Purdue University Purdue e-Pubs

Computer Science Technical Reports

Department of Computer Science

2004

Towards the Design of Robust Peer-to-Peer Communities

Ioannis Ioannidis

Ananth Y. Grama

Purdue University, ayg@cs.purdue.edu

Report Number: 04-016

Ioannidis, Ioannis and Grama, Ananth Y., "Towards the Design of Robust Peer-to- Peer Communities" (2004). *Computer Science Technical Reports*. Paper 1599. http://docs.lib.purdue.edu/cstech/1599

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

TOWARDS THE DESIGN OF ROBUST PEER-TO-PEER COMMUNITIES

Ioannis Ioannidis Ananth Grama

Department of Computer Sciences Purdue University West Lafayette, IN 47907

> CSD TR #04-016 May 2004

Towards the Design of Robust Peer-to-Peer Communities

Ioannis Ioannidis and Ananth Grama Purdue University, W. Lafayette, IN, 47907, {ioannis,ayg}@cs.purdue.edu

May 3, 2004

Abstract

Peer-to-peer (P2P) sharing of resources, and technologies for facilitating resource sharing have witnessed tremendous advances in the recent past. As these technologies become commonplace, emphasis must be placed on the survivability of such communities in the face of non-cooperative peers (freeriders, malicious users). While incentive-based approaches provide possible solutions, similar problems in ecological populations are solved by complex social interactions that have evolved over the ages. Evolutionary biology has addressed these problems and numerous models of cooperation between selfish organisms have been proposed to explain how factors such as altruism, guilt, and the sense of justice have evolved in spite of harsh life-or-death conditions. These studies provide blueprints for essential computational techniques in support of stable, scalable, robust, and highly cooperative P2P communities. In this paper, we present a range of stable models of social interaction, their relevance to P2P communities, the associated computational bottlenecks in the context of P2P networks, and motivate the need for the next generation of structured and unstructured resource sharing networks.

1 Introduction

Conventional P2P and pseudo P2P networks such as Napster, Gnutella, and Kazaa have enabled millions of users to reach across the world and share their music collections with complete strangers in a chaotic and somewhat inefficient manner. Even as Napster was essentially shut down, these communities organized themselves in more chaotic, less efficient, and exceedingly popular ways. The initial explanation for this phenomenon seemed simple: there is collective benefit in sharing.

The tendency to cooperate towards achieving a common goal appears natural. Unfortunately, so is the tendency to stop cooperating and to take advantage of the benevolence of one's peers. Small communities can withstand attacks by selfish members by identifying and isolating them. The fear of being expelled from such communities is indeed a major deterrent to non-cooperation. Conventional P2P communities, however, are generally too large and the underlying information flow mechanisms too weak, for such deterrents to be effective. The challenge of designing a robust P2P network therefore involves incorporation of effective methods for identification and elimination of non-cooperating peers. In this paper, we examine the computational basis for building a stable and scalable P2P community that is large, maintains a high degree of anonymity, and is robust to groups of non-cooperating peers.

Issues of stability and scalability have been of principal interest to research on P2P networks. Scalability was among the first problems addressed in the context of structured P2P networks, since fast routing, and efficient placement and location of re-

sources are critical performance parameters. As resource sharing becomes widespread, uncooperative users become a threat to the integrity of such networks. Even if the cost of sharing a resource is small, it can be shown (using various models of cooperation and behavior) that a small number of non-cooperative peers can prove fatal. This is because cooperative peers that encounter such users tend to become non-cooperative, themselves, or leave the network. This reinforces negative behavior and eventually leads to the collapse of the network.

To prevent this scenario a number of incentive-based mechanisms have been proposed. These mechanisms involve various forms of pricing ([1], [2]), based on the intuition that if one accrues credit only by sharing a resource and must expend credit to acquire a resource, freeriders will be unable to attack the network. The primary challenges associated with architecting such networks include scalable protocols for handling currency and robustness to collusion. These challenges continue to be areas of active research interest. One potential drawback of incentive-based schemes is that such schemes may be less appealing to new users, who, contrary to their intentions, might be perceived as non-cooperative until they can build enough credit.

This position paper addresses an alternate approach to building stable networks. Experience with current resource sharing networks indicates that incentives other than acquiring a desired resource are not essential to the survivability of such communities. Large networks such as Kazaa (Kazaa provides a pricing utility, but it is optional, and most peers choose not to use it), Gnutella, iMesh, and WinMX appear to do well, effectively harnessing selfish behavior to build a cooperative base. This mechanism is not unlike the development of altruism between organisms, which has been studied extensively in evolutionary biology. In this context, Hamilton ([3]) provides an excellent example of cooperative behavior based on entirely selfish motives. Here, the appearance of herds is explained on the basis of a population of animals trying to minimize their individual probability that they will be attacked by a predator. As predators tend to attack prey that is closest to them, animals that are preyed upon try to cluster so that the area from which an attack can come (their cell in the Voronoi diagram of the members of the population) is minimized. This comes at a cost though, as their grazing area is minimized as well. Since the cost of being killed is much higher than that of losing some food, the formation of herds is favored.

