Peer-to-Peer Networks: An Evolutionary Game Theoretic Approach

CS532a Multiagent Systems Final Project Report

Andrew Adam
University of British Columbia
Vancouver
BC, Canada
andyadam@cs.ubc.ca

ABSTRACT
I consider the problems of peer-to-peer file sharing networks, such as *freeriding*, *whitewashing* and *traitors*. In respect to this I attempt to analyse ideas from Evolutionary Game Theory and apply them here, drawing upon recent work from Feldman et al.[2][3]. Evolutionary Game Theory is applicable since it deals with strategies in large populations.

Keywords
Research, Multiagent Systems, Game Theory, Evolutionary Game Theory, Network Design, Peer-to-Peer

1. INTRODUCTION
Peer-to-Peer (P2P) file sharing systems allow computer users to share files with each other around the world without ever meeting or interacting with the people they are sharing with. This has been highly publicised in recent years first for those who share music files and now those who share movies, thanks to increases in network bandwidth. An example of this would be Napster, shut down for illegal copyright issues and now a legal download site, or more recently BitTorrent (which uses a central server called a tracker which coordinates the actions of the peers, and seeds that at any time have a complete copy of the file to upload.)

In these systems the file service is performed by the systems users. In order to persuade users to upload their files to other users instead of just downloading many have to be given some sort of incentive. Some systems require a high reputation to maximise download potential, and a high reputation is achieved by the amount of uploading that the user has done. However with a lot of systems this can be circumvented by hacking the client program to stop uploading, or in cases where the option is voluntary, then the user chooses to switch the option off. This is known as the *free-rider* problem and was studied before P2P systems.

Economics and Game Theory can be applied to this problem as it is one of shared resources and payoffs to users. Game Theory has emerged in the last 60 or so years as a powerful challenger to the conventional methods of studying economics. The main purpose is instead of considering situations where agents are making decisions based on reactions to prices (dead variables), to think about strategic decisions based on reactions to other agents actions (live variables). A strategy can either be pure, as in to always play a particular move, or mixed, to play different moves over some probability distribution.

Various equilibrium concepts can arise from this analysis, such as the Nash Equilibrium, which is reached when each agents action is a best response to every other agents actions. These concepts are important when analyzing games, finding a stable point within them where all agents are happy and do not wish to deviate from their current strategy.

Evolutionary Game Theory has been around since the 1970s when John Maynard-Smith [5] and George R. Price related the concepts of Nash Equilibria in an evolutionary environment. A strategy is an Evolutionarily Stable Strategy (ESS) if a population of individuals homogenously playing this strategy is able to outperform and eliminate a small amount of any mutant strategy introduced into the population.

There are a few principle examples of how agents react with each other in large populations. [1] Firstly we have the Prisoners Dilemma. According to rational reasoning an individual agent should always defect against the agent he is playing, as it will result in a better payoff for them. However, if the game is repeated a strategy can be adopted, such as Tit-for-Tat, where the fear of retaliation in the future causes agents to retaliate in the present. In well-mixed populations, the fitness of agents can be measured the average payoff over a certain number of iterations, and in this case the co-operators are doomed.

Another example is the Snowdrift (Hawk-Dove) game. In this game two agents can either stay in the warm car, or shovel the snowdrift. If either or both shovels then they both get home, else they wait till spring the snow to thaw. This is different to Prisoners Dilemma as this time each agents best action depends on the other agents behaviour. In this case co-operators and defectors can co-exist.

I will give an introduction to Evolutionary Game Theory and then analyse how this could be applied to peer-to-peer systems.

