

## How to cheat BitTorrent and why nobody does

Simon Patarin and David Hales  
University of Bologna

ECCS 2006, Oxford, UK



Information Society  
Technologies

- BitTorrent is a peer-to-peer (P2P) file sharing system
- Widely used (10's of millions of users)
- Relies on client programs running on user machines to cooperate to spread files fast
- Clients cooperate by uploading pieces of a file to others as well as download from them

## Why Does it Work?

- But why do clients cooperate when their are incentives to act selfishly?
- It is widely assumed that BT maintains cooperation via its tit-for-tat (TFT) method
- This means that clients only upload to others who download
- TFT has been demonstrated to produce cooperation in abstract games such as the Prisoner's Dilemma (PD)

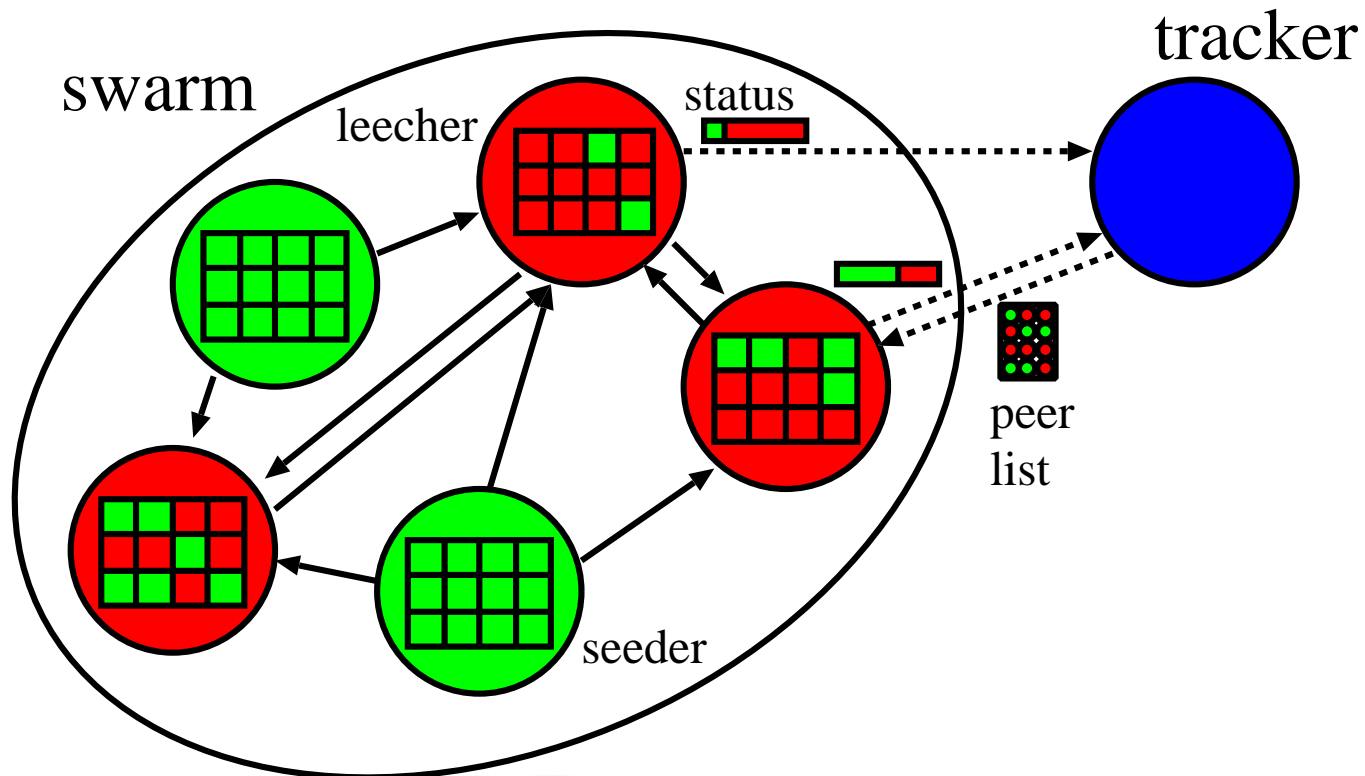
## But is this true?