In evolutionary biology such communal behavior is explained by units of selection attempting to maximize their benefit, even when said behavior can only be described as suicidal at first look ([4]). Game theory has been influential in this line of research and has provided the necessary analytical tools. The paradigm of the Repeated Prisoner's Dilemma is used to study the success and stability of behaviors in the presence of alternate strategies and with a scoring model for all possible outcomes of an interaction. We describe this paradigm in more detail in Section 3, but note that we assume that there is a cost c > 0 associated with allowing another peer to access a resource and a benefit b > 0 from successfully accessing a resource.

A score is associated with each peer based on the outcome of the interactions he has been involved in. Consequently, in the context of a single interaction, it is preferable for a peer to avoid sharing so as to maximize his score. When a peer enters the network his score is set to 0, even if this is not the first time he has entered the network. Upon termination of a session, the probability that he will use the network again, and the extent of use are proportional to the score achieved during the last session. In other words, a peer that accesses a large number of resources and who was not burdened by other peers accessing resources from it, is more likely to return than a peer who did not access any resources but served a large number of other peers. This is analogous to the evolutionary assumption that the expected number of offspring an organism leaves is proportional to a score, or fitness, it accumulates over its lifetime. A poor strategy results in a few, if any, descendants, similarly to a dissatisfied peer, who will stop using the network or will change his strategy the next time he joins the network.

The idea of applying game theoretic models to P2P networks is, by no means, new. However, it has generally been assumed that unless direct incentives are employed, a network is vulnerable to noncooperating peers (freeriders) and can only be supported by benevolent users who do not mind sharing their resources ([5], [6], [7]). An interesting and natural concept that can substitute incentives is that of memory. Users can remember the behavior of those they have interacted with and can communicate their experience across the network. This is a method commonly used in other communities (for eg., eBay), where the record of a user is published for all to sec. This model has worked well in deterring malicious users. A similar model has also been proposed in the context of P2P networks ([8]). However, maintaining a satisfactory, large history in a decentralized manner is difficult in large networks. Furthermore, it is difficult to sanction non-cooperative peers when they can change their identities at no cost and effectively erase their record. In Section 5, we discuss how shared histories are vulnerable to collusion, misconception, and deliberate misinformation, even if problems associated with changed identities are solved. Research in this direction, though, has derived some strong results that can be used to architect P2P networks, as long as one remains aware of the limitations of history-based strategies.

An alternate, scalable and powerful set of strategies is based on observation. In Section 4, we discuss a strategy called Observer Tit-For-Tat (OTFT). According to this strategy, users derive information about the behavior of their co-players by observing interactions they do not personally participate in. When approached by a player they know has not cooperated in the past, they ignore him. Pollock and Dugatkin ([9]) have analyzed the performance of OTFT and their results, when applied to P2P networks, show that OTFT can defeat freeriders under a minimal set of constraints.

We, finally, present a third class of strategies – contrite players (Section 5), that can correct misunderstandings over the intentions of players or errors in protocol implementation and execution. The main intuition behind these strategies is remorse. A player that has mistakenly decided not to cooperate, for example, due to communication failure, can repair his status by unconditionally cooperating in the next round, demonstrating his contrition. Contrite strategies alleviate the problem of vendettas when histories are used as guides and can be as good as the latter.

The rest of this paper is organized as follows: in Section 2 we briefly summarize related work; we formally define notions of stability, equilibria, and optimality in Section 3; we present various strategies for achieving these desired criteria in Sections 4-6, along with associated overheads in P2P networks; and finally motivate the development of technologies in support of robust communities in Section 7.

2 Related Work

As has been mentioned, there has been some work on applying game theory to the study of P2P networks. The tendency is to explicitly motivate cooperation by incentive mechanisms. Golle et al. ([1]) propose micropayments as the incentive. While this is an intuitive approach, enforcing micropayments in a secure manner is a challenging problem in itself ([10]). Lai et al. ([8]) have studied, experimentally, the impact of shared histories. Their work is closer in spirit to ours, since they show that cooperative environments can form without incentives. As we have briefly discussed, and will expand on later, shared histories suffer from scalability and robustness (fidelity) problems. However, they do provide several desirable characteristics.

We cannot, and do not intend, to present a comprehensive list of references for related work in evolutionary biology. An excellent summary of the work in this field can be found in [11]. We would like to note that one of the earliest models of selfish behavior leading to cooperation was described by Trivers

in 1971 ([12]). Dawkin's work ([13]) popularized the notion that ruthless competition is the driving force behind every hereditary trait, including those that are characterized as altruistic.

Finally, an interesting experiment was conducted by Axelrod in 1984 ([14]). A tournament of computer programs was conducted. The game was one of Prisoner's Dilemma, and the goal of the experiment was to see what strategies would succeed. The result was that a strategy called Tit-For-Tat (TFT) dominated. We will describe TFT in Section 4, but we note that the theoretical analysis for TFT predicts its dominance in the setting of the experiment ([15]). Although the setting was very different from that of P2P networks, as the population was small and the identities were persistent, the fact that competition through automated interfaces conforms to the theoretical results of evolutionary biology is encouraging.

