

# Cooperation under Scarcity: The Sharer's Dilemma

Michael Rogers<sup>1</sup> and Saleem Bhatti<sup>2</sup>

<sup>1</sup> University College London, London WC1E 6BT, UK

m.rogers@cs.ucl.ac.uk

<sup>2</sup> University of St. Andrews, Fife KY16 9SS, UK

saleem@cs.st-andrews.ac.uk

**Abstract.** Many researchers have used game theory to study the problem of encouraging cooperation in peer-to-peer and mobile *ad hoc* networks, where resources are provided collectively by the users. Previous work has modelled the problem as either a multi-player social dilemma or a network of two-player prisoner's dilemmas, but neither of these approaches captures a crucial aspect of the problem, namely *scarcity*: when resources are limited, players must not only consider how to establish and sustain cooperation with each opponent, but how to allocate resources among their opponents in order to maximise the total cooperation received.

This paper presents a new game theoretic model of cooperation under scarcity, *the sharer's dilemma*, and a simple *expected utility strategy* that is shown to perform well against a wide range of opponents. The expected utility strategy can easily be applied to file sharing networks to create an incentive for users to contribute resources.

## 1 Introduction

Any system with infrastructure that is provided collectively by the users faces the problem of encouraging users to contribute resources as well as consuming them. In other words, users must cooperate with one another. Cooperation in peer-to-peer and mobile *ad hoc* networks has received a great deal of attention from researchers in recent years, and many incentive mechanisms to encourage cooperation have been proposed. However, in this paper we will argue that such mechanisms have so far lacked a suitable theoretical foundation: some use *ad hoc* methods that may be vulnerable to manipulation by selfish users, while others are based on game theoretic models that fail to capture the problem of allocating scarce resources among neighbours. We develop a new game theoretic model of this problem and show that a strategy based on strict utility-maximisation creates an incentive for selfish users to cooperate.

The next section of this paper reviews existing models of cooperation in networks, including the prisoner's dilemma and multi-player dilemmas. Section 3 introduces our new model, the sharer's dilemma, and describes the expected utility strategy. Section 4 presents the results of simulations comparing a number of strategies. We discuss the limitations of our findings in section 5, and conclude the paper in section 6 with some directions for future work.

## 2 Game Theoretic Models of Cooperation

If a communication network is viewed as a group of self-interested individuals interacting according to rules specified by the protocol designer, then *game theory* provides tools for modelling the behaviour of rational participants, and *mechanism design* can be used to create protocols that reward cooperation, encouraging rational participants to behave in ways that benefit the network [1,2,3].

It might seem reductive to regard the participants in a communication network as simple egoists – even economists no longer believe that people are motivated purely by self-interest [4,5,6] – but here we are concerned with the *behaviour of nodes* rather than the *intentions of users*: game theory is not appropriate for modelling all human interactions, but it is well suited to modelling those interactions in which humans delegate routine decisions to software, reducing complex social considerations to a choice between programmatic ‘strategies’. Even if most participants have the best interests of the network at heart, game theory allows us to assess the network’s vulnerability to exploitation by a selfish minority. Mechanisms that prevent free riding may also be able to prevent denial-of-service attacks in which malicious users exhaust the resources provided by others.

### 2.1 The Prisoner’s Dilemma

Simple games can embody surprisingly complex problems, and perhaps no simple game has received more attention than *the prisoner’s dilemma*, a single-round game for two players, each of whom chooses between two actions, *cooperation* and *defection*, and receives a payoff that depends on the choices of both players:  $T$  is the ‘temptation’ payoff for unilateral defection,  $R$  is the ‘reward’ payoff for mutual cooperation,  $P$  is the ‘punishment’ payoff for mutual defection, and  $S$  is the ‘sucker’ payoff for unilateral cooperation [7].

The dilemma arises because  $T > R > P > S$ , which means a rational player will defect regardless of her opponent’s choice. The players cannot escape the dilemma by communicating about their intentions, because a rational player will claim that she intends to cooperate, but will then defect. Thus rational players always defect, leading to a suboptimal payoff  $P < R$  for both players.

