When Can an Autonomous Reputation Scheme Discourage Free-riding in a Peer-to-Peer System?

Nazareno Andrade[1], Miranda Mowbray[2], Walfredo Cirne[1], Francisco Brasileiro[1]

[1] *Universidade Federal de Campina Grande*
{*nazareno,walfredo,fubica*}*@dsc.ufcg.edu.br*

[2] *HP Laboratories Bristol*
miranda_mowbray@hp.com

Abstract

We investigate the circumstances under which it is possible to discourage free-riding in a peer-to-peer system for resource-sharing by prioritizing resource allocation to peers with higher reputation. We use a model to predict conditions necessary for any reputation scheme to succeed in discouraging free-riding by this method. We show with simulations that for representative cases, a very simple autonomous reputation scheme works nearly as well at discouraging free-riding as an ideal reputation scheme. Finally, we investigate the expected dynamic behavior of the system.

1 Introduction

Peer-to-peer systems can be an effective and robust way of sharing resources. However, the effectiveness of several existing peer-to-peer systems is diminished by widespread free-riding [1, 15, 16]. A peer that is a free-rider consumes resources donated by others but does not donate any resources itself. If there is a nonzero cost of donation, and the system does not discriminate between free-riders and other peers, then peers have an economic incentive to become free-riders, thus reducing the resources available for donation in the community, and diminishing the utility of the system as a whole.

One potential solution is to introduce a reputation scheme to the system. The interactions between peers affect their reputation in a way designed so that free-riders are unlikely or unable to build up a high reputation. When a peer has a resource to donate, and there are several peers requesting this resource, the peers with higher reputation are given priority. The idea is that the advantage that this gives to peers who donate resources may be enough to overcome the disadvantage given by the cost of donation.

This use of reputation differs from the classic use of reputation [14] to enhance the quality of transactions in peer-to-peer systems such as eBay [6], or to marginalize untrustworthy peers as in the systems surveyed by Ooi et al. [12]. If reputation is used to discourage free-riding then the choice to interact with a peer with high reputation is made in order to reward the peer for its previous behaviour, rather than to enhance the expected quality of the immediate transaction.

We are particularly interested in the circumstances under which it is possible to discourage free-riders using an *autonomous* reputation scheme. In an autonomous reputation scheme, peers use only local information to prioritize other peers. As such, they can only access reputation information involving peer-to-peer interactions in which they themselves have participated. The reputation of a given peer will in general be different in the eyes of different peers, and there is no attempt to reconcile these local reputations to create a global assessment. As a result, autonomous reputation schemes are relatively simple to implement, and do not require a cryptographic infrastructure or centralized storage to guarantee integrity of data retrieved from other peers, as is the case for some other reputation schemes that assign a single global reputation value to a peer. Autonomous reputation schemes are used in several peer-to-peer resource-sharing networks [2, 4, 7, 8].

An alternative way of using a reputation scheme to discourage free-riders would be not to give peers with low reputation low priority access to donated resources, but to refuse to donate resources altogether to peers with reputation below some chosen limit: if a peer had resources to donate but only peers with low reputation requested them, then the resources would remain undonated. However, this alternative would cause bootstrapping problems for an autonomous reputation scheme, because a new collaborator entering a system with autonomous reputation can only show that it is not a free-rider by donating resources, and can only detect that another peer is not a free-rider by being donated resources by that peer. Hence in this paper we consider the effect when the reputation scheme is used to prioritize donations rather than to ban certain kinds of donation completely.

In this paper we explore the design space for a peer-to-peer system in which it is possible to discourage free-riding by prioritizing resource donations using an autonomous reputation scheme.

In [2] we described an extremely lightweight autonomous reputation scheme. This was designed to promote equitable resource sharing in OurGrid, a peer-to-peer system that we are currently developing for sharing CPU cycles for bag-of-tasks applications [3]. For this autonomous reputation scheme, the reputation of peer $P1$ in the eyes of peer $P2$ is equal to the total value of resources that $P2$ has donated $P1$, minus the total value of resources that $P1$ has donated $P2$ – or is zero if this value is negative. Through simulations of representative cases, we will demonstrate the effectiveness of this reputation scheme in discouraging free-riders when the system does not exhibit eager consumption, that is, when peers in consuming state have a limit on the amount of resources that they can use with positive utility. Non-eagerness is realistic if the resource being shared is for example CPU time for applications that are not easily parallelizable, or is access to a particular software application.