3 Stability, Equilibria and Optimality in P2P Networks

In this section, we discuss some of the game theoretic concepts commonly used in evolutionary biology. A detailed introduction of these concepts is beyond the scope of this short paper and formal definitions can be found in many of the papers whose results we discuss. The interaction between two users, or players, is modeled using the Prisoner's Dilemma game. When a round of the game is played, each player decides whether to cooperate or defect (as the name suggests, the players are prisoners and declining to cooperate with your co-player implies that one defects to the captors). If both players cooperate, they receive a score of R. If both defect, they receive a score of P. If one of them cooperates, but the other defects, the defector receives a score of T and the cooperator, a score of S. For the game to make sense, we must have T > R > P > S. Multiple rounds of the game can be played and a cumulative score for each player is maintained. In this case, we must have $2 \cdot R > T + S$. The reason is that if T+S is larger than $2 \cdot R$, the players can maximize their scores by alternately cooperating and defecting and the game becomes trivial.

Nowak and Sigmund ([16]) showed that when the players make their decisions alternately, so that one of the players knows his co-players decision in the current round before deciding, the additional condition of T - P = R - S must hold. This is obviously the case in P2P networks due to the asymmetry of one of the users initiating an interaction by requesting a resource. Moreover, if two peers decide they do not want to allow each other to download, they have only a negligible cost and no benefit and we can set P=0. In this case the benefit of deceiving another user into allowing you to access a resource, without reciprocating, is R-S. Since the benefit S of being deceived should be negative and the benefit R of mutual cooperation should be positive, it holds that T > R > P > S in our adaptation of the Prisoner's Dilemma to P2P networks. We refer to R as benefit and S as cost, although the latter is actually the inverse of cost. We do this to be consistent with the conventional analytical framework, which is indifferent to the sign of S.

A second useful concept is one of equilibrium. Roughly, an equilibrium is reached when both players do equally well. The concept makes sense only if multiple rounds of the game are played. Equilibria are not extremely relevant to our problem for two reasons. The first is that an equilibrium can be reached when both players adopt defection as their strategy, or, equivalently when the network is made entirely of non-cooperative peers. The second reason is that an equilibrium can be disturbed by an invasion of other strategies. For example, a population consisting of defectors (peers that do not cooperate with anyone) is in an equilibrium. However, if a number of new members decide to cooperate with other cooperators and form a cluster, they can defeat the defectors.

To address these problems, the concept of evolutionarily stable strategies (ESS) is introduced. According to Maynard Smith's definition ([17]), a strategy Σ is ESS against a strategy Σ' , if the benefit $A(\Sigma|\Sigma)$ of adopting Σ in a pure Σ population is

strictly larger than the benefit $A(\Sigma'|\Sigma)$ of adopting Σ' , or, if $A(\Sigma|\Sigma) = A(\Sigma'|\Sigma)$ and $A(\Sigma|\Sigma) > A(\Sigma'|\Sigma')$. The second condition ensures that an ESS cannot be defeated by non-superior strategies that invade in clusters. For a strategy to be useful in P2P networks, it must at least be ESS against ALL D (defect-against-all strategy).

The concept of ESS formalizes the requirement for robust strategies. However, it does not address the need for a highly cooperative environment. Pareto-optimal strategies are those that, when adopted by the entire population, maximize the collective benefit. We note that an ESS that always cooperates with those that cooperate is the ideal strategy for P2P networks, as it is robust, and, when it prevails, all interactions are successful.

There are other properties of a strategy that can make it desirable. If the initial population consists largely of freeriders, it is difficult for cooperators to contact each other and, consequently, they might leave the network. An ESS is not guaranteed to succeed in invading a population. There is usually a threshold on the size of the invading population for an ESS to prevail. However, this threshold is hard to predict. A cooperative strategy that needs a small initial population has a better chance of prevailing.

Low implementation cost is a highly desirable feature of any strategy. Evolutionary models do not incorporate the cost of adopting a strategy. If memory is part of a strategy, we would like its capacity to be minimal. A strategy that floods the network to collect information is undesirable, for obvious reasons. The results of Section 4 are spectacular for this reason, because they imply that stability can be reached, in spite of the negligible implementation cost.

It is also desirable that a strategy can succeed under short generations. The usual assumption is that a further round is played between two specific players with probability w, so that the expected number of interactions between this pair is $\frac{1}{1-w}$. If the population is made of n+1 players, the expected number of interactions during the lifetime of a player is $\frac{n}{1-w}$. However, this need not be the case, as we will see in

Section 6. Image Scoring has been analyzed for the minimum number of interactions during the lifetime of a player that guarantees that this strategy will succeed. In the context of P2P networks, a shorter lifetime implies that a new user who is hesitant, will be easier to lure.

4 The Importance of Reputation

Tit-For-Tat (TFT) is a simple, cooperative strategy in which players cooperate when interacting with strangers. However, once a player defects on them, they never cooperate with him again. Let w be the probability that two players interact in the future. TFT can be ESS against ALL D (defectors), however, it needs high values of w to be so. In large P2P networks, the probability that one encounters the same player again is small, especially if non-cooperative peers periodically change their identities. When a stranger is encountered, a TFT player is willing to pay a price to learn how the stranger behaves. If there is a good chance of interacting with the new player in the future and the price paid for that bit of information is reasonable, a TFT player will be successful. When w tends to 0, no matter how low the price, the information collected regarding a player is going to be worthless. Defending against ALL D players in this scenario is impossible.