The prisoner’s dilemma has been used to model a wide range of situations in nature and society where the benefit of cooperation is greater than the cost. Wahl and Nowak [8] describe the prisoner’s dilemma in terms of the cost of cooperating,  $c$ , and the benefit of receiving cooperation,  $b$ . The restriction  $b > c > 0$  leads to the payoff structure described above. Roberts and Sherratt [9] describe the dilemma using a single parameter,  $k = b/c$ .

Public goods problems, social dilemmas [10] and reciprocal altruism [11] find natural expression in the form  $b > c > 0$ , but not all prisoner’s dilemmas can be expressed in this way: for example, many studies use the payoffs  $T = 5$ ,  $R = 3$ ,  $P = 1$ ,  $S = 0$ . In this paper we will only consider dilemmas that arise from the costs and benefits of cooperating and can therefore be expressed in the form  $b > c > 0$ .

### 2.2 The Shadow of the Future

Although rational players always defect in the single-round prisoner’s dilemma, it may be possible to establish cooperation if the game is repeated for more than one round.

Players who expect to interact for many rounds must consider the long-term effects of their short-term decisions: automatic defection is no longer necessarily the best strategy, because players have the chance to recognise cooperative opponents and gain a higher payoff through mutual cooperation (although each player will still be tempted to defect once her opponent has started to cooperate).

Simple strategies for the repeated prisoner's dilemma include Tit For Tat, which cooperates in the first round and thereafter copies its opponent's action from the previous round [12]; Win Stay Lose Shift, which cooperates in the first round and thereafter repeats its previous action if it receives cooperation, or switches to the other action if it suffers defection [13]; and Stochastic Tit For Tat, which cooperates with a probability equal to the fraction of rounds in which the opponent has cooperated [14]. All of these strategies create an incentive for rational opponents to cooperate, while minimising losses against uncooperative opponents.

### 2.3 Multi-player Dilemmas

When a dilemma involves more than two players, the model must specify which actions affect which players. If each player chooses one action that affects all her opponents, the situation is a *social dilemma* [10]; encouraging cooperation is harder than in two-player games, because it is not possible to cooperate with cooperators while defecting against defectors [15,16].

In a *networked social dilemma*, each player chooses one action that affects her neighbours in a spatial lattice or other network [17,18,19,20]. Cooperation can succeed if cooperative players' interactions with other cooperators sufficiently outnumber their interactions with defectors.

Finally, if each player can choose a different action for each opponent, the situation can be modelled as a *network of two-player games*. Several networked variants of the prisoner's dilemma have been developed [21,22,23,24], all based on the assumption that a player's choices and payoffs in her pairwise games are independent. However, if the payoffs represent the costs and benefits of cooperating, we may ask whether the assumption of pairwise independence is always realistic: there may be situations in which a player has limited resources for cooperation, but can allocate them freely among her pairwise interactions. In such cases no two-player strategy indicates how best to allocate her scarce resources.

We argue that many of the situations modelled as networked dilemmas fit this pattern: players can allocate their resources unevenly, cooperating more with some opponents than with others. Hunters sharing food, animals grooming one another, and network nodes uploading files are all faced with opportunities to strengthen or weaken cooperative relationships by choosing how much to share, and with whom.

## 3 The Sharer's Dilemma

To explore the problem of allocating resources in networked dilemmas we propose a simple extension to the prisoner's dilemma, incorporating scarcity into the game by

limiting the number of times a player can cooperate in each round. We call this new game *the sharer's dilemma*. The prisoner's dilemma can be viewed as a special case in which the limit is high enough that cooperation with every opponent is possible in every round.

As with the prisoner's dilemma, many variants of the game are possible, but here we will only consider the simplest case: in each round, a player can either cooperate with one of her opponents or defect against them all. While simple, this starting point captures the essential problem of cooperation under scarcity: when resources are limited, the problem is not only how to establish and sustain cooperation with each opponent, but how to prioritise opponents in order to maximise the total cooperation received.