2. BACKGROUND

[4] Golle et al. begin by characterizing the problem in a more formal game-theoretic manner. The file sharing scenario is modeled by a game that takes place within a defined time period. n agents participate in the system and are denoted $a_1, ..., a_n$, and each agent a_is strategy, denoted $S_i = (\sigma, \delta)$, consists of 2 independent actions:

1. **Sharing:** Agents select what proportion of files to share. Sharing takes three levels: σ_0 (none), σ_1 (moderate) or σ_2 (heavy).
2. **Downloading:** Each agent must also determine how much to download the network in each period. Models are downloaded with agents choosing between three levels: δ_0 (none), δ_1 (moderate) or δ_2 (heavy).

Each agent has a utility function which is based on a number of factors, such as Amount Downloaded (agents get happier the more they download) and Bandwidth Used (a cost to agents associated with uploading files to the network), there are others. Each agent also has a quasilinear utility function, meaning that each of the factors above can be mapped onto a dollar value for that agent, be it positive or negative, creating a total utility value for that agent. Agents are simplified to have the same utility function for each of the five factors described, so that agents can be considered to have the same *type*, and all agents are deemed to be economically rational.

The idea is to find some sort of equilibrium which satisfies the goals that individual agents want to achieve whilst also satisfying the group of agents in a maximal way. The joint strategies of all agents is denoted $\Sigma = \{S_1...S_n\}$, and for example, A *Nash Equilibrium* over these strategies would occur if every agent would be strictly worse off should he change his strategy.

Several alternative mechanisms are discussed:

1. **Micro-Payment Mechanisms**: This approach attempts to balance the system with what the agents take and contribute, by rewarding the user for every upload and charging for every download. The number of files a user downloads δ is tracked by the central server and also the number of files uploaded v. At the end of each period, each user is charged an amount $C = g(\delta - v)$. However the paper suggests that users dislike micro payments as it forces each download to be an individual buying decision.
2. **Quantized Micro-Payment Mechanisms**: This is similar to the above except that the user purchases a block b of files. However there are problems in this model where agents could gain via collusion, i.e. downloading between themselves for profit.
3. **Rewards for Sharing**: This approach continues to penalize downloads but rewards agents based on the proportion of material they have available for sharing rather than the number of uploads they provide. This can be done via a point system, where a user gets points by purchasing them with money or by providing contributions to the network. However the paper states that the drawback with point systems is that they result in a degenerate equilibria in which all agents download at the highest level and share nothing at all (although as shown by the original Napster this does not necessarily mean that it will happen.) Also, it is difficult to give points for files that are available, when an agent may attempt to cheat by claiming to be unavailable as soon as another agent chooses to download from them.

Feldman et al. [2] discusses the problem of free-riding, and attempts to solve it through the use of Generalized Prisoners Dilemma (GPD). It states that P2P systems impose a unique set of challenges, such as high populations (100000+) with high turnover, asymmetry of interest (each user wants something from someone else in a cycle), and zero-cost identity (the ability to change identity at no cost.) To address this they introduce scalable and robust techniques, and a new decision function they call Reciprocative. They have five different tradeoffs:

1. Discriminating Server Selection: Cooperation requires familiarity between entities, but the high turnover and populations make this unlikely, so each peer will keep a private history of actions, and discriminating server selection will be used.
2. Shared History: This results in a higher level cooperation than private history, but the cost is a distributed infrastructure to store it.
3. Maxflow-based Subjective Recognition: Shared history makes it possible to collude. The paper shows that a maxflow algorithm that computes reputation subjectively promotes cooperation in spite of collusion between a third of the population.
4. Adaptive Stranger Policy: Zero-cost identities can be tackled by treating strangers based on the history of other strangers.
5. Short-term History: History means that a previously well behaved peer could turn traitor and exploit other peers, so histories are kept short-term.

GPD provides the general form for an asymmetric payoff matrix that preserves the dilemma of the problem. It has the properties that mutual cooperation leads to higher payoffs than mutual defection or one player defecting, but defection weakly dominates at the individual level for the player who decides whether to cooperate or defect.