Pollock and Dugatkin ([9]) analyze a strategy called Observer Tit-For-Tat (OTFT). This strategy is designed to defend against ALL D by collecting information actively and without paying the price of an uncooperative transaction. An OTFT player behaves like TFT towards strangers. Instead of maintaining a private history of past interactions, however, he observes an interaction between two other players. If the interaction is successful, he classifies the two participants as cooperators. If not, he classifies them as defectors, irrespective of what the reason for the failure could be. If a player that has been observed and classified in the past tries to interact with the OTFT

player, he cooperates if the request comes from a cooperative player and defects if the request comes from a defector. A player classified as a defector may well be another OTFT player that defected on someone he perceived as a defector, but this does not matter. A number of scenarios are analyzed, where the population is in an equilibrium, one of TFT, OTFT, and ALL D dominate and the other two strategies attempt to invade. We summarize these results and extend them to the special case of P2P networks.

 Proposition 1: In a population saturated with simple game evolutionarily stable TFT, TFT is also evolutionarily stable against the mutation class {OTFT, ALL D} for sufficiently high w < 1; further, some Prisoner's Dilemma games exist which preclude OTFT invasion for any w > 0.

The necessary condition for OTFT to invade is

$$\max\{\frac{T-R}{T-P}, \frac{T-R}{R-S}\} < w < 1 - \frac{2R - (T+S)}{P-S} < 1.$$

For a P2P network we can set T = R - S and P = 0 and derive the simpler condition

$$\frac{S}{S-R} < w < \frac{R+S}{S}$$

which is easy to verify that it cannot hold for S < 0 < R. This proposition states that if somehow all the users of a P2P network played TFT, a user that switches to OTFT would eventually give up or change back to TFT. This, however, does not imply that OTFT is useless. One of the problems with TFT is that it needs to keep a history of past interactions, which leads to scalability problems. As we shall see, OTFT needs to keep only a constant size memory and can eliminate non-cooperative peers under certain constraints.

A generalization of ALL D is a strategy called ROVER ([18]). If the population is fragmented into groups of stationary cooperative players, a ROVER player moves from group to group, behaving as an ALL D while he is in a specific group and leaving when everybody in the group knows he is uncooperative. In the rest of the analysis, ALL D is replaced by its more powerful ROVER counterpart. We denote by q the probability that a player classified by an OTFT player, as defector, is indeed a ROVER. This probability is at least 1/2 and approaches 1 as ROVERs become more frequent.

 Proposition 2: In a population polymorphic for ROVER and TFT (ROVERs and TFTs coexist), TFT resists invasion by OTFT for sufficiently high w < 1. However, the threshold for w is lower and there are more games that do not preclude invasion.

The necessary condition for OTFT to invade is

$$w \le 1 - \frac{(1-q)[2R - (T+S)]}{q(P-S)},$$

which simplifies, for P2P networks, to

$$w \le 1 + \frac{(1-q)R}{qS}.$$

This is certainly feasible and the threshold approaches 1 with q, signifying that in a population where non-cooperative peers are a problem, OTFT can do better than TFT.

- OTFT does not require infinite memory, unlike TFT. Remembering only the last observed interaction is enough, therefore, OTFT is more scalable. A larger memory helps, however it does not change the two propositions. A positive effect of larger memory is that a player seen defecting in two interactions, has a much larger chance of truly being a defector. This implies that q grows with memory size.
- Proposition 3: OTFT is ESS when TFT is not; moreover, in principle w can be 0 and OTFT is still ESS.

A new variable z is introduced for this proof. z is the probability that someone engaged in mutual defection is seen and remembered by a future, yet presently unencountered, player. Freeriders have z=0, but all OTFTs have the same z value. Assume that OTFT is coded into the P2P clients, and that a few users have been able to alter the code to behave as freeriders. The condition for OTFT to be ESS is

$$w > -\frac{S}{R-S} - \frac{z}{1-z} \cdot \frac{R}{R-S}.$$

Since w = 0, the condition for OTFT to be ESS in the P2P setting reduces to

$$z > \frac{S}{S - R}$$
.

We make a further assumption on the Prisoner's Dilemma for P2P networks. We assume that the benefit of accessing a resource is larger than the cost of allowing another user to access a resource. This implies that R > -S and the stability condition is always satisfied if z > 1/2.

 Proposition 4: For sufficiently high z < 1, OTFT requires less initial clustering than TFT to invade a population saturated with ALL D.

This is an important result, since it can be shown that a scattered and small population of OTFTs can invade any network overwhelmed by freeriders and establish a cooperative environment. We introduce a new variable, c, to show this. c is the probability that an interaction carried out by an OTFT will be with a fellow OTFT, or the clustering of OTFTs. The condition for OTFT to successfully invade in a P2P setting is

$$c > -\frac{(1-z)S}{R-S}.$$

We want the right side of the inequality to be as small as possible. If it is 0, only two OTFT players, in principle, can eliminate a population of freeriders, no matter how large the network is, as there is always a small, but positive, probability that they will meet. The right side becomes 0 only when z=1, however, c can be reduced if the cost of allowing a resource access is small.

