set of the seeder, then $I \subseteq N$ contains the peers which are interested. Then for every peer $p \in I$ we calculate the PI score as follows:

$$PI(p) = \sum_{i=1}^{|I|} V_i(p). \quad (1)$$

- The function $V_i(p)$ returns 1 if peer i voted for peer p and 0 if not.
- If $PI(p_n) = PI(p_m) \rightarrow$ the peer that has waited the longest gets a higher priority.
- If a vote contains peers to whom the seeder is not connected to, it will add this peer to a list of potential peer candidates.

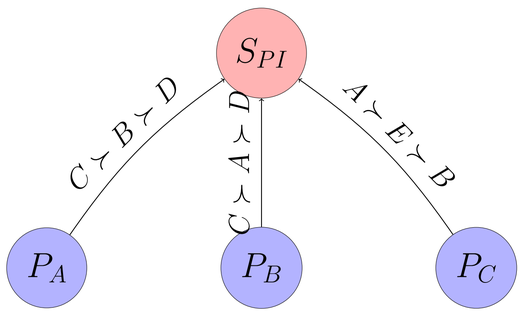

Figure: Example of how leechers P_A–P_C send their votes to the seeder S_{PI}.

Additional Security Features

- A peer has to send a vote in order to get voted.
- A peer gets disconnected and blacklisted if the vote contains more than n peers, the IP address of the requesting peer, or repeated peers.

Other Seeding Algorithms

Fastest Upload (FU) favors the peers who upload the fastest. This algorithm is still used by around 29.5 % of the clients.

Round Robin (RR) is a well-known algorithm, especially as a scheduling algorithm for processes in OS. In context of BitTorrent, it gives each peer a constant number of pieces and rotates every n pieces.

Longest Waiter (LW) sorts all unchoked and interested peers according to their time they were last unchoked. LW unchokes the u peers who have waited the longest. Each unchoked peer gets at least two rounds until LW unchokes them.

Anti Leech (AL) prefers peers when they only have a few pieces or when they have nearly all pieces.

Performance Evaluation

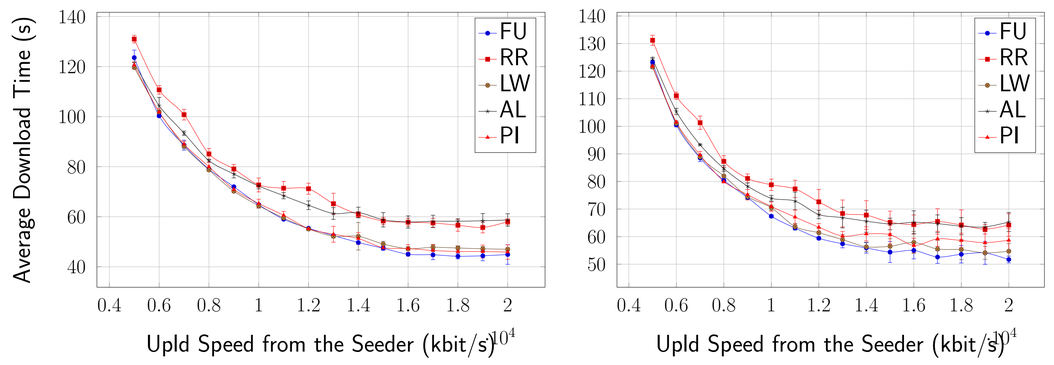

(a) Leechers have no upload or download limit (b) Leechers have different download limits and no upload limit

Figure: Effects of the different seeding algorithms on the average download speed in different environments. The upload speed of the seeder was gradually increased. The error bars show 99 % confidence intervals.

Security Analysis

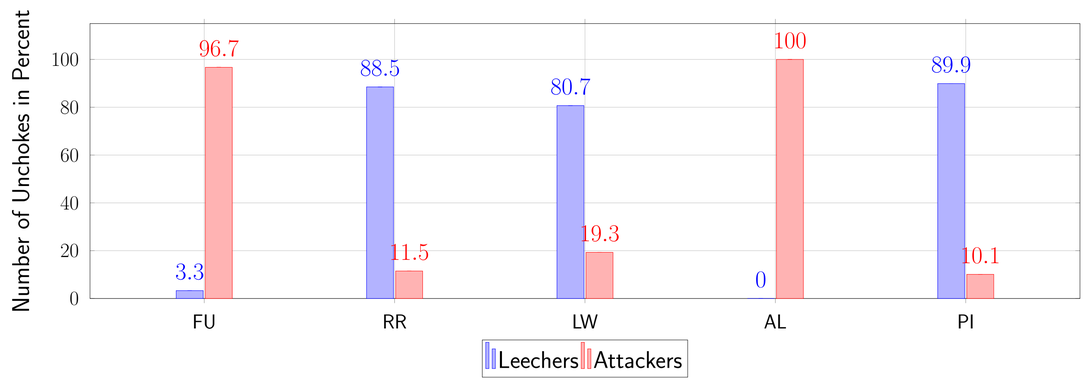

Figure: The unchoke ratio of attackers and leechers of the different seeding algorithms without the optimistic unchoke slot. The data was produced with 1 seeder that has 5 Mbit/s upload limit, 29 leechers with 900 kbit/s download limit and 3 malicious peers. All values are average values of ten iterations. The 99 % confidence intervals of all values is < 0.05.

Stability Evaluation

Table: The number of connected peers of the seeder.

Seeding Algorithms:	FU	RR	LW	AL	PI
Number of Connected Peers:	5.1	5.1	4.8	4.9	29.0

Conclusions

- Our results support our hypothesis that PI is more robust against bandwidth attacks and does not lose performance compared to other algorithms.
- In our experiment, PI was even faster than RR and AL.
- If votes contain unknown peers, the seeder saves them to a candidate list which can be used for later contact.
- Summarized, our novel choking algorithm in seed state implements the incentive mechanism in BitTorrent consequently through all peers.