Why Does BitTorrent Work So Well?

Simon Patarin and David Hales

Keywords: BitTorrent, computational sociology, altruism

BitTorrent (BT) is currently the most popular file-sharing peer-to-peer (P2P) clients [1]. Several reports have shown that it accounts for the majority of peer-to-peer traffic, up to one third of the global Internet traffic. However, leaving the media attention aside for one moment, we ask the question: Why does BT work so well?

BT attempts to build robustness to freeloading (i.e. downloading without uploading) by implementing a tit-for-tat-like strategy (TFT) within its protocol. It is often believed that this strategy alone is responsible for the high levels of cooperation found. The TFT strategy was championed in Axelrod's book "the evolution of cooperation" [2]. BT works by groups of peers (called swarms) with an interest in downloading a specific file coordinating and cooperating to speed-up the process [3]. Each peer in the swarm stores pieces of the file. Cooperating peers download and upload required pieces. If a peer stops uploading it will tend to be "choked" by other peers, meaning they stop uploading to it. This implements the TFT-like process. So-called "seeder" peers store the whole file: if a swarm contains no seeders, it may lead to a situation in which pieces of the file are missing from the swarm as a whole. Since seeders have nothing left to gain, the system requires some altruistic behavior from peers.

We argue that the TFT strategy is not an adequate explanation of the high levels of cooperation found within BT because: (i) the TFT strategy can be bettered by other less cooperative strategies; (ii) identity can be faked by modifying the client thus circumventing TFT; (iii) unconditional altruism is required for BT to operate in any case. Given that such loopholes exist: why is the system not dominated by freeloaders?

Hypothesis: Group Selection We hypothesize that BT may resist freeloaders and support altruism, at least in part, in a way that has not been previously fully comprehended. Ironically, this process relies on what is commonly believed to be a weakness of BT — the lack of integrated meta-data search. One consequence of this is to partition the BT network into numerous isolated swarms — often with several independent swarms for an identical file — which is one of the necessary conditions for a kind of novel group selective process that has been recently identified in similar simulated systems both in the context of computational sociology and simulated P2P file sharing.

Essentially, if users move between swarms (leave one swarm and enter another) based on the quality of the service they receive, then this process means that swarms containing many freeloaders will tend to "die" as peers leave the swarm for better swarms. Swarms that contain altruists will tend to grow since they support a quality service. Similar models have been advanced in computational sociology [4], [5], [6].

A further implication of the hypothesis is that, given the choice, users may choose unconditional altruism rather than the more restrictive reciprocal approach [7] because of the same group selective process has been shown to select for pure altruism — peers acting for the benefit of the group to their own individual cost.

Conclusions An awareness of some results from social scientific work, particularly the emerging area of computational sociology, can help to inform the analysis of existing working systems "in the wild". This is valuable because such systems currently demonstrate some of the desirable features required of future software systems. Following the empirical line, one way to test our hypothesis would be to implement and distribute a modified BT client that allows users to select pure altruism over the more restrictive reciprocal protocol currently implemented.

REFERENCES

[1] "Bittorrent," official website, http://bittorrent.com/.

[2] R. Axelrod, *The evolution of cooperation.* N.Y.: Basic Books, 1984.

[3] B. Cohen, "Incentives build robustness in bittorrent," in *1st Workshop on the Economics of Peer-2-Peer Systems*, 2003.

[4] D. Hales, "Cooperation without space or memory: Tags, groups and the prisoner's dilemma," in *Multi-Agent-Based Simulation, LNAI*, no. 1979, 2000, pp. 157–166.

[5] R. Riolo, M. D. Cohen, and R. Axelrod, "Cooperation without reciprocity," *Nature*, no. 414, pp. 441–443, 2001.

[6] D. Hales, "From selfish nodes to cooperative networks – emergent link based incentives in peer-to-peer networks," in *P2P2004*, 2004.

[7] R. Trivers, "The evolution of reciprocal altruism," *Q. Rev. Biol*, no. 46, pp. 35–57, 1971.

Simon Patarin and David Hales are with the Department of Computer Science of the University of Bologna, Mura Anteo Zamboni, 40127, Bologna, Italy. e-mail: `patarin@cs.unibo.it`, `dave@davidhales.com`