The variable z gives rise to a set of interesting possibilities. While w and c are beyond the control of the designers of a P2P network and will approach 0 as the network grows, there are ways to control z. Suppose that an OTFT user A is asked by a stranger to allow a download. A may spawn a "dummy", unconditionally cooperative user B, who attempts to download from the stranger. If the stranger seems uncooperative to B, A will be non-cooperative to the stranger. If the stranger is an OTFT himself, he will spawn his own "dummy" C, which will cooperate with B. With both parties assured of the other's good intentions, the original transfer can go through. This makes z = 1 for OTFT players. For structured P2P networks such as PASTRY and CHORD, an alternative may be for a user A to store an aggregate of all his interactions at node hash(A) (the rendezvous point of A). This aggregate may be the result, success or failure, of A's last interaction, or the percentage of successful interactions to avoid manipulation of the results. An OTFT user B contacted by A can query hash(A) for the aggregate, and, based on this observation determine A's status. This is indeed not dissimilar from how eBay ratings work, only, here we need to distribute these ratings over the network. Again, this makes z=1.

These strategies are merely suggestions based on a formal analysis of an easily implementable and very natural strategy. OTFT seems to correspond to Alexander's view of the prerequisite for indirect reciprocity as "direct reciprocity occurring in the presence of interested audiences" ([19]). Admittedly, non-cooperative peers may try to collude and deceive cooperative players by, for example faking, successful interactions. Cooperative players may respond by finding ways to make their observations more accurate. This "arms race" is a common phenomenon in nature, especially among predator and prey. The positive aspect of this is that freeriders must expend

more resources. If the cost of allowing a resource access is small, freeriders will benefit from converting to cooperative strategies.

5 Repentance and Error Correction

Boerlijst et al. ([20]) examine the problem of cooperation in a model that allows mistakes to happen and be corrected. Suppose that a user has assigned all the burden of enforcing a cooperative strategy to the P2P client. There are two things regarding the behavior of the client that could go wrong. Either the software has an error and occasionally refuses to cooperate when it is supposed to, or, due to a network error it misunderstands the intentions of a cooperative player and refuses to cooperate. These errors are not detected by the user, who might have been able to resolve the misunderstanding. If the cooperative players use a TFT-like strategy in which a player is blacklisted if he refuses to cooperate, a simple error of this nature may lead to an incorrect vendetta. If the network supports shared historics, as in [8], the result for users whose transactions incurred such errors can be dramatic, with large parts of the network being informed of their defection and refusing to ever cooperate with them. To alleviate such problems, Boerlijst et al. describe and analyze a set of errorcorrecting strategies.

We start with the simplest of these strategies, Contrite Tit-For-Tat (cTFT). cTFT assumes there is a standing associated with each player. A player starts with good standing, but loses it if he defects on a player with good standing. cTFT players behave just like TFT players, until an error occurs. If a cTFT player accidentally defects on a player of good standing, he loses his good standing. In the next round, the cTFT player with bad standing cooperates unconditionally with a co-player, however, the co-player defects without losing his standing. After this round of repentance, the cTFT will recover to a good standing. As shown by Sugden ([21]), cTFT is as good as

TFT for invading a population of defectors and it is ESS in a broader sense than that discussed in Section 3: if there is a small, but positive, probability that a move is misimplemented, no other strategy does better than cTFT.

The authors experiment with the behavior of mixtures of strategies. They conducted a set of experiments where they set S = 0, P = 1, R = 3 and T = 5.5. The effect of these payoff values is that there is a high temptation to defect. The strategies used were ALL D, cTFT, REMORSE (another error correcting strategy that cooperates if and only if it was in bad standing in the previous round or both players cooperated), and GRIM (a strategy that starts as cooperative, but turns into ALL D the first time it encounters a defection). All simulations start with an ALL D population. A small probability of mutating to another strategy allows the appearance of cooperative strategies. Each time, cTFT and RE-MORSE eliminated the uncooperative strategies and established a highly cooperative environment.

Finally, a strategy called Prudent PAVLOV (pPAVLOV) is discussed. Two pPAVLOV players will cooperate indefinitely, if there are no errors. If a pPAVLOV player mistakenly defects, both players enter a D_1 state and they both defect. In the next round, a D_0 state is entered and they both defect again. They resume cooperation in the next round. pPAVLOV is ESS for P2P networks if

$$wR + w^2R > -S.$$

Assuming that the w^2R term is negligible, we have that pPAVLOV is ESS if $w > -\frac{S}{R}$. In the next section we will see a very similar stability condition emerge in the case of incomplete histories. pPAVLOV has two huge advantages over other strategies. The first is that it can correct errors in both the implementation of the strategy and in the perception of a move. The second is that it is scalable.