We define a free-rider as a peer that does not contribute resources to the system, and a *collaborator* as a peer that does contribute resources. We say that *the system works* at time t if at that time there is a disincentive for collaborators to change their strategy to free-riding: in other words, if the expected utility for a collaborator is greater than the expected utility the collaborator would have if it changed strategy. Free-riders always have an expected utility at least as great as the expected utility that they would have outside the system, and so if the system works then the expected utility for collaborators in the system is greater than their expected utility if they left the system.

In this paper we assume that peers in consuming state can be donated resources by any collaborator in donating state. This implies that the resources are interchangeable, which can be the case if the system shares generic CPU time, or bandwidth, or storage. However the analysis of this paper may not extend to peer-to-peer systems sharing less generic resources, such as data files, because a peer requesting a specific file will not in general be able to receive it from any peer currently donating resources, only from those peers currently donating resources that have a copy of the file. However, measurements of large scale peer-to-peer file-sharing systems [11, 16] have found that a large percentage of all requests are for a relatively small number of files. For each one of these popular files, we can consider the peer-to-peer virtual subsystems consisting of requests for and donations of the file. Within each of these subsystems the resources are interchangeable. It is intuitively reasonable that if each of these virtual subsystems satisfies conditions that allow a particular reputation system to drive out free-riding, then by using the reputation system for the prioritization of donations in the file-sharing system as a whole it should be possible to discourage peers with typical resource requirements from free-riding. A more precise analysis of the circumstances in which a reputation scheme can be used to discourage free-riding in a file-sharing system is beyond the scope of this paper.

The rest of this paper is structured as follows. After a brief discussion of related work, we describe our autonomous reputation scheme and the design parameters for our system. Next we analyze the conditions on these parameters for the system to work at a fixed time, and use this analysis to predict system behaviour for some representative scenarios. We then simulate these scenarios for our reputation scheme to check that the predictions are met, and compare the performance of our scheme with that of an ideal reputation scheme in the simulated scenarios. Finally, we investigate the dynamic behaviour of the system if peers change their strategies according to their own economic interest.

2 Related Work

There recently has been an increasing amount of research in the area of reputation schemes for peer-to-peer networks, particularly for file-sharing networks. However, most of this research is on schemes that are not autonomous. For analyses of the effect on free-riding of various non-autonomous reputation schemes, see for example [5, 9, 13].

The file-distribution system BitTorrent [4] uses an autonomous tit-for-tat mechanism to decide to whom to upload, at what bandwidth. However BitTorrent does not use long-term reputation records, because the community it serves is very dynamic and long-term peer relationships are unlikely.

Autonomous reputation schemes are used in the peer-to-peer file-sharing networks eMule [7] and GNUnet [8], but we are not aware of any analytical studies of the effect of these schemes on the amount of free-riding in these networks.

One previous paper, by Lai et al., [10], analyses the effect on free-riding of an autonomous reputation scheme. In the scheme analysed, if peer A receives a request from peer B, peer A has made $r > 0$ requests to B in the past, and B has cooperated with c of these, then A cooperates with B's current request with probability c/r. If resource contention is possible, Lai et al's scheme may result in some available resources being wasted. For example, suppose that an available resource is requested by just one peer, but some past requests from the resource owner were rejected by that peer because they were for contended resources. Then the available resource may not be donated.

3 System description

In this section we describe our autonomous reputation scheme, and the design parameters for the peer-to-peer resource-sharing system that uses it.

3.1 The autonomous reputation scheme

In our autonomous reputation scheme, which was introduced in [2], each peer $P1$ keeps a local record of $V(P1, P2)$ and $V(P2, P1)$ for each peer $P2$ with which it has interacted, where $V(P1, P2)$ is the total value of resources that have been donated from $P1$ to $P2$ in the past. Each time it makes or receives a donation from $P2$, $P1$ updates this record, and recalculates its local reputation score for $P2$, which which is equal to $max\{V(P2, P1) - V(P1, P2), 0\}$. We write $r_{P1}(P2)$ to denote this score.