In each round each agent plays one game as a client and one as a server, and at the end of each round the player may either mutate to a new strategy, learn, turnover (leave the system), or stay the same. If a player mutates then they switch to a randomly picked strategy. If they learn, they switches to a strategy they believe will bring a higher score, and if they keep the same name they will be identified as a traitor. If a player suffers turnover, they leave the system and are replaced by a newcomer who uses the same strategy as the exiting player.

Each agent has to decide how to act and this is where the Reciprocative function comes into play. A simple decision might be to 100% cooperate, a naive choice by a user who does not realise they are being exploited. The probability that a Reciprocative player cooperates with a peer is a function of its normalized generosity, i.e. the benefit an agent has provided compared to what it has consumed. It uses its own generosity when judging other agents generosity to ensure Reciprocative agents do not defect on each other. There are further details of the incentive schemes discussed, and also types of cheating or attacks that may occur. To conclude the paper states that it tackled the problems of P2P systems and that its Reciprocative decision function results in cooperative behaviour and improves overall system performance.

They conclude that their Reciprocative function supports cooperative behaviour and improves overall system performance, and that the usage of Shared History and Discriminating Server Selection can address the challenges of large population size, high turnover and asymmetry of interest.

Feldman et al. [3] also consider the problem of *whitewashing*, where users leave the system and rejoin with new identities in order to avoid reputational penalties. Each agent has a *type*. When *generosity* (the average type) is low then intervention is required to keep the system running. Each decides whether to contribute or free-ride based on the relationship between the cost of contribution and their type. If x is the fraction of users who contribute, then a rational user of type t_i will contribute if:

$$1/x < t_i$$

With too many non-contributors the system will collapse. They go on to model a mechanism for the penalty imposed on newcomers which only degrades performance when turnover is high.

Wang and Li [9] attempt to use a Cournot Oligopoly for the same problem, and Ranganathan et al. [8] use a Multi-Person Prisoners Dilemma approach. Nowak and Sigmund [6] use Evolutionary Game Theory whilst scoring images.

3. EVOLUTIONARY GAME THEORY

Evolutionary Game Theory can be described, at least at a basic level, as a series of 1-on-1 encounters between players in a game. Maynard Smith termed this as *pairwise competition*. The players in the game become the strategies themselves, and then the game considers the overall fitness effects on each strategy after all possible contests are played, in proportion to their likelihood. Each of these strategies are attempts at modelling behaviour in biological populations but the ideas can also be used here.

There can be contests between individuals using the same behavioural strategy or between individuals of different strategies. We need to work out the *fitness* of each actor in the game. To do this we need to include the value of the resource, the chances of winning the resource, the chance of a loss and the costs of winning or losing. Benefits are assigned by using a relative but yet arbitrary scale of value, a *common currency.*

From this payoffs have the general form:

*Payoff(to 1 vs. 2) = chance of win * (resource value - cost of win) + chance of loss * cost of loss*

However fitness also depends on the frequency of other behaviours in the game. If certain behaviour is infrequent within the game then it is less likely that another agent will interact with that type of behaviour, so that behaviour has less impact on the overall fitness for that behaviour.

If fitness is denoted by W and expected fitness by E, overall fitness to a particular strategist in a particular type of contest is

*Change in W(S1) = E(to S1 vs S2) * freq(encounter)*

3.1 Evolutionary Stable Strategies

Game Theory can be used to predict situations where:

1. one behaviour is more fit than all known alternatives[1]
2. a specific mix of behaviours where none are more fit than the other

This results in *evolutionary stasis*; there is no change in relative frequency of strategies over time, it is an *evolutionary stable strategy (ESS).*

1. **Pure ESS**: where one strategy totally out competes all the others. Regardless of its frequency it is always more fit than any other known alternative, and it is immune to invasion from other known strategies. Thus any alternative that appears by mutation or immigration will not be able to increase and will become extinct.
2. **Mixed ESS**: where two strategies permanently coexist. This could be achieved by individuals playing one strategy all the time where the two strategies are at equilibrial frequencies. Alternatively, it could be caused by everyone playing a mixed strategy where each of the behaviours is performed at the equilibrial frequencies.

