CS698W: Game Theory and Collective Choice **Jul-Nov 2017**

Lecture 11: August 25, 2017

Lecturer: Swaprava Nath *Scribe(s): Ameya Loya*

11.1 Game Theory Ecosystem

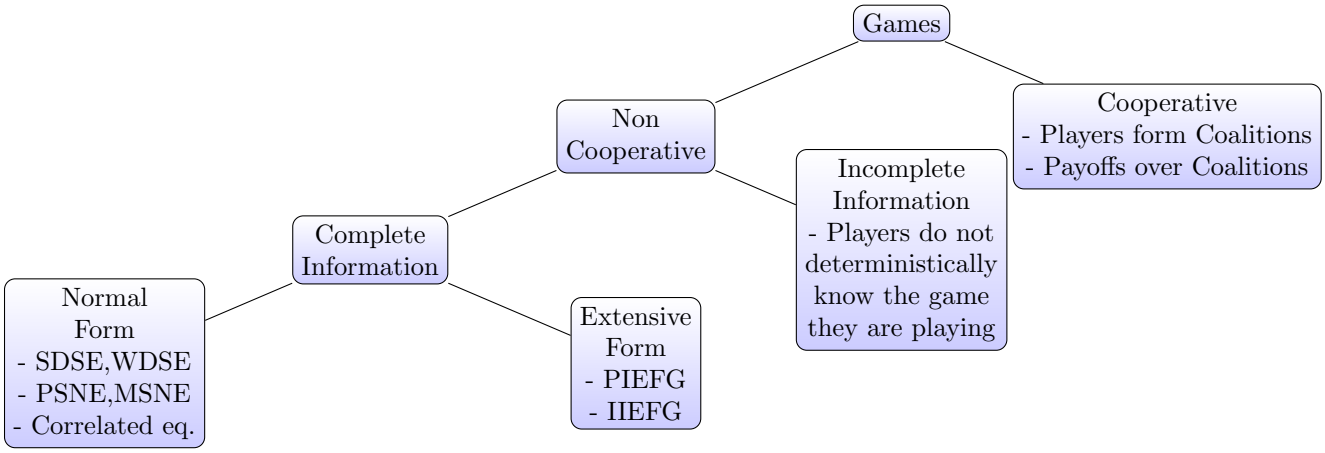

Things not covered: Repeated Games, Stochastic Games,

11.2 Game Theory in Practice: Peer to Peer File Sharing

Conventional server-client models use single server to which every client is connected. However, this is inefficient since the download speed is bottlenecked by the server's bandwidth – while there is enough download/upload bandwidth with the clients that is unused. Peer-to-Peer (denoted P2P) is a technology to mitigate this inefficiency. It is a completely decentralized network in which each client is capable of acting as a server for another, and helps keep the download speed unaffected by the number of users (Figure 11.1). The advantages of P2P is that it is (a) highly scalable, and (b) resilient to server failures.

11.2.1 Basic Terminology

1. **Protocol**: Messages that can be sent, actions that can be taken over the network.
2. **Client**: A particular process for sending messages, taking actions.
3. **Reference Client**: A particular implementation of P2P.

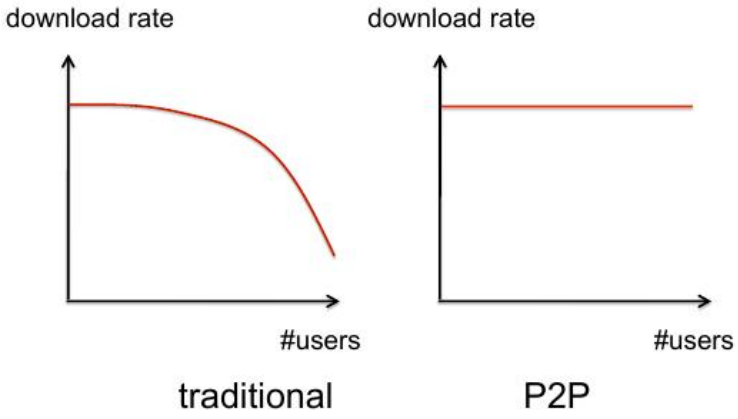

Figure 11.1: P2P innovation

11.3 Implementations of P2P Protocols

Early implementations of P2P technology was Napster and Gnutella.