When peer $P1$ has spare resources that are requested by more than one other peer, it uses its local reputation scores to prioritize its donations of these resources, giving highest priority to satisfying the requests of the requesters P for which $r_{P1}(P)$ is largest.

As an illustration, suppose that peers $P1$, $P2$, $P2$ have not interacted before. We have $r_{P1}(P2) = r_{P1}(P3) = 0$. If $P2$ donates a resource with value R to $P1$, $r_{P1}(P2)$ will increase to R (whereas $r_{P1}(P3)$ will remain zero). If then $P1$ has a spare resource and has to choose between a request for that resource from $P2$ and a request from $P3$, $P1$ will choose to donate the resource to $P2$. $P1$ would make the same decision if $P3$ had interacted with $P1$ before but $r_{P1}(P3)$ was smaller than R.

3.2 System model

We consider a peer-to-peer system comprised of a set of collaborators and free-riders. At a fixed time t, a peer can be either in consuming or in non-consuming state. When in non-consuming state, collaborators donate their resources, while free-riders go idle. The design parameters that we consider for the peer-to-peer system are:

- Eagerness. We assume that for each peer there is a maximum value $C > 0$ of the utility of resources that can be consumed during a unit time interval when the peer is in consuming state. Thus, C limits the amount of resources that can be useful for a peer. The value of C is fixed for a given peer, but may vary between peers, with average value $\overline{C}$.
- The probability ρ of a peer being in consuming state. We assume that at a given time each peer has an independent probability ρ of being in consuming state.
- Cost of donation. The utility lost to the donator as a result of donation is a constant v times the utility gained by the recipient as a result of the donation, with $0 < v < 1$. If resources are available for donation but are not donated, no utility cost associated with these resources is incurred by the resource owner.
- Value of donation. When a collaborator is not in consuming state, it has resources of value D available to donate to the system. Note that D models performance as well as theoretical capacity. For example, if a peer has 1000 spare CPU cycles to donate, but its performance is bad and in practice the value it delivers when donating 1000 cycles is only as good as if it had no performance problems and donated 800 cycles, then D for the peer will be 800, not 1000. We assume that the value of D is fixed for a given peer, but can vary between different peers, with average value $\overline{D}$.
- The proportion f_t of peers that are free-riders at time t. The value f_t lies between 0 and 1. At time t, $N.f_t$ peers will be free-riders and $N.(1 - f_t)$ collaborators, where N is the total number of peers in the system.

For our analysis and simulations we will assume that all the values of the variables other than f_t are fixed over time.

The protocol for donation of resources is that collaborators that are not in consuming state donate all the resources that they have available as long as there are peers in consuming state prepared to consume them. Peers with high reputation are given priority in donations. We assume that the granularity of resources is low enough that a donating peer with at least as many spare resources as a consuming peer requests is able to give exactly the amount of resources requested, if the donating peer wishes to do so. Any resources left over, after all peers in consuming state have been donated the maximum amount of resources that they are prepared to accept, are not donated. Collaborators with resources to donate at a particular time do not have to donate them all to the same peer: they can donate resources to several different requesting peers. Free-riders that are not in consuming state stay idle.

4 Analysis

In this section we calculate the values of design parameters for which an ideal reputation system succeeds in discouraging free-riding, and use approximations to give predictions for the behaviour of sample scenarios.

Recall that we say that the system works at time t if at that time there is a disincentive to collaborators to change their strategy to free-riding. Define the *advantage to collaborators* at time t as the expected utility gain to a collaborator as a result of being in the system minus the expected utility gain to a free-rider. This is a measure of how much free-riding is discouraged at time t. It will in general be a function of f_t. The system works at time t with $f_t = f$ if and only if either $f \in (0, 1)$ and the advantage to collaborators is positive at $f_t = f$, or $f = 0$ and the limit of the advantage to collaborators as $f_t \to 0$ is positive.

Initially we pick a fixed time t, and calculate whether the system works at that time. Later on (in Section 6) we will discuss the dynamic behaviour of the system.