Let us consider an example. Figure 1 shows the payoffs of two strategies, A and B. The population begins entirely of individuals who use strategy A. If a is the frequency of strategy A and b is the frequency of strategy B then

- A vs A occurs at the frequency a^2
- A vs B occurs at the frequency $2ab$

[1] When a strategy is said to be *uninvadable*, it is with relation to *known* strategies. It is potentially vulnerable to a new strategy that may come along.

		Opponent Strategy	
		A	B
Focal Strategy	A	E(A,A) = 0	E(A,B) = 1
	B	E(B,A) = -0.5	E(B,B) = 0.5

Figure 1: Payoffs for A versus B strategies.

- B vs B occurs at the frequency b^2

So from figure 1 we can see that when A plays A the payoff is 0. If a single invader of type B appears, all of his interactions will be with players of strategy A and the payoff will be of -0.5 (i.e. the B strategist will lose fitness as a result of this interaction.) So in this case B cannot invade A. This can be formalized as

$$E(A, A) > E(B, A)$$

which means that A is *stable* against B, a *Pure ESS* vs B. If this rule was not true but

$$E(A, A) = E(B, A) \text{ and } E(A, B) > E(B, B)$$

is then it would also be a Pure ESS. We can also have Mixed ESS, where two strategies have an intermediate point where both have the same fitness.

3.2 Hawks and Doves

The Hawks and Doves game is a classic introduction to Evolutionary Game Theory, like the Prisoners Dilemma is to traditional Game Theory. **Hawk** and **Dove** are simplified behavioural strategies. A Hawk against a Hawk is a very aggressive affair, they will fight for the resource and the loser is injured. Doves never fight, but if they face another Dove they display their feathers and each has an equal chance of winning the resource. This can be explored further in many Evolutionary Game Theory texts and originally comes from *Evolution and the Theory of Games* [5].

A brief analysis shows that Hawks do extremely well when it is rare. For example, with only one Hawk against a population of Doves, the Hawk always wins, but as Hawks increase they begin to do very poorly. The payoffs to each strategy can be drawn on a graph based on their payoffs, and where they intersect is a Mixed ESS. Neither is a Pure ESS, as each population can be invaded by a small number of the other strategy.

3.3 Further Strategies

The **Bourgeois** strategy is a possible addition to the Hawks and Doves game. Previously we did not consider that a resource may have already be owned, and that neither player was the current owner. We can introduce the strategy Bourgeois, a strategy associated with *respect for ownership*. The Bourgeois strategy fights like a Hawk when it has ownership of a resource, but displays like a Dove when attempting to gain the resource.

Fixed Cost Strategies can also be used. However it can be proven [7] that none of these strategies are Pure ESS's as a strategy where someone pays more than the current fixed cost can invade, but also there are an infinite number of Mixed ESS.

It is possible for different strategies to co-exist in a system, without an ESS. In a *disequilibrium* one strategy is more fit but there has not yet been sufficient time to reach equilibrium. If the *environment changes*, favouring one strategy then the other, then an ESS may not be reached. This cyclical variation is apparently common in nature [7]. It is also possible that a Pure ESS strategy might not be the best choice to an agent in a certain condition. If an individual is not fit, and performing the ESS strategy is costly, then it may wait until it is fitter. This type of behaviour is called *coping*. Even if both strategies persist it is not a Mixed ESS.

4. PEER-TO-PEER NETWORKS AND EGT

What do we need to consider when building a peer-to-peer system? Each users attempt to maximise their own utility effectively lowers the utility of the whole system. We need to model the different types of the players in the game. What actions can the players have at any point? Each user can either stay the same, change to a new type, or leave the system.

We need to consider how to store and deal with history (local/global), how peers choose who to interact with, how to do deals with traitors and other unhelpful users, those who collude, those who change identity to avoid reputation penalties and how to reward contributors.