Napster (1999 - 2001)

1. Centralized database
2. Users download music from each other

Gnutella (2000 -)

1. Get list of IP addresses of peers from set of known peers (no server)
2. To get a file: Query message broadcast by peer A to known peers
3. Query response: sent by B if B has the desired file (routed back to requestor)
4. A can then download directly from B

In both these protocols, the downloader's strategies were not taken into account. If we consider that data upload is costly and every agent gets a positive utility if he downloads some file that is useful to him, then the upload-download dynamics can be represented as a normal form game – where every agent can either upload and download (Share) or never upload but only download (Free Ride). The setup is represented in Table 11.1. Clearly, the incentives with these two naïve protocols lead to scenario where free-riding is

		Person 2	
		Share	Free Ride
Person 1	Share	$2, 2$	$-1, 3$
	Free Ride	$3, -1$	$0, 0$

Table 11.1: A P2P file sharing game

a dominant strategy for both players (attested by the study of Adar and Huberman (2000), Figure 11.2). In Gnutella, 85% peers were observed free-riding by 2005; it had less than 1% of worldwide P2P traffic by 2013. Few other P2P systems met similar fate. Therefore it is unlikely that peers will share files with such protocols – hence a client with incentivizing properties was needed.

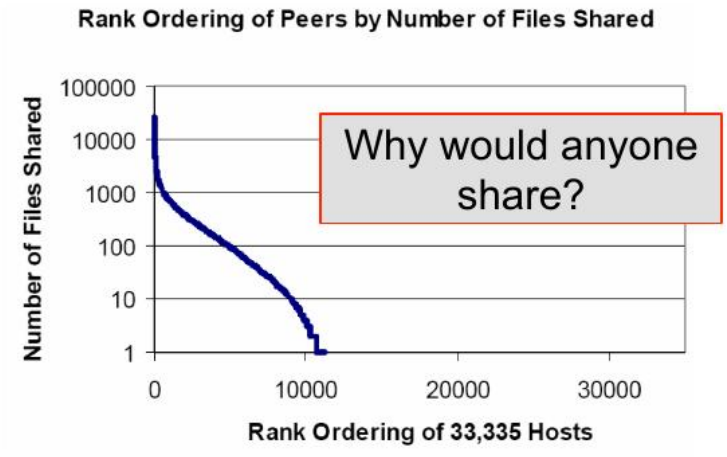

Figure 11.2: Sharing statistics. Image courtesy: Adar and Huberman (2000).

BitTorrent Protocol (2001 -)

BitTorrent was the first protocol that broke this deadlock with incentives for uploading. It breaks the file into multiple pieces and different pieces are delivered to different agents. Among the agents the exchange of the pieces follow a repeated game and the strategy that is followed is an adaptation of the tit-for-tat policy for repeated prisoner's dilemma game. The principle was:

> "If you let me download, Ill reciprocate."

The schematic of BitTorrent is shown in Figure 11.3.

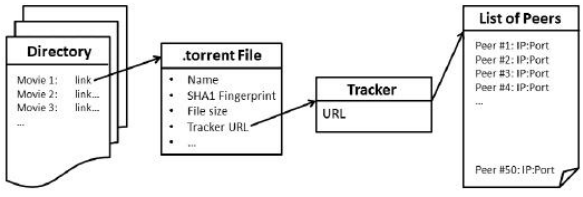

Figure 5.4.: Starting a download process in the BitTorrent protocol: 1) A user goes to a searchable directory to find a link to a .torrent file corresponding to the desired content; 2) the .torrent file contains metadata about the content, in particular the URL of a tracker; 3) the tracker provides a list of peers participating in the swarm for the content (i.e., their IP address and port); 4) the user's BitTorrent client can now contact all these peers and download content.

Figure 11.3: BitTorrent details. Image courtesy: Parkes and Seuken (2017).

BitTorrent Optimistic Unchoking Algorithm Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past).
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period.
- If a peer j uploaded to i at a rate $< r$, choke j in the next period.
- Every three time periods, optimistically unchoke a random peer from the neighborhood who is currently choked, and leave that peer unchoked for three time periods.