4.1 Analysis for fixed time

For this subsection we will assume that the system uses an ideal reputation scheme that is able to identify free-riders perfectly, that is, any free-rider always has a lower reputation than any collaborator.

Suppose at a fixed time t the total value of resources offered for donation is x_d (and that this is greater than zero), the total value of resources requested by collaborators in consuming state is x_c, and the total value of resources requested by free-riders in consuming state is x_f. We distinguish three cases, a *famine* of donations, a *glut* of donations, and the *middle* case.

The condition for famine is $x_d \leq x_c$. If this holds, then free-riders receive no donations, so gain utility zero by being in the system, whereas the set of collaborators gains a total utility $(1 - v).x_d > 0$ by being in the system. Therefore the advantage to collaborators is positive, and the system works at time t.

The condition for glut is $x_d \geq x_c + x_f$. If this holds, then all peers who make a request at time t will be donated all the resources they request. The expected utility gain for a peer resulting from the resources it is donated depends on C, but does not depend on whether the peer is a collaborator or a free-rider. On the other hand, a collaborator has an expected utility cost resulting from the resources it donates. So a collaborator can increase its overall expected utility by changing its strategy to free-riding. Therefore the advantage to collaborators is negative, and the system does not work at time t.

The condition for the middle case is that there is neither famine nor glut, ie. $x_c < x_d < x_c + x_f$. In this case the total utility gain by the set of collaborators is $x_c - v.x_d$, and the total utility gain by the set of free-riders is $x_d - x_c$. If $x_c \leq v.x_d$, then clearly the advantage to collaborators is non-positive and the system does not work at time t. Suppose $x_c - v.x_d$ is positive and $f_t \in (0, 1)$. Then the advantage to collaborators is

$$\frac{(x_c - v.x_d)}{(1 - f_t).N} - \frac{(x_d - x_c)}{f_t.N} \tag{1}$$

which is a monotonically increasing function of f_t that tends to minus infinity as $f_t \to 0$. So the system does not work at time t for $f_t = 0$, and works for $f_t \in (0, 1)$ if and only if the value of this function is positive.

So far we have assumed that there is non-eager consumption. But a similar argument can be used if there is eager consumption. Eager consumption can be modeled by putting $x_c = \infty$. This implies that the condition for famine holds, whatever the values of the other variables, and hence the system works if there is eager consumption.

In this subsection we have not assumed that the allocation of resources to requesting collaborators favours those collaborators who have donated more. However, supposing it does, if there is famine and a collaborator acquires some new spare resources additional to its original resources of value D, then the collaborator has an incentive to donate these new resources to the community, provided that doing so will not move the system out of the famine condition.

We now use the results of this analysis to pick some representative scenarios for the system parameters, and make predictions for the behaviour of the system for these scenarios.

4.2 Predictions for sample scenarios

The mean values of x_d, x_c and x_f can be expressed in terms of the design parameters as $(1 - \rho).\overline{D}.(1 - f_t).N$, $\rho.\overline{C}.(1 - f_t).N$, and $\rho.\overline{C}.f_t.N$ respectively. We can estimate whether the system will work or not at a fixed time for a given set of parameter values, by determining whether the system will work for the mean values of x_d, x_c and x_f. This is only an estimate, because the actual values fluctuate statistically about these values, but this is a reasonable approximation to make because small changes in these values will result in small changes in the utilities we calculate. (The approximation is less accurate if D varies widely between peers.)

The scenarios we choose are the ones where the parameter values satisfy $\overline{D} = 10$, $C = 9D$ for each peer, $C = D$ for each peer, or $C = D/10$ for each peer (recall that D may vary from peer to peer); $\rho \in \{0.1, 0.5, 0.9\}$; $f_t \in \{0.25, 0.5, 0.75\}$; and $v \in \{0.1, 0.4\}$. This makes a total of 3 x 3 x 3 x 2 = 54 sets of parameter values.

We have chosen these values to be realistic, to include both low and high realistic values, and to include some scenarios where the mean values of x_d, x_c and x_f are on the borderline between different cases.

