

# How to cheat BitTorrent and why nobody does

Simon Patarin, David Hales

## I. INTRODUCTION

BitTorrent [1, 2] is currently “king” of the popular file-sharing clients. Some reports have claimed that it accounts for the majority of peer-to-peer traffic on the internet others that Hollywood is doomed. Bram Cohen, the original inventor of BitTorrent, is now under high demand on the invited talks circuit.

However, leaving the media attention aside for one moment, we ask the question “why is BitTorrent so popular”. File-sharing systems come and go, Napster, Gnutella, Kazaa, e-Donkey. What makes one more popular than another? The key to understanding this, we believe, is to understand the quality of the user perceived experience when running such clients. Essentially, if a user feels they are getting what they want from the client then they will tend to stick with it and perhaps even influence others to try it. So what do users want from a client? This, of course, is a very complex question. It is certainly the case that users want quick access to files that they desire but there are also other factors such as ease-of-use and social or psychological needs. For example, the inclusion of a chat option to talk to those with files of common interest may add value for some users. The feeling of being part of an online “community”, often of a counter cultural hue, may satisfy certain social and psychological needs of users even when the impression is often illusory [3].

BitTorrent 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” [4]. BitTorrent works by groups of peers (called swarms) with an interest in downloading a specific file coordinating and cooperating to speed-up the process [2]. 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.

In this paper we argue, that by structuring the entire BitTorrent population into disconnected torrents (or swarms) and leaving the meta-data search outside of the BitTorrent system this may have the side-effect of promoting peer altruism. We argue that it may be this side-effect rather than the TFT strategy that is responsible for discouraging free-riders on BitTorrent.

In the following sections we first discuss the TFT approach in general and highlight its inherent weaknesses. We then advance our hypothesis concerning what might be stopping the exploitation of these weaknesses. Finally we conclude with a general discussion concerning our hypothesis and future work that might test it.

## II. TIT-FOR-TAT IS NOT THE “BEST STRATEGY”

It is generally believed that BitTorrent maintains high levels of cooperation because it uses a protocol based on the tit-for-tat (TFT) game strategy - as championed in Axelrod’s, now classic, book “the evolution of cooperation” [4]. Axelrod held computer tournaments in which

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: simon@patarin.info, dave@davidhales.com

programs, submitted by different researchers, repeatedly played the canonical game of cooperation - the Prisoner's Dilemma (PD).

#### A. *The Prisoner's Dilemma*

In the iterated PD game two players each play a number of rounds. In each round each selects a move from two alternatives (C or D) then the round ends and each player receives a score (or pay-off). Figure 1 shows a so-called 'pay-off matrix' for the game. If both choose the 'cooperate' move then both get a 'reward' — the score R. If both select the 'defect' move they are 'punished' — they get the score P. If one player defects and the other cooperates then the defector gets T (the 'temptation' score), the other getting S (the 'sucker' score). When these pay-offs, which are numbers representing some kind of desirable utility (for example, money), obey the following constraints:  $T > R > P > S$  and  $2R > T + S$  then we say the game represents a Prisoner's Dilemma (PD). When both players cooperate this represents maximising of the collective good but when one player defects and another cooperates this represents a form of free-riding. The defector gains a higher score (the temptation) at the expense of the cooperator (who then becomes the 'sucker').

#### B. *Tit-for-tat does well but can always be bettered*

Axelrod found that in a round-robin tournament the TFT strategy did best on average against the other submitted strategies. The TFT strategy is very simple, it starts by selecting a cooperative move and then for subsequent moves copies the last move made by its opponent. So if a TFT player meets another TFT player they cooperate all the time. However, if TFT meets another player who behaves "selfishly" and defects, then it can be punished with defection in future interactions. The so-called "shadow of the future" can therefore effect behaviour in the present. Cooperation resulting from TFT-like behaviour has been termed "reciprocal altruism" [5] and has been used to ex-

plain various animal and human cooperative behaviour.

However, what is less often mentioned is that Axelrod's work does not prove or demonstrate that TFT is the best strategy or that it can not be bettered by other less cooperative strategies. As stated by Axelrod [4] - *no strategy is best irrespective of the other strategies in the population*. This makes intuitive sense when we appreciate that the performance of a particular strategy in a population is a social phenomena not an individual one. To give an extreme example, if all players in a population always defected and we inserted a single TFT player, then that player would obviously do worse than all the other players since it would initially cooperate with each new player it met and consequently be suckered!