Strategies from the prisoner's dilemma can be adapted to the sharer's dilemma by specifying how to choose between neighbours when more than one neighbour is eligible for cooperation.

### 3.1 The Expected Utility Strategy

If a player expects that cooperating with an opponent will result in a higher level of cooperation in return, she can weigh the expected benefit of her opponent's reciprocation against the cost of cooperating, and compare the incentives offered by different opponents. This idea is the basis of our *expected utility strategy* for the sharer's dilemma.

A player using the expected utility strategy estimates the benefit of cooperating with each opponent, under the assumption that all the benefit received from the opponent so far is a result of reciprocation – in other words the cooperation received can be attributed to the cooperation given. The benefit of all the cooperation received in previous rounds divided by the cost of all the cooperation given in previous rounds is the *expected benefit per unit of cost* of cooperating in the current round. A player maximises her expected benefit by cooperating with whichever opponent will provide the greatest benefit in return.

When comparing her opponents in this way, a player does not need to know the cost or benefit of cooperation from the opponent's point of view – she only needs to estimate the cost to herself of cooperating, and the benefit to herself of the resulting reciprocation, so costs and benefits may be subjective.

If it is possible to measure costs and benefits in the same units then the cost of earning reciprocation can be subtracted from the expected benefit, and a player may defect if the cost of cooperating exceeds the expected benefit. However, even if costs and benefits are not commensurable, a player can still use the expected utility strategy to maximise her benefit by earning the most reciprocation per unit of cost.

Like any cooperative strategy, the expected utility strategy faces the problem of *bootstrapping*: when two players first meet, one or both of them must risk cooperating without knowing how much reciprocation (if any) will result. In the prisoner's dilemma, Tit For Tat and Win Stay Lose Shift take the simple approach of always cooperating in the first round, but this may not be possible in the sharer's dilemma due to the limit on the amount of cooperation per round. The expected utility strategy could assign a high expected benefit to first-time interactions, but this might be vulnerable to exploitation by *whitewashers* who can continually change identities [25];

alternatively, the benefit could be estimated using the average benefit of previous first-time interactions [26].

Uncertainty about the duration of the game can be incorporated into the strategy by applying a *discount factor* to future payoffs, reducing the expected benefit of reciprocation if games tend to be short-lived [25]. The discount factor need not be the same for all opponents; if old players can be expected to outlive new players, as in many peer-to-peer networks [27,28,29], then it may be appropriate to use a heavier discount factor for new opponents.

## 4 Simulations

This section describes simulations to compare various strategies for the sharer's dilemma. Our model is a population of  $n$  players connected uniformly at random so that each player has  $d$  neighbours on average. In each round of the game, each player either cooperates with one of her neighbours, increasing the neighbour's payoff by  $b$ , or defects, increasing her own payoff by  $c$ . The players make their choices in a random order each round.<sup>1</sup>

We simulate two strategies adapted from the prisoner's dilemma: the first, Tit For Tat (TFT), cooperates with a randomly chosen neighbour that cooperated with it in the previous round, or defects if no neighbours cooperated. The second, Stochastic Tit For Tat (STFT), cooperates with one randomly chosen neighbour, choosing each neighbour with a probability proportional to the fraction of rounds in which the neighbour has cooperated. To bootstrap cooperation, TFT and STFT treat new neighbours as if they cooperated in the previous round.

We also simulate a strategy based on BitTorrent's incentive mechanism, which uses reciprocation to encourage peers in a file sharing network to upload [30]. Each peer maintains an *active set* of connections, with all other connections 'choked' (nothing is uploaded). The peer updates its active set periodically, unchoking those connections that have recently provided the best download speeds, so neighbours that upload more quickly are more likely to be unchoked. To bootstrap cooperation, the active set also includes one randomly chosen connection.