Our prediction, using the estimate given by taking the mean values and applying the analysis of the previous subsection, is that among these 54 sets of parameter values, assuming perfect identification of free-riders, the system will work just for the 36 sets of parameter values that satisfy $C = 9D$, or $C = D$ and $\rho = 0.9$, or $C = D$ and $\rho = 0.5$, or $C = D/10$ and $\rho = 0.9$. For the scenarios satisfying one of the first three alternatives there is famine for the mean values, and the scenarios satisfying $C = D/10$ and $\rho = 0.9$ the mean values are in the middle case with positive advantage to collaborators.

Clearly, if the system will not work for an ideal reputation system that has perfect identification of free-riders, it should not work for a weaker reputation scheme. Our autonomous reputation scheme does not in general give perfect identification of free riders, so for this scheme we predict that the system will work for a subset of the 36 sets of parameter values identified above.

5 Simulations

We now turn to simulations for the design parameters above. The simulator simulates some aspects of a real implementation (specifically, the fluctuations over time of amounts donated and requested) that are ignored by the analytical model. However it is not a total P2P system simulator, since for example it does not deal with topology issues. We aim to investigate the effectiveness of our autonomous reputation scheme in providing a positive advantage to collaborators in the scenarios for which the analysis of the last section predicts that it is possible for a reputation scheme to do so. In order to provide a reference system, we simulated an ideal reputation scheme that perfectly identifies all free-riders.

In our simulations, the timeline is in turns, and at each turn each peer has an independent probability ρ of being in consuming state. We ran the scenarios described in Subsection 4.2 with the value of donation $D = 10$ for all peers, using our autonomous scheme and using the ideal reputation scheme.

For both reputation schemes the advantage to collaborators in the simulations was positive for 35 of the 36 scenarios the analysis had predicted it would be positive. Also, for 29 of these 35 scenarios the behavior of the the autonomous sceme in the simulations was close to the behavior of the ideal reputation scheme. Figure 1 illustrates the comparison of the advantage for collaborators between a system using the autonomous reputation sceme and the ideal reputation scheme for some of these scenarios, where the eagerness level $C = D$ and the cost of donation $v = 0.4$.

The scenarios where the difference between our autonomous reputation scheme and the ideal reputation scheme was significant were all the scenarios in which $C = 9D$ and $\rho = 0.1$. This difference is illustrated in Figure 2. The scenarios in which $C = 9D$ and $\rho = 0.1$ are on the border between the famine and the middle case, so the statistical fluctuations in x_d have a greater impact on them. Indeed, it was in these scenarios that the advantage to collaborators found in the simulation using the ideal reputation scheme differed most from the the values predicted by our analysis, and also differed most from the values found in the simulation of our autonomous reputation scheme.

The sole scenario in which the system did not work using our autonomous reputation scheme when the analysis predicted that it would using a perfect reputation scheme was the one with $C = 9D$, $\rho = 0.1$, $f = 0.25$ and $v = 0.4$. This scenario is also in the border between the famine and the middle case, and setting $f = 0.25$ and $v = 0.4$ is enough to give a negative advantage to collaborators under our autonomous scheme. Note that these two parameter values define a scenario where there is a relatively large cost of donating resources, and few free-riders to share the resources they manage to get. The ideal reputation scheme, however, did not perform much better in the simulation of this scenario: its advantage for collaborators stays fluctuating around zero when the system reaches steady state. So, according to our definition, even when using a reputation scheme that perfectly identifies free-riders, the system did not work all the time in this scenario.

As a second step, we introduced some new scenarios where peers do not have the same D (and, hence, not the same C) or the same ρ. We investigated the cases where either D or ρ is given by the uniform distributions $U(1, 19)$ or $U(0.1, 0.9)$, respectively.