This lesson was starkly demonstrated in the recently organised 20th anniversary IPD computer tournaments [6] in which a strategy based on collusion with other submitted strategies controversially outperformed TFT. By submitting numerous strategies that acted as "masters" and "slaves", using a set of initial game interactions to signal which was which, the Southampton team could outperform cooperative TFT interactions by allowing the "masters" to constantly sucker the "slaves". The slave strategies performed very badly of course, but their sacrifice boosted the "masters" to much higher scores than TFT could ever obtain.

#### C. *The "shadow of the future"*

The TFT strategy works, as we discussed above, by punishing bad behaviour in the future: *if you cheat me today, I will cheat you tomorrow*. If you short-change your regular newsagent today, he might charge you double tomorrow. Given that individuals know there is a good chance of interacting in the future, it is often rational to not cheat in the present. This kind of reasoning has been termed the "shadow of the future" [4].

This appears intuitive in human and animal systems since individuals are incredibly good at recognising individuals in the future. When I

	<b>Cooperate</b>	<b>Defect</b>
<b>Cooperate</b>	R, R	S, T
<b>Defect</b>	T, S	P, P

Fig. 1. A payoff matrix for the two-player single round Prisoner’s Dilemma (PD) game. Given  $T > R > P > S \wedge 2R > T + S$  the Nash equilibrium is for both players to select Defect but both selecting Cooperate would produce higher social and individual returns. However, if either player selects Cooperate they are exposed to Defection by their opponent — hence the dilemma

go into my newsagent I know that I am dealing with the same person I was yesterday, and the chances are so does the shop owner. It’s costly and often impractical to try to pretend to be someone else in face-to-face interactions. However, this is not the case in automated online interactions between many kinds of peer clients over a network.

The technique of fixing a non-fakeable identity is a big issue in security for distributed online systems therefore. Where the stakes are high enough, when security is essential, then trusted third parties can be used to issue certificates of identity that can be verified - this is a mechanism operated by many online systems.

However, within distributed and open systems where certification would be costly or impractical other ad-hoc methods are used to attempt to bind an identity to a peer. Obviously, if these mechanism are violated then it is comparatively easy for one peer to fool other peers into believing that they are a different peer from the one they met before. A single peer can therefore collect multiple identities - often termed a “Sybil Attack” [7] and therefore escape the “shadow of the future”.

#### D. Faking Identity in BitTorrent

In the context of BitTorrent, identity is signalled to the tracker and to other peers using a 20-byte string. A unique identity is generated by the client for each swarm it participates in. This identity is used for every interaction with the tracker and is sent to other peers during the handshake at the beginning of each connection. The tracker, when it returns a list of peers, sends the identity of each of them, in addition to the address and port to connect to.

When a client initiates a connection, it checks whether the identity it receives in the handshake matches the one it has obtained from the tracker. If this is not the case, it drops the connection. The recipient of the connection, however, cannot perform this sort of checking as it has only a limited view of the peers in the swarm. Moreover, trackers do not provide an interface to perform online identity checks.

A client that would want to fake its own identity could do it very easily. As long as trackers do not allow online identity checks (based on IP addresses and ports), it is sufficient to have one identity to interact with the tracker and with the connecting peers, plus one distinct identity for every other peer it connects to. Given the current liberalism of the current implementations it is not even necessary to remember the identity used with a given peer and the (fake) identity could be created randomly for each new outgoing connection. In the eventuality where trackers and (legitimate) clients would become more cautious, it would be necessary to open a new port for each identity used and register to the tracker using these parameters. As long as trackers support clients connecting from behind a firewall, this subterfuge cannot be defeated.

### III. WHY DOES BITTORRENT WORK?

Given that, currently at least, BitTorrent is “king” of the file-sharers and yet, it would seem, is so easy to cheat, why does it work? Also, as stated previously, the system relies to some extent on the pure altruism of “seeders” who have nothing to gain from continuing to serve the file. Why is there so much cooperation going on? Why doesn’t selfish behaviour swamp the system? One obvious answer is that

people are simply more cooperative and altruistic than a worst case kind of economic rationality would suggest - that people don't always act selfishly when they could. This is certainly true and has been demonstrated in experimental settings [8] and identified even in real competitive markets [9]. But there seems little reason why such behaviour would preferentially attach to only the BitTorrent client. This kind of altruistic behaviour would benefit other systems too - like e-Donkey etc. But we know that selfishness in such systems is very high [10]. So why does BitTorrent appear to be doing so well?