Strategies based on standings are generally less scalable because they need to tag every co-player. pPAVLOV, on the other hand, needs to keep track of where it stands with respect to a co-player,

only if there has been some degree of interaction. Strangers are assumed to be cooperative. This leads to pPAVLOV being suckered every third round by ALL Ds, but as long as it is ESS it does not matter, ALL Ds are eventually eliminated. In terms of P2P networks, a user needs to keep track of other users only until he leaves the network. The next time he enters the network, he can start with a clean slate. This is a scalable scenario, since the history information can be cleared when it exceeds a manageable size. Every time the history is cleared, it is akin to a player dying. Unfortunately, this affects the value w, as the probability that a player will be seen again before his record is cleared is extremely low. Since we need w to be larger than -S/R for stability, pPAVLOV may not be stable in a P2P setting. As network connections become faster and cheaper, the cost of allowing a resource access becomes small and pPAVLOV becomes stable. Furthermore, it could be that after a certain session length, users interact mostly with peers they have seen recently. If this is the case, users that stay in the network for a long time can discard information about peers they have not heard from in a while, without hurting their pPAVLOV strategy. Such cooperative and persistent users can be influential in stabilizing a cooperative network.

6 Building One's Image

Kazaa implements an interesting policy: each user has a standing score. This score improves with each allowed download. Periodically a small quantity is subtracted, so that users keep sharing. This is different from an incentive-based mechanism, as users are not obligated to consult their peers' standing before deciding to cooperate with them. However, many of them will be reluctant to cooperate with users in poor standing. This standing score cannot be turned off, unlike Kazaa's pricing policy. There are two disadvantages to this scoring. The first is that it needs a trusted authority supervising each interaction so that colluding users cannot artificially elevate their scores. The second is that proficient users may be able to

alter the code and become undetectable freeriders. Despite these shortcomings, scores are a useful tool for cooperative users. Nowak and Sigmund analyze a similar policy in the context of indirect reciprocity and how it may have evolved.

Nowak and Sigmund first conducted experiments on a multivalued image score metric ([22]). In these experiments, players started with a score of 0. Random interactions between a recipient and a donor changed the individual scores. If a donor refused to cooperate, a unit was deducted from his score. If he cooperated, a unit was added. The maximum score was 6 and the minimum was -4. A negative score indicated an uncooperative player. Each player decided to cooperate or not, as a donor, based on a value k, fixed during a player's lifetime and inherited by his offspring. The strategy was to cooperate with recipients that had a score no less than k. A negative k indicated a cooperative strategy. A positive k, a non-cooperative one. k = 0 indicated a discriminatingly cooperative strategy, in which only those users appearing to be freeriders are denied cooperation. A lifetime lasted, on average, 2.5 interactions. After a fixed number of interactions occurred between players of a generation, all players died. The next generation consisted of the previous generation's offspring. The number of children a player left was proportional to his fitness. Each time a player was the recipient in a successful interaction, his fitness increased by b. Each time a player was the donor in a successful interaction, his score decreased by c, 0 < c < b. The initial population started with a uniform mixture of strategies. After just 20 generations, only cooperative strategies were significantly represented, with discriminating cooperators being the majority and unconditional cooperators (k = -4) making up the rest of the population. After 150 generations, the few existing freeriders had eliminated all the unconditional cooperators. The freeriders, in turn, had been eliminated by discriminating cooperators, who became the only existing strategy.

The same experiment is repeated, however, random mutations are introduced. The result was a cyclical effect. Discriminating cooperators would win, but due to mutations, less strict cooperators would appear, leading to a short proliferation of freeriders. Discriminating cooperators would reestablish a cooperative environment. A less encouraging experiment followed. In this setting, only the recipient, the donor and ten random observers updated the scores after an interaction. One player could have different scores in the eyes of different players. After 100 generations, only 15% of the population were cooperators. This percentage fell as the population increased.

The problem translates immediately to P2P networks. Maintaining perfect score information eliminates freeriders, however, this is expensive and might not be practical. Maintaining imperfect information can be efficient, but is ultimately futile. A remedy to this problem is as follows: a player has two evolutionary variables k and h. He cooperates only with players whose score is at least k and only if his own score is less than h. This strategy, called AND, is successful in eliminating freeriders (players with k > 0), even with incomplete information. However, there is an associated price. The prevalent strategies are reluctant to cooperate. With complete information, the prevalent strategies were those with $k \leq 0$ and h = k + 1. The higher the mean h value, the more cooperative the environment, but with complete information, only 55% of the interactions are successful. With incomplete information, there is some uncertainty on how accurate the information a player possesses is and the prevalent strategies exhibited a somewhat higher differential between h and k. Yet, only 57% of the interactions was successful.

A compromise between eliminating freeriders and allowing a more cooperative environment can be the OR strategy. According to this strategy, a donor cooperates if the recipients score is at least k or the donor's own score is less than k. Nowak and Sigmund experimented with this strategy and with incomplete information, the trend was towards cooperative strategies ($k \le 0$), with k values uniformly distributed. In this case, 80% of all interactions were

successful.