- But the BT protocol can be easily exploited
- We propose another possible explanation for cooperation
- Based on new ideas from computational sociology - group selection mechanism
- If this hypothesis is true it could have implications for future protocol design

- Outline of BitTorrent Protocol
  - How it works
  - How to cheat
- Possible mechanisms other than tit-for-tat that prevent cheating
- Outline of group-selection hypothesis
- Conclusions / Future work

- Torrent file: meta-data about the data shared
  - Data is split in smaller pieces called chunks
- Tracker: centralized manager
  - Keep track of all peers in the swarm
  - Keep track of the status (up/down ratio) of each peer
- Peers: clients
  - Seeder: holds a full copy of the data
  - Leecher: holds only a part of the data (initially nothing)

- Get a list of other peers in the swarm from the tracker
- Ask peers their list of chunks and tell them what is yours



- Limited number of peers you can interact with simultaneously
- **Tit-for-tat**: upload data to the peers that let you download most (*choking*) and stop cooperating with peers that don't upload to you (*anti-snubbing*)
- Also cooperate with peers picked at random (*optimistic unchoking*)
  - Newcomers get 3x more chance to be picked in this phase
- Seeder mode: upload to those that download the fastest (spread the data faster)

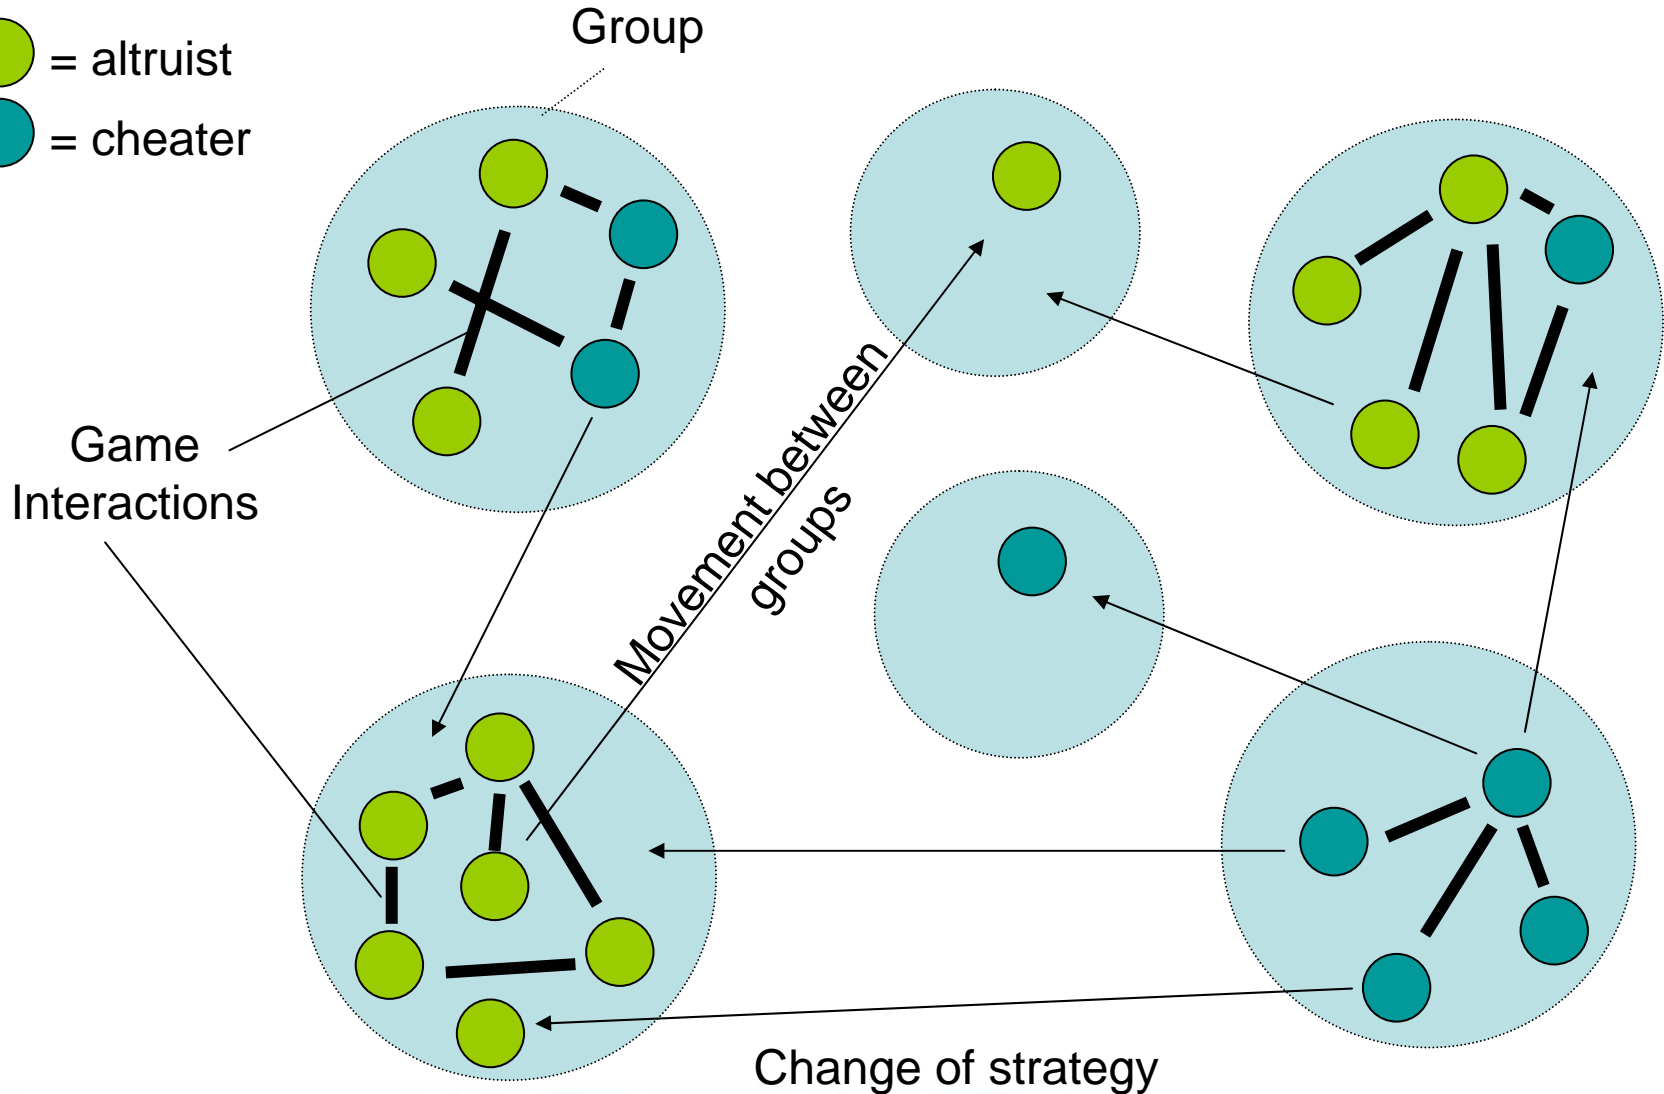


- Identifiers dynamically generated by the client
- Identifier provided in each communication with the tracker
- Identity controlled when connecting to a peer (as given by the tracker in the peer list)
  - No control possible when receiving a connection (the whole swarm is not known)
- Easy to create fake identities
  - Keep a 'public' identity when communicating with the tracker and receiving connections
  - Create new identities when connecting to other peers