As we have discussed above, the claim that it is due solely to the TFT-like interaction protocol can not be the full explanation since it can not account for the pure altruism of seeders, isn't the best strategy in any case and can easily be subverted with a little hacking of the client.

In the following sections we sketch an alternative theory that might explain the phenomena. Firstly, we need to discuss very briefly the way BitTorrent handles meta-data.

#### A. Leaving Meta-Data to the Users

One aspect that place BitTorrent apart from many other file-sharing systems is the way that meta-information concerning content is distributed. BitTorrent, in fact, does not distribute meta-information at all. In order to download a particular file using BitTorrent the user must supply the details of the specific file which are included in a .torrent file. How the user gets this .torrent file is not a concern of the BitTorrent client.

Hence, in order to get .torrent files the user must use other mechanisms - such as the web, where sites exist that run trackers and list active torrent files. There are many such websites (the now defunct Supernova.org became very popular in its day) catering to different tastes and audiences. Some are subscription only, not available to the public, most contain active message boards and appear to support a kind of user community in which people ask favours from others (such as re-seeding a swarm). Users can also, e-mail torrents they have found between

their social networks.

#### B. Torrents as Tribes

This "leave it to the user" approach to meta-data is sometimes considered a weakness of BitTorrent, however, *we hypothesise this is actually a key strength of the system and helps to support altruism and cooperation.*

It does this in two ways. Firstly, by fostering cooperative in-groups of like-minded users with common interests and isolating them, to some extent, from casual users finding meta-data with simple queries - to find and register onto some torrent websites is time-consuming and requires users who know what they are looking for so reducing the casual user (who may be less altruistic). Secondly, and more significantly, each individual torrent swarm is logically isolated from all other torrent swarms, even those sharing the same files running on different trackers.

What this latter point means is this: if a less than cooperative client, a bad guy, enters a swarm then that swarm will, on the whole, perform less well than other swarms composed of all good-guys. It is true that the bad guy will do well at a cost to the other swarm members. However, users may decide to manually leave the swarm if they feel they are not getting a good enough performance from it. Almost all BitTorrent clients show the user the upload vs. download ratios for the swarms they are members of thus allowing them to monitor the quality of the service they are receiving from the swarm. Anecdotally, it is not infrequent for users to remove themselves from swarms that deliver a poor ratio.

This latter process, if indeed practised often enough by enough users, would lead to a kind of user driven group selection process. This process has already been observed and simulated in the context of human and peer-to-peer systems [11–14]. We categorise such systems as "tribal systems" since they promote cooperation by the dynamic formation and dissolution of "tribes" or groups of peers. By partitioning the population into a finite set of such tribes, those tribes that offer better performance

to their members tend to prosper at the expense of those filled with free-riders or bad guys.

The interesting thing about this process is that it can be driven by purely selfish and greedy behaviour at the peer level but produce high levels of altruism at the “tribe level”.

The best way to think of this is as a kind of recruitment process in which peers wish to increase their own benefit from the system. Assuming peers can move between “tribes” (in our case, exit and enter swarms) they will tend to be attracted to those tribes that can offer them more benefit (in this case, download speed). We would expect this to be impaired by free-loaders. Therefore tribes with fewer free-loaders tend to grow whereas tribes with many tend to shrink - eventually “dying” if all peers leave it. This process has been observed to occur in a number of simulation models in which peers behave in a selfish and greedy way [14]. It is a kind of group selection process which supports high levels of altruistic behaviour and is robust to invasion by free-riders.

We argue that the current architecture of BitTorrent appears to provide the right selective environment for this process to occur. This may offer a more plausible explanation for the high levels of altruism that appear within the system at present.

### *C. Implications and Hypotheses*

If we are correct, this view allows us to make a number of implications or hypotheses that might appear counter-intuitive.

Firstly, if peer users were given the freedom to switch their clients between pure altruism and pure selfishness this may actually improve performance of the system since pure altruism would predominate. In the context of the choking algorithm this would mean the implementation of both a more altruistic sharing rule and a less altruistic rule than the TFT-like one. The group selective process should then select highly altruistic swarms since peers would leave ones full of selfish free-riders.

Secondly, if meta-search systems became fused into the BitTorrent protocol this might

have a deleterious effect on altruism and co-operation. Why? Because by allowing all peer clients to quickly locate and simultaneously participate in many related swarms at the same time, the group selective process becomes weakened since the swarms effectively become fused into one large swarm. Additionally, this removal of the need for communities to be created around trackers, may remove some of the psychological advantages.