In EGT we keep track of the frequencies of each type of behaviour. These models specify the global behaviour of a system. It will tell us how many users adopt a strategy at any time, and how that strategy evolves over time. However we need to keep track of how individual entities act with each other and also which individuals switched strategies, rather than just the frequencies themselves. Newcomers to the system may be treated in a different way, even if they play the same strategy as those who cooperate, to try and combat the problem of zero-cost identities (changing identity to avoid reputation penalties.)

Analysis of the interactions between the players shows that a Hawk/Dove model may be inappropriate in this case. An invader to the system could be a Hawk entering a realm of cooperating Doves (who don't face off against each other for the resource), and so if the Hawk could invade then it would work in this respect. But also it is unlikely that a Hawk/Hawk face-off would result in negative utility to one of the non-co-operators (one of their computers in damaged??) So perhaps it makes more sense to think about large populations playing Prisoners Dilemma against each other, where those who fail to cooperate receive a greater payoff and that is their dilemma. Therefore we can use a payoff matrix as modelled in [2], except more generic (see Figure 2).

In this case each user alternates between being a client and a

		Server	
		Provide Service	Ignore Request
Client	Request Service	g / h	0 / 0
	Don't Request	0 / 0	0 / 0

Figure 2: Possible payoff matrix for a peer-to-peer system.

server in the game. The only meaningful decision is made by the Server, who decides whether or not to cooperate or not. Typical scores might be o = 5 and p = -1. It is clear that if the client doesn't request anything then the payoff for both should be zero. However If the Server chooses to cooperate then the client will receive x (presumably positive) and the server will receive y (presumably negative). It is important to consider whether the Client player is aware of the Server's action when he defects, as this may effect how he is treated in the system later.

How can we model the performance of the system? The benefit received by a user of the system can be given by: [3]

$$Q = \alpha x^{\beta}$$

where $\beta \leqslant 1$ and $\alpha > 0$ are positive constants. This means user benefit increases with more contributors but gives diminishing returns. The performance of the system is the difference between the average benefit received by all users and the average contribution cost by all users, and is $Q = \alpha x^{\beta} - 1$.

The contribution level x is the fraction of users whose generosity (i.e. type) exceeds the current contribution cost, $1/x$. So the fraction of users that contribute can be given by:

$$x = Prob(t_i \geqslant 1/x)$$

We can now add a penalty mechanism to the network. It is difficult to exclude users who freeride from the network, as whilst freeriding behaviour may be observable, the innate types of a user may not be. The penalty is introduced to those users who are labelled as freeriders.

The effectiveness of any penalty scheme needs to be counterbalanced by the fact that those who freeride may be able to escape penalty by *whitewashing*, by repeatedly changing their identity to escape punishment. Anyone with a new identity will be indistinguishable from newcomers who want to contribute, and penalising everyone who is new can result in legitimate users not wanting to join in the first place.

The users in the system can be given 4 simplistic types to represent their status in the population:

- EC - existing contributors
- EF/WW - existing freeriders/whitewashers

	NC Not Penalized	NC Penalized
% Penalized	(1-d)(1-x)	d+(1-d)(1-x)
% Not Penalized	(1-d)x+d	(1-d)x

Figure 3: Penalization equations.

- NC - new contributors
- NF - new freeriders

The EF/WW group is deemed to be changeable depending upon the conditions of changing identities. If identities are unchangeable or highly penalised on change then they will act as freeriders - given the chance they will attempt to change their identity to escape penalty and will act like whitewashers.