- Hacked clients exist but BT still works well enough to be widely used, why?
  - BT community have a niceness “norm”
  - Uploading cost is low due to flat-rate cost of most internet connections
  - “Closed trackers” which use IP for “persistent identity” and “ratio” accounting
  - *A novel form of “group selection” based on clients “moving between” swarms*

- Recent evolutionary “tag” models of altruism (Holland, Riolo, Hales)
  - Structure population into dynamic groups
  - Individuals move between groups to maximise their payoffs
  - Groups containing cheaters tend to die out
  - Groups containing altruists grow
  - Even when individuals can behave selfishly

● = altruist  
● = cheater



- If the groups are seen as BT swarms
- Clients move from swarms offering bad performance to those offering good
- Then, swarms containing cheaters will tend to “die out” as altruist move away
- A swarm containing all cheaters is useless even to cheaters!

- Are such dynamics driving BT altruism?
- We currently don't know
- Potentially experiments measuring what BT clients do could be performed
- Implication if true: P2P protocols can rely on unconditional altruism so long as a healthy “group ecology” is maintained
- We have already modelled such protocols

- Hales, D. & Patarin, S. (2005) How to cheat BitTorrent and why nobody does. University of Bologna, Dept. of Computer Science, May 2005, Technical Report UBLCS-2005-12.
- Hales, D. (2000) Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. In Moss, S., Davidsson, P. (Eds.) Multi-Agent-Based Simulation. Springer-Verlag.
- Hales, D. & Arteconi, S. (2006) SLACER: A Self-Organizing Protocol for Coordination in P2P Networks, IEEE Int. Sys. 21(2):29-35.
- Riolo, R., Cohen, M. D. & Axelrod, R. (2001) Cooperation without reciprocity. Nature, 414:441–443.
- see: [www.patarin.info](http://www.patarin.info), [www.davidhales.com](http://www.davidhales.com)