To model this behaviour in the sharer's dilemma, the BitTorrent (BT) strategy cooperates in each round with a randomly chosen member of its active set, which contains one randomly chosen neighbour and the  $s$  other neighbours that have provided the most cooperation in recent rounds. A new active set is chosen every  $r$  rounds, using an exponential moving average to measure the cooperation received from each neighbour.<sup>2</sup>

Finally, we simulate the expected utility (EXU) strategy described in the previous section, which cooperates with whichever neighbour has provided the highest benefit/cost ratio in previous rounds, unless the highest ratio is less than 1 (meaning that the cost exceeds the benefit), in which case EXU defects. Cooperating with a neighbour lowers its benefit/cost ratio, while receiving cooperation raises its ratio. To bootstrap

<sup>1</sup> In all the simulations presented here,  $n = 1,000$ ,  $d = 10$  and  $c = 1$ . Three different values of  $b$  are simulated, as explained in the text.

<sup>2</sup> The results presented here use  $s = 1$  and  $r = 5$ , which appear to give the best payoff for the BT strategy in this setting.

cooperation, EXU assigns new neighbours a benefit of  $b$  and a cost of  $c$ , as if they have cooperated once and received cooperation once.

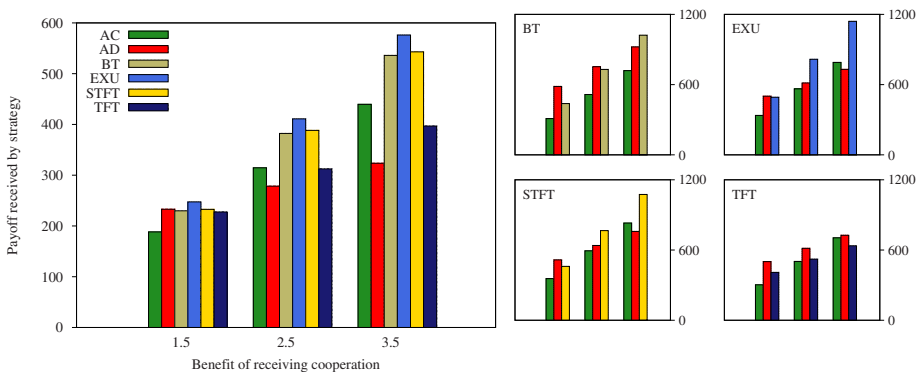
#### 4.1 Fixed Population Proportions

Each simulation consists of 20 independent runs of 2,000 rounds each. At the end of each round a randomly chosen player is removed and replaced with a new player using the same strategy, who is connected to  $d$  randomly chosen neighbours. To allow the initial conditions to fade, no measurements are taken during the first half of the run. The payoff received by each strategy is averaged over the second half of the run.


In the first set of simulations the population contains equal proportions of the four reactive strategies described above and the simple strategies Always Cooperate (AC) and Always Defect (AD). We vary the strength of the dilemma by simulating three different values of  $b$ , the benefit of cooperation; the cost  $c = 1$  is held constant. The first frame of Figure 1 shows the payoff received by each strategy for  $b = 1.5$ ,  $b = 2.5$  and  $b = 3.5$ . When  $b$  is close to  $c$ , AD narrowly outperforms all the reactive strategies except EXU. Increasing  $b$  favours the cooperative strategies at the expense of AD. For all values of  $b$ , EXU receives the highest payoff.

We also evaluate each reactive strategy separately against AC and AD, as shown in the small frames of Figure 1. Each population contains equal proportions of AC, AD and the strategy being tested. Once again the outcome depends on the value of  $b$ . For  $b = 1.5$ , none of the reactive strategies does better than AD, although EXU comes close. However, as  $b$  increases, all the reactive strategies except TFT do better than AD.

The poor performance of TFT can be ascribed to its short memory: even in a perfectly cooperative population, the probability of receiving cooperation from a given neighbour in a given round is only  $1/d$ . TFT only considers the previous round, so it is often unable to distinguish between cooperators and defectors. The other reactive strategies evaluate their neighbours over longer periods.