To consider how this effects the users of the system we need to look at the contribution costs and threat. For example any users who intend to leave the system at the end of the current period do not care about the penalty that they will receive for freeriding. The system can be denoted by two different values, the contribution level of those who are leaving and also of those who are staying. R denotes contribution cost and T denotes threat, and we get two contribution equations:

$$x_{leaver} = Prob(t_i \geqslant R)$$

$$x_{stayer} = Prob(t_i \geqslant R - T)$$

With turnover rate d the average contribution in the system is none

$$x_{average} = dx_{leaver} + (1-d)x_{stayer}$$

This contribution level of users who stay is always greater than or equal to those who leave. User contribution is determined by the ratio between those who get full benefit and those who get reduced benefit. We can either penalize all newcomers, or we can penalize just those who have shown freeriding/whitewashing behaviour.

If we have two contrasting systems, one where everyones identity is permanent, and one where identities are changed for free, we can show that the cost with free identities are less, but the benefits to all users are also reduced. Feldman et al. [3] claim that a scenario can be reached where the penalty level can threaten users but yet no penalty is actually imposed.

Back to the actions of users of the system, each player should be able to learn from the performance of different strategies, though their own performance and observation of others. Whether they decide to use this information is up to them and if they do they will become a *traitor*. This is where a user builds up reputation by cooperating and then defects to a new non-cooperative strategy that results in better utility for them.

5. CONCLUSIONS

If we want to create a peer-to-peer network with ideas taken from Evolutionary Game Theory then I believe there are useful ideas that can be applied. However, we must also consider many things that are much more specific to the operation of peer-to-peer networks than the dynamics of populations, and focus upon individual entities actions in the population rather than the overview of population strategies.

The populations do however follow those found in more biological circumstances. Users can mutate by changing their strategy, either randomly or through learning and turnover can be modelled through the life and death of cells. If we can create a system which has an Evolutionary Stable Strategy where users cooperate (types NC/EC), and can't be invaded by freeriders and whitewashers (types NF/EF/WW) then we have solved the problem of peer-to-peer systems.

Figure 2 shows a game which players in a peer-to-peer system play, and I believe this is the best way to model such interactions in this system. It makes sense to have a client and a server, with no other players involved in an individual transaction. Even if multiple players are uploading different parts of a file to a client this can be modelled by pairwise client/server games.

Overall, more analysis need to be done before I could successfully construct a peer-to-peer system, and it seems certain areas are more applicable than others. Finding an ESS which supports cooperation is the key.

6. FURTHER WORK

I have provided an overview of how Evolutionary Game Theory could be applied to peer-to-peer systems. In order to decide whether a particular set of behaviours could produce an ESS requires data over a number of generations of the fitness and persistence of that behaviour.

Obviously whether the behaviours will come to an ESS or not will be affected by the mechanism of the peer-to-peer network. Incentive schemes need to be settled upon, and penalty schemes. A mechanism design approach could complement the game theoretic ideas.

A deeper analysis of the Evolutionary Game Theory materials may yield a game based idea that is even more suitable than Prisoners Dilemma.

Whatever theory we work out would eventually need to be tested in practice to prove that the theory is sound and that we haven't missed any factors that might effect the system, such as effects of hardware.

7. REFERENCES

[1] J. M. Alexander. Evolutionary game theory. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Summer 2003.

[2] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks, 2004.

[3] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems, 2004.

[4] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. *Lecture Notes in Computer Science*, 2232:75+, 2001.

[5] J. Maynard-Smith. Evolution and the theory of games. *American Scientist, 64(1), pp. 41-45.*, January 1976.

[6] M. A. Nowak and K. Sigmund. Evolution of indirect reproicity by image scoring. *Nature 393 (1998), 573-577.*

[7] K. Prestwich. Game theory. *http://www.holycross.edu/departments/biology/kprestwi/behavior/ESS/pdf/games.pdf.*

[8] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster. to share or not to share an analysis of incentives to contribute in collaborative file sharing environments. *Workshop on Economics of Peer-to-Peer Systems*, June 2003.

[9] W. Wang and B. Li. To play or to control: a game-based control-theoretic approach to peer-to-peer incentive engineering. *International Workshop on Quality of Service*, June 2003.