The protocol is forcing file sharing to be a repeated game by fragmenting the files, which is a repeated Prisoners' Dilemma. Strategy of the seeder (the reference client – the one that is the default BitTorrent client) is tit-for-tat (TfT). A TfT strategy for a prisoner's dilemma starts with cooperate action, and if the opponent defects at any stage, it keeps of defecting from the next stage until the opponent plays cooperate again.

Though BitTorrent is extremely popular (Approx 85% of P2P traffic in US is BitTorrent), it does have some strategic vulnurebilities.

Attacks on BitTorrent With the BitTorrent protocol, an adversary may consider the following actions to decide how this can be gamed.

- How often to contact tracker?
- Which pieces to reveal?
- How many upload slots, which peers to unchoke, at what speed?
- What data to allow others to download?
- Possible goals: minimize upload, maximize download speed.

11.3.1 BitThief

BitThief does not perform any chokes or unchokes of remote peers, and it never announces any pieces. In other words, a remote peer always assumes that it interacts with a newly arrived peer that has just started downloading. Thus, a BitThief client is able to download without ever uploading any time. However, if it stays in one neighborhood for long, the other peers figure out that it is never uploading and will choke it. Therefore, the client asks the tracker more frequently for new peers and grows the neighborhood rapidly.

This protocol seems to have an easy fix. One can modify the tracker to block such peers that asks for new peers more rapidly.

Ref: Locher et al., "Free Riding in BitTorrent is Cheap", HotNets 2006

11.3.2 Strategic Piece Revealer

This is another attack on BitTorrent that picks the piece of the file to share in order to maximize its importance in its neighborhood. Peer A is 'interested' in peer B if B has a piece that A does not. Strategic Piece Revealer (SPR) client strategically shares only those files that are the most common piece the reciprocating peer does not have and saves its own 'rare' pieces. The BitTorrent reference client always use "rarest-first" to request, i.e., it will announce the rarest piece first and share that if requested. Threfore, by sharing the most common piece an SPR client will protect its monopoly over the rare pieces and keep others interested. The improvement in terms of the reduction in download time is shown in Figure 11.4.

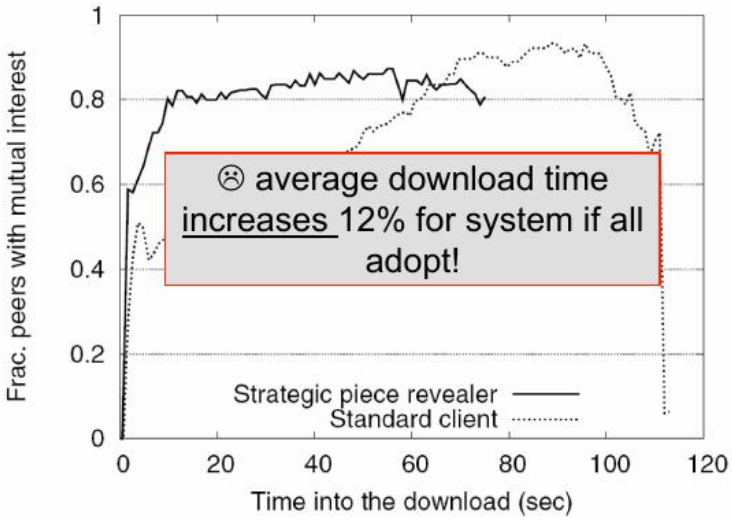

Figure 11.4: Strategic Piece Revealer.

Ref: Levin et al., "BitTorrent is an Auction: Analyzing and Improving BitTorrents Incentives", SIGCOMM 2008

11.4 Summary

- P2P demonstrates importance of game-theory in computer systems.
- Early systems were easily manipulated.
- BitTorrent's innovation was to break files into pieces, enabling TitForTat.
- Still some vulnerabilities, but generally BitTorrent is a very successful example of incentive-based protocol design.

Notes

This lecture has been adapted from the relevant lecture notes of CS186 Harvard (Instructor: David C. Parkes) and the content and images are from the unpublished book of Parkes and Seuken, "Economics and Computation", (2017).