**Fig. 1.** Fixed population proportions. The payoff received by each strategy as a function of  $b$ , the benefit of receiving cooperation, averaged over 20 runs. The large figure shows a population containing all six strategies; the small figures show populations containing AC, AD and each of the four reactive strategies.



**Fig. 2.** Evolution of the population. The number of players using each strategy as a function of the number of rounds, averaged over 50 runs, for  $b = 1.5$  (left),  $b = 2.5$  (top right) and  $b = 3.5$  (bottom right).

All the results shown in Figure 1 are significant at the 99% level using the two-tailed Mann-Whitney U test.<sup>3</sup>

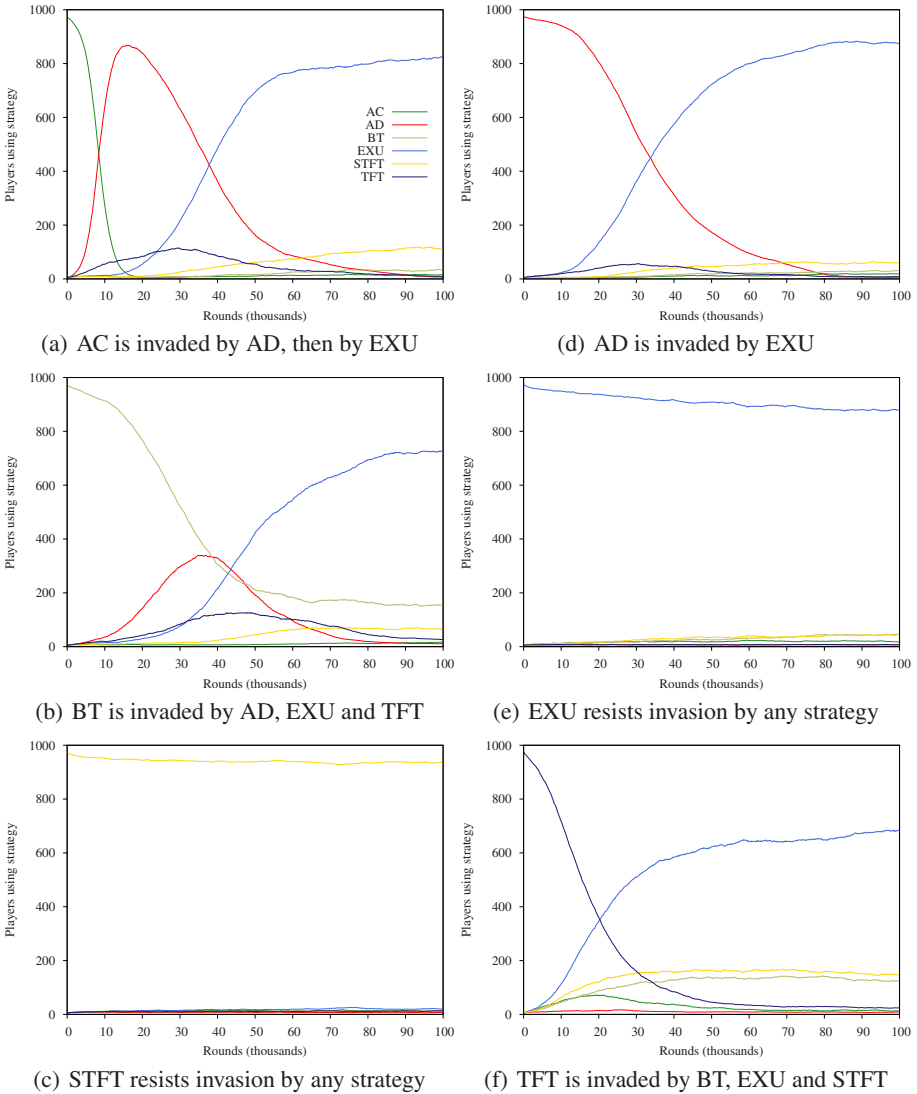
## 4.2 Evolutionary Simulations

To further investigate how each strategy fares in a mixed population, we now allow the population proportions to evolve. As before, at the end of each round a single player is chosen uniformly at random and replaced with a new player, but now the new player's strategy is chosen using the *roulette wheel method*: the probability of the new player adopting each strategy is proportional to the total payoff received by that strategy in the previous round. Thus the mixture of strategies in the population evolves according to the payoffs received: a player's payoff can be interpreted as her *reproductive fitness* [31].