The authors observe an intriguing phenomenon during their experiments. Suppose all players adopt a simple, k = 0 strategy, but start with a random score. The question whether this population converges to a cooperative one is non-trivial. Experimentally, it is observed that there is a threshold on the fraction of players that start with negative scores and still the population converges to all-out cooperation. This threshold is 0.7380294688360.... If more players start with negative scores, the population converges to a non-cooperative one. Therefore a population in which more than 73% of the players appear to be freeriders, although their policies are not that of a freerider, still converges to a cooperative network. This observation does not translate directly to P2P networks, because the cooperative policies help convergence. However, it is a good indication that a P2P network can survive a large number of freeriders, even without incentives.

6.1 Predictive Tools

AND and OR strategies may alleviate scalability problems by disseminating information, but they do not solve them. In a structured P2P network, the result of an interaction can be communicated to a constant number of random users. However, as the network grows, this number must also grow to enforce cooperation. Nowak and Sigmund formally analyze their Image Scoring strategy when there only two possible scores, good and bad ([11]). Their results may help resolve the scalability bottlenecks.

It is easy to see that when there is complete information about the standings, the strategy that cooperates only with those having a good standing is ESS. It is also possible to derive the stability condition when score information is incomplete. Suppose the probability that one knows the co-players score is q. The stability condition is

$$q > \frac{c}{b}$$

This is reminiscent of the stability condition for pPAVLOV, since c is equal to -S and b is R. This is

also exactly identical to Hamilton's rule for altruism through kin selection, which states that cooperation through kin selection works whenever the coefficient for relatedness is larger than the cost to benefit ratio ([23]). The problem in P2P networks is that if we just try to disseminate information about the standing of a player, the sheer size of the network will force q to be almost 0. Furthermore, if freeriders can change identities at no cost, this inequality will never be satisfied.

A solution to this problem is to give up on trying to disseminate standings information and help cooperative users predict the behavior of a stranger. In real life, when we meet someone for the first time, we try to read certain signs that clue us in on his character. We can apply the same principle to P2P networks. While in [11] q is assumed to be derived from the interaction history of a player, it could be an oracle that gives this information. Can we build such an oracle for P2P networks? When a stranger asks for a resource, certain characteristics can be identified. How many resources is he sharing? What is the nature of these resources? Are there other users trying to access resources from him? If not, can I access something at random from him? If the network is fast and the cost c is small, the prediction accuracy does not need to be high for the cooperative strategy to be stable. Non-cooperative peers may try to deceive this tool, however, this increases the cost of an uncooperative strategy and discourages such behavior. We believe that in existing, unstructured file sharing networks such as Gnutella and WinMX, there are users that implement such strategies and try to infer the nature of a stranger. Software designed to make their predictions more accurate and easier to arrive at could be a powerful aid against non-cooperative peers.

A second interesting aspect relates to the robustness of this strategy. Image Scoring is ESS, but once there are no non-cooperative peers, the cost of keeping track of the scores of other users is unnecessary. The population can drift to a mixture of discriminating cooperators and unconditional cooperators. It is possible to derive the dynamics of this population in the event of an invasion by freeriders. Suppose the fraction of population that is made of unconditional cooperators is p and the rest (1-p) is made of discriminating cooperators. Depending on the values of b, c, q, and the expected lifespan, there is a threshold for p above which the population cannot resist an invasion by non-cooperators. On the other hand, if p is below the threshold, the invasion will reduce the number of unconditional cooperators and noncooperators will proliferate. However, in the end they will be eliminated. The new population will consist of p' unconditional cooperators and 1-p' discriminating cooperators, with p' being smaller than p. In other words, after an unsuccessful invasion, the network is even more adverse to non-cooperators. As a consequence, occasional bursts of non-cooperators should be expected as the users become complacent, however, the network will quickly emerge to a higher state of cooperation.

7 Concluding Remarks

The results we have derived and summarized here represent only a fraction of the literature in population dynamics. Yet, a quick examination reveals that the problem of handling non-cooperative users can be solved without resorting to explicit incentives. Certainly, if the information maintained at each peer regarding the rest of the network is complete, eliminating non-cooperative peers with either of the three classes of strategies discussed in this paper is easy and even error correction is possible. Similarly, if there is a centralized authority that maintains scores and if these scores cannot be tampered with, standings are very effective. The challenge is to design P2P networks in which users maintain only a limited amount of information about the state of the network, each user can implement the policies independently, and stabilization to a cooperative state is guaranteed.

OTFT can be a useful strategy for large networks in which non-cooperating peers are common. If observations are effective, this is a scalable and resilient strategy. However, in a network where noncooperators are rare, the cost and delay of simulating a series of interactions can be an overkill. pPAVLOV is more desirable in these cases. As we have discussed, it has the property that it corrects errors in both implementation and perception of a move. It can be difficult to implement pPAVLOV in a scalable manner while keeping it ESS, however, stability when freeriders are rare is not crucial. A combination of OTFT and pPAVLOV holds considerable benefit. If non-cooperators appear only rarely, the most desirable strategy is pPAVLOV. If pPAVLOV is not ESS for the specific network, the frequency of noncooperative peers encountered will increase. If it rises above a threshold, the client may swith to the more expensive, but more resilient, OTFT strategy, until non-cooperators become rare again. This process can be made entirely transparent to the user. Implementing efficiently these two strategies is an open problem and experiments on real networks should reveal fascinating behavior. In theory though, OTFT, and pPAVLOV can stabilize any network.