In both these cases, it would appear that there are so many unknowns that the only way to test these hypotheses would be via empirical experimentation in the “wild” of the BitTorrent ecology. This, of course, is not impossible since BitTorrent is an open protocol. Modified clients would need to be released and gain some reasonable level of take-up. The data would need to be collected from both clients and trackers.

It would appear that a modified client that tested the first hypothesis would be practical to implement since it requires only modest changes to the choking and peer identification algorithms. This may be the subject of future work.

Testing the second hypothesis would be a much more complex task. However the evolution of the BitTorrent protocol may come to test this for us since the last versions of the official client now integrate meta-data search.

There are of course certain kinds of “ethical” issues in experimenting in this way. If we were wrong in our first hypothesis then clients capable of selfish cheating behaviour might come to dominate the system and destroy its usefulness - though we think is very unlikely.

## IV. CONCLUSION

We have argued that the success of BitTorrent is unlikely to be due purely to the use of the tit-for-tat inspired protocol, as is often claimed. We argue that the real driving force behind the high cooperation might be the by-product of the lack of meta-data search within BitTorrent. This results in the creation of a number of disconnected “tribes” at both the swarm and the tracker level. The users are active in the tribal

dynamics by selecting those tribes that best satisfy their needs hence tribes filled with free-riders will tend to die out. We compare this to existing simulations of tribal dynamics in both human and peer-to-peer systems.

We advance a number of hypothesis which our theory suggests including the production of both cheating peers and unconditionally altruistic peers. We argue that releasing such peers into the “wild” of the BitTorrent ecology would not damage BitTorrent but could actually increase the system level performance because unconditional altruism would tend to be selected and predominate.

In order to fully test these hypotheses we would need to construct and release such clients into the “wild” and collect data from them. This may be the subject of future work.

It’s a sobering thought to consider that, perhaps, the most bandwidth hungry applications on the internet today work by complex social mechanisms we don’t yet understand. However, this is less sobering when one realises that since peer-to-peer systems are really just computationally supported human social systems then we should expect the same issues to arise as we observe within human social systems - namely the question “what’s going on?”.

#### ACKNOWLEDGEMENTS

This work would not have been possible without perceptive discussions with many people particularly those in the Bologna group including: Mark Jelasity, Ozalp Babaoglu and Alberto Montresor. This work partially supported by the EU within the 6th Framework Program under contract 001907 (DELIS).

#### REFERENCES

- [1] “Bittorrent,” official website, <http://bittorrent.com/>.
- [2] B. Cohen, “Incentives build robustness in bittorrent,” in *1st Workshop on the Economics of Peer-2-Peer Systems*, 2003.
- [3] L. Strahilevitz, “Charismatic code, social norms, and the emergence of cooperation on the file-swapping networks,” *Virginia Law Review*, vol. 89, 2003.
- [4] R. Axelrod, *The evolution of cooperation*. N.Y.: Basic Books, 1984.

- [5] R. Trivers, “The evolution of reciprocal altruism,” *Q. Rev. Biol.*, no. 46, pp. 35–57, 1971.
- [6] W. M. Grossman, “New tack wins prisoner’s dilemma,” *Wired Magazine*, Oct. 13, 2004.
- [7] J. Douceur, “The sybil attack,” in *Proc. of the IPTPS02 Workshop*, Cambridge, MA (USA), Mar. 2002.
- [8] W. Poundstone, *Prisoner’s Dilemma*. N.Y.: Doubleday, 1992.
- [9] A. P. Kirman and N. J. Vriend, “Evolving market structure: An ACE model of price dispersion and loyalty,” *Journal of Economic Dynamics and Control*, vol. 25, no. 3/4, pp. 459–502, 2001.
- [10] E. Adar and B. Huberman, “Free riding on gnutella,” *First Monday*, vol. 5, no. 10, 2000.
- [11] 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.
- [12] R. Riolo, M. D. Cohen, and R. Axelrod, “Cooperation without reciprocity,” *Nature*, no. 414, pp. 441–443, 2001.
- [13] K. Sigmund and M. Nowak, “Tides of tolerance,” *Nature*, no. 414, pp. 403–405, 2001.
- [14] D. Hales, “From selfish nodes to cooperative networks – emergent link based incentives in peer-to-peer networks,” in *P2P2004*, 2004.