To prevent any strategy from becoming extinct, if the strategy of the player being replaced is used by five or fewer players, the new player always adopts the endangered strategy. This allows strategies that are unsuccessful under certain conditions to re-emerge later in the game, and prevents strategies from being eliminated by random drift.

Figure 2 shows the evolution of the population for  $b = 1.5$ ,  $b = 2.5$  and  $b = 3.5$ , averaged over 50 independent runs of 100,000 rounds each. AC is quickly pushed to the edge of extinction: by cooperating equally with all its neighbours, it wastes resources on AD that could have been used to earn reciprocation from reactive neighbours. TFT has

<sup>3</sup> Student's t test is not suitable for making comparisons between strategies because the samples are not independently and randomly drawn from normally distributed populations. The Mann-Whitney U test was chosen because it makes fewer assumptions about the population distribution.



**Fig. 3.** Invasion simulations. The number of players using each strategy as a function of the number of rounds, averaged over 50 runs, for  $b = 1.5$ . EXU can invade any strategy except STFT; EXU and STFT both resist invasion.

trouble distinguishing between cooperators and defectors due to its short memory, and it too is quickly defeated.<sup>4</sup> Once AC has been eliminated, the other reactive strategies all outperform AD. It might seem paradoxical that any strategy would benefit from the

<sup>4</sup> This is not a weakness of the Tit For Tat strategy as such, but only of our method of adapting it to the sharer's dilemma. STFT, which is also based on the Tit For Tat strategy from the prisoner's dilemma, does not have the same weakness.



elimination of altruists, but in evolutionary simulations it is relative payoffs, rather than absolute payoffs, that are decisive: the defeat of AC makes the environment harsher for the reactive strategies, but harsher still for AD, which receives most of its cooperation from AC.

Despite their selfishness, BT, EXU and STFT succeed in establishing almost full cooperation: after 20,000 rounds the fraction of players cooperating in each round is always above 95%. As in the simulations with fixed population proportions, EXU outperforms all the other strategies for all values of  $b$ .

To test the significance of the results we use the two-tailed Mann-Whitney U test to compare the number of players using each strategy at the end of the 50 runs. All the differences between strategies are significant at the 99% level, with the exception of AD and TFT, which are not significantly different when  $b = 1.5$  or  $b = 2.5$ .

### 4.3 Invasion Simulations

The previous sections have shown that the expected utility strategy performs well when played by a substantial fraction of the population, but we would also like to know whether it is suitable for use in populations dominated by other strategies. To find out, we simulate the evolution of six populations, each containing 975 players of one strategy and five players of each of the other strategies. As before, no strategy is allowed to drop below five players. The results are averaged over 50 runs of 100,000 rounds each.

The results for  $b = 1.5$  are shown in Figure 3. EXU is able to invade any strategy except STFT. This is a strong result – clusters of Tit For Tat players can invade other strategies in networked variants of the prisoner’s dilemma [12], but in our simulations there are no clusters: new players are connected to every existing player with equal probability.

STFT and EXU can each resist invasion by any other strategy. In 50 longer runs of 1 million rounds each, EXU stabilises at 80% of the population when it starts as the dominant strategy, while STFT stabilises at 90% when it is initially dominant. This shows that each strategy does better against itself than it does against the other, making both strategies *evolutionarily stable*, at least among the strategies considered here [31].