A successful strategy that requires the user's involvement may depend on predicting the peer's behavior. It is hard to imagine that prediction can be fully automated. However, providing the user with the right tools can make it easy to employ. The advantage of prediction is that it can defend against non-cooperators when changing ones identity has no cost. Furthermore, it is stable even if the user decides whether to change his strategy or give up the network only after two interactions. A kind of balance on the usage of predictive tools should also evolve. Accurate tools are more expensive than inaccurate ones. Browsing the contents of the shared folder of a peer says a lot about his behavior, however, simply assessing the number of resources he brings to the network is faster and can be easily automated. Depending on how expensive it is to allow a download and how frequent non-cooperators are, users will lean towards one of these solutions.

An uncooperative strategy that is of interest in the context of P2P networks that are divided into groups

of users is that of the the briefly discussed ROVER ([18]). This division can be the result of a hierarchical organization aimed at efficiency or of groups forming on common interests. In any case, resilience against ROVER, not just ALL D, can be a desirable property for a cooperative strategy.

Finally, we would like to note that although in our analysis we have assumed that the probability a peer will meet the same peer twice and the clustering of cooperative peers are negligible quantities, this is not necessarily the case in real networks. It is reasonable to assume that if a user downloaded a file from a certain peer, he will return to the same peer. Estimating these values could reveal that strategies such as pPAVLOV, that depend on a non-negligible w to be ESS and scalable, can be scalable, as well as ESS.

References

- P. Golle, K. Leyton-Brown, I. Mironov, M. Lillibridge, "Incentives for Sharing in Peer-to-Peer Networks", ACM Conference on Electronic Commerce (EC), 2001.
- [2] B. G. Horne, B. Pinkas, T. Sander, "Escrow Services and Incentives in Peer-to-Peer Networks", ACM Conference on Electronic Commerce (EC), 2001.
- [3] W. D. Hamilton, "Geometry for the Selfish Herd." Journal of Theoretical Biology 31, 295-311, 1970.
- [4] A. Zahavi, "Altruism as a Handicap The Limitations of Kin Selection and Reciprocity", Journal of Avian Biology 26, 1-3.
- [5] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, D. S. Wallach, "Security for Structured Peer-to-Peer Overlay Networks", Fifth Symposium on Operating System Design and Implementation (OSDI), 2002.
- [6] J. Douceur, "The Sybil Attack", First International Workshop on Peer-to-Peer Systems (IPTPS), 2002.

- [7] E. J. Friedman, P. Resnick, "The Social Cost of Cheap Pseudonyms", Journal of Economics and Management Strategy (10) 2, 2001.
- [8] K. Lai, M. Feldman, I. Stoica, J. Chuang, "Incentives for Cooperation in Peer-to-Peer Networks", First Workshop on Economics of Peer-to-Peer Systems, 2003.
- [9] G. B. Pollock, L. A. Dugatkin, "Reciprocity and the Evolution of Reputation", Journal of Theoretical Biology, 159, 25-37, 1992.
- [10] S. Micali and R. Rivest, "Micropayments revisited", In B. Preneel, editor, Progress in Cryptology CT-RSA 2002, volume 2271, Lecture Notes in Computer Science. Springer-Verlag, February 18-22 2002.
- [11] M. A. Nowak, K. Sigmund, "The Dynamics of Indirect Reciprocity", Journal of Theoretical Biology 194, 561-574, 1998.
- [12] R. Trivers, "The Evolution of Reciprocal Altruism", Quarterly Review of Biology 46, 35-57, 1971.
- [13] R. Dawkins, "The Selfish Gene", Oxford: Oxford University Press, 1976.
- [14] R. Axelrod, "The Evolution od Cooperation", originally 1984, reprinted 1989 by Penguin, Harmondsworth.
- [15] R. Axelrod, "The Emergence of Cooperation Among Egoists", Am. Pol. Sci. Rev. 75, 306-318, 1981.
- [16] M. A. Nowak, K. Sigmund, "The Alternating Prisoner's Dilemma", Journal of Theoretical Biology 168, 219-226.
- [17] J. Maynard Smith, "Evolution and the Theory of Games", Cambridge: Cambridge University Press, 1982.

- [18] L. A. Dugatkin, D. S. Wilson, "ROVER: A Strategy for Exploiting Cooperators in a Patchy Environment", American Naturalist 138, 687-701, 1991.
- [19] R. D. Alexander, "The Biology of Moral Systems", Aldine de Gruyter, New York, 1987.
- [20] M. C. Boerlijst, M. A. Nowak, K. Sigmund, "The Logic of Contrition", Journal of Theoretical Biology 185, 281-293, 1997.
- [21] R. Sugden, "The Economics of Rights, Cooperation and Welfare", Blackwell, Oxford, 1986.
- [22] M. A. Nowak, K. Sigmund, "Evolution of Indirect Reciprocity by Image Scoring", Nature 393, 573-577, 1998.
- [23] W. D. Hamilton, "The Evolution of Altruistic Behavior", American Naturalist 97, 354-356 (1963).