## 5 Discussion

This paper is only a preliminary exploration of the sharer’s dilemma – many interesting aspects of the game remain to be investigated. For example, we have assumed that all players possess equal resources for cooperation, but the dynamics of cooperation may change when some players have more resources than others. Piatek *et al.* [32] have shown that when high-capacity BitTorrent nodes are able to participate in multiple swarms they benefit from allocating their resources between swarms, to the detriment of the low-capacity nodes in each swarm. The sharer’s dilemma gives us a theoretical framework for investigating whether such issues apply to cooperation under scarcity in general. We can simulate variation in capacity by updating the players at different rates, with high-capacity players making their choices more frequently than low-capacity players, allowing them to cooperate more often in a given period of time.

In the simulations presented here we have also assumed that all players assign the same subjective cost to cooperating and the same subjective benefit to receiving cooperation. This restriction is not required by the model, and we would like to explore the effect of variation within the population: for example, free riding might be an appealing strategy to players who consider the cost of cooperating to be high, while altruism might appeal to those who consider the cost to be low. However, when players can receive different payoffs from the same outcomes, it becomes difficult to compare the success of different strategies in a meaningful way, and the use of evolutionary simulations becomes problematic; we will need to consider new ways of comparing strategies before we can explore subjective payoffs.

A third simplifying assumption concerns network structure: we have only considered random graphs where a new player is connected to each existing player with equal probability. This rules out the formation of clusters, for example, which might help to establish cooperation in otherwise hostile networks; on the other hand, if players are able to choose their neighbours, defectors might be able to exploit the first-time cooperation of reactive strategies. The structure and dynamics of the network are clearly relevant to the outcome of the game, so when using the sharer's dilemma to model any scenario we will need to make sure that we are modelling the network, as well as the individual players, realistically.

## 6 Conclusions and Future Work

We have seen that a simple extension to a well-known game can provide a new perspective on the problem of cooperation in networks: incorporating scarcity into the prisoner's dilemma reframes the problem of cooperation as a problem of prioritisation and suggests new strategies based on maximising expected utility. The sharer's dilemma provides a game theoretic model for many situations in nature and society where the benefit of cooperation is higher than the cost, and where resources for cooperation are scarce.

It is easy to see how the strategies described in this paper could be applied to peer-to-peer file sharing networks such as BitTorrent; our simulations, though simplified, show that the expected utility strategy performs well in a mixed population of other strategies, indicating that it may be possible to deploy it incrementally in existing networks.

We are also interested in the possibility of using the expected utility strategy in multi-hop networks, such as peer-to-peer overlays and mobile *ad hoc* networks, to create an incentive for nodes to forward messages. However, this will require a more complex utility model that can incorporate actions whose outcome depends on the choices of other nodes. We are currently investigating the use of Savage's concept of *subjective expected utility* to choose between actions with uncertain outcomes [33].

## References

1. Gibbons, R.: Game Theory for Applied Economists. Princeton University Press, Princeton (1992)
2. Feigenbaum, J., Shenker, S.: Distributed algorithmic mechanism design: Recent results and future directions. In: Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, pp. 1–13 (2002)

3. Shneidman, J., Parkes, D.: Rationality and self-interest in peer to peer networks. In: Kaashoek, M., Stoica, I. (eds.) IPTPS 2003. LNCS, vol. 2735, pp. 139–148. Springer, Heidelberg (2003)
4. Bolton, G., Ockenfels, A.: ERC: A theory of equity, reciprocity, and competition. *American Economic Review* 90(1), 166–193 (2000)
5. Fehr, E., Schmidt, K.: A theory of fairness, competition and cooperation. Working Paper 4, Institute for Empirical Research in Economics, University of Zurich (April 1999)
6. Fehr, E., Gächter, S.: Altruistic punishment in humans. *Nature* 415(6868), 137–140 (2002)
7. Kuhn, S.: Prisoner's dilemma. In: Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy* (2003), <http://plato.stanford.edu/archives/fall2003/entries/prisoner-dilemma/>
8. Wahl, L., Nowak, M.: The continuous prisoner's dilemma: I. linear reactive strategies. *Journal of Theoretical Biology* 200, 307–321 (1999)
9. Roberts, G., Sherratt, T.: Development of cooperative relationships through increasing investment. *Nature* 394(6689), 175–179 (1998)
10. Dawes, R.: Social dilemmas. *Annual Review of Psychology* 31, 169–193 (1980)
11. Trivers, R.: The evolution of reciprocal altruism. *Quarterly Review of Biology* 46(1), 35–57 (1971)
12. Axelrod, R.: *The Evolution of Cooperation*. Basic Books, New York (1984)
13. Nowak, M., Sigmund, K.: The alternating prisoner's dilemma. *Journal of Theoretical Biology* 168(2), 219–226 (1994)
14. Nowak, M., Sigmund, K.: Oscillations in the evolution of reciprocity. *Journal of Theoretical Biology* 137(1), 21–26 (1989)
15. Félegyházi, M., Hubaux, J., Buttyán, L.: Equilibrium analysis of packet forwarding strategies in wireless ad hoc networks - the dynamic case. Technical Report IC/2003/68, EPFL (November 2003)
16. Félegyházi, M., Hubaux, J., Buttyán, L.: Nash equilibria of packet forwarding strategies in wireless ad hoc networks. *IEEE Transactions on Mobile Computing* 5(5), 463–476 (2006)
17. Nowak, M., May, R.: Evolutionary games and spatial chaos. *Nature* 359(6398), 826–829 (1992)
18. Nowak, M., Bonhoeffer, S., May, R.: Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Sciences USA* 91, 4877–4881 (1994)
19. Epstein, J.: Zones of cooperation in the demographic prisoner's dilemma (1997), Santa Fe Institute Working Paper 97-12-094
20. Ohtsuki, H., Hauert, C., Lieberman, E., Nowak, M.: A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 441, 502–505 (2006)
21. Grim, P.: The greater generosity of the spatialized prisoner's dilemma. *Journal of Theoretical Biology* 173(4), 353–359 (1995)
22. Cohen, M., Riolo, R., Axelrod, R.: The role of social structure in the maintenance of cooperative regimes. *Rationality and Society* 13, 5–32 (2001)
23. Axelrod, R., Riolo, R., Cohen, M.: Beyond geography: Cooperation with persistent links in the absence of clustered neighborhoods. *Personal and Social Psychology Review* 6(4), 341–346 (2002)
24. Nowak, M., Sigmund, K.: Evolutionary dynamics of biological games. *Science* 303, 793–799 (2004)
25. Feldman, M., Chuang, J.: The evolution of cooperation under cheap pseudonyms. In: 7th International IEEE Conference on E-Commerce Technology, Munich, Germany (July 2005)
26. Lai, K., Feldman, M., Stoica, I., Chuang, J.: Incentives for cooperation in peer-to-peer networks. In: Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA (June 2003)

27. Stutzbach, D., Rejaie, R., Sen, S.: Characterizing unstructured overlay topologies in modern P2P file-sharing systems. In: Internet Measurement Conference (IMC 2005), Berkeley, CA, USA (October 2005)
28. Bustamante, F., Qiao, Y.: Friendships that last: Peer lifespan and its role in P2P protocols. In: 8th International Workshop on Web Content Caching and Distribution, Hawthorne, NY, USA (September-October 2003)
29. Guha, S., Daswani, N., Jain, R.: An experimental study of the Skype peer-to-peer VoIP system. In: Proceedings of the 5th International Workshop on Peer-to-Peer Systems (IPTPS 2006), Santa Barbara, CA, USA, pp. 1–6 (February 2006)
30. Cohen, B.: Incentives build robustness in BitTorrent. In: Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA (June 2003)
31. Smith, J.M.: Evolution and the Theory of Games. Cambridge University Press, Cambridge (1982)
32. Piatek, M., Isdal, T., Anderson, T., Krishnamurthy, A., Venkataramani, A.: Do incentives build robustness in BitTorrent? In: Proceedings of the 4th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2007), Cambridge, MA, USA, pp. 1–14 (April 2007)
33. Savage, L.: The Foundations of Statistics. Wiley, Chichester (1954)