When D was given by a uniform distribution with mean 10 there was the same overall behavior as when D was set equal to 10 for all peers, and making ρ different for different peers had only a slightly greater impact. Although the mean value of ρ was equal to 0.5 in all our scenarios, the difference between the performance of the system in providing incentives for collaborations using our autonomous reputation scheme and using an ideal reputation scheme was greater in the scenarios where different peers had different values of ρ. More specifically, the statistical fluctuations in x_d made the difference in performance greater for the scenarios where $C = D$ and $f = 0.25$ and made the system not work when using our autonomous reputation scheme in the two scenarios where $C = D$, $v = 0.4$, $\rho = U(0.1, 0.9)$ and $f \in \{0.25, 0.5\}$. Once again, these scenarios are on the border between the famine and middle cases. The statistical fluctuations in x_d arising from the differing values of ρ regularly pushed the system into the middle case, in which our autonomous reputation scheme is less efficient at rewarding collaborators. When combined with the high cost of donation $v = 0.4$, the effect was that the advantage to collaborators was negative for our autonomous reputation scheme in these scenarios.

Still, the system using our autonomous reputation scheme performed similarly to one using an ideal reputation

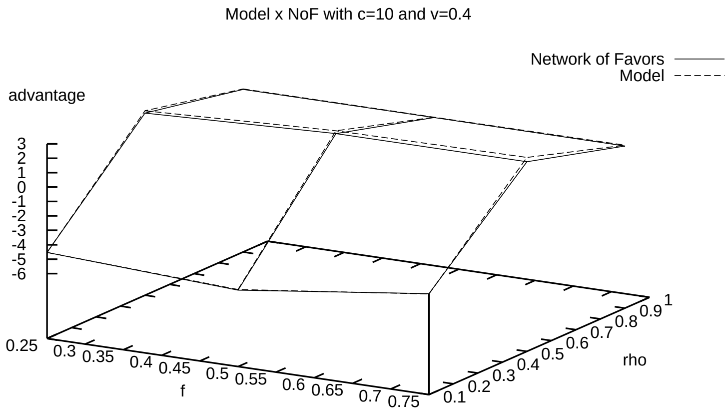

Figure 1. Advantage to collaborators in the scenarios where $C = D$ and $v = 0.4$, varying ρ and f.

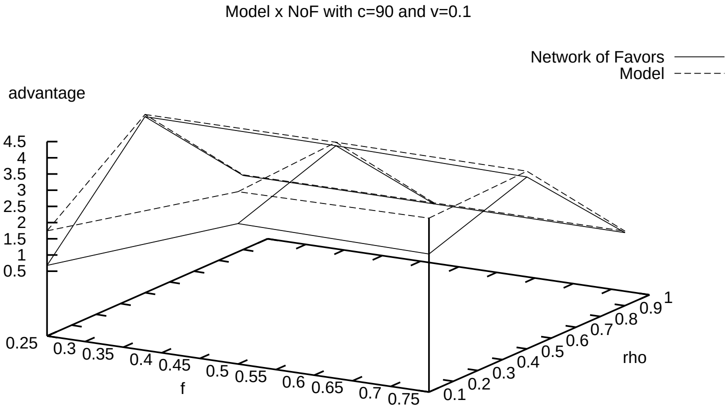

Figure 2. Advantage to collaborators in the scenarios where $C = 9D$ and $v = 0.1$, varying ρ and f.

scheme in almost all scenarios of our experiment where we predicted that any reputation had a chance of being effective (including 35 out of our original 36 scenarios with fixed ρ and D). With the exception of three scenarios where our autonomous reputation scheme proved to have a slightly different tolerance for non-contention of resources, it made the system work whenever an ideal reputation scheme would. Moreover, for the great majority of scenarios, there was only a very small difference between the measured advantage for collaborators in a system using our autonomous reputation scheme and in a system using the ideal reputation scheme.

This shows, at least for the scenarios that we simulated, that in most of the cases where it is possible to use an ideal reputation scheme, it is also possible to use our autonomous scheme without a great loss in performance. With the exception of the three border scenarios with large donation costs, although it sometimes did not perform as well as the ideal reputation scheme, our autonomous scheme still managed to provide a positive advantage to collaborators whenever the ideal reputation scheme did so. We had imagined that more complex centralized reputation schemes with global assessments of reputation would give a significantly greater advantage to collaborators than our autonomous reputation scheme, but this appears not to be the case.

Note that our comparisons between the two reputation schemes were all made after the system using our autonomous reputation scheme had reached a steady state. A system using our autonomous scheme requires some time to reach a steady state in which it has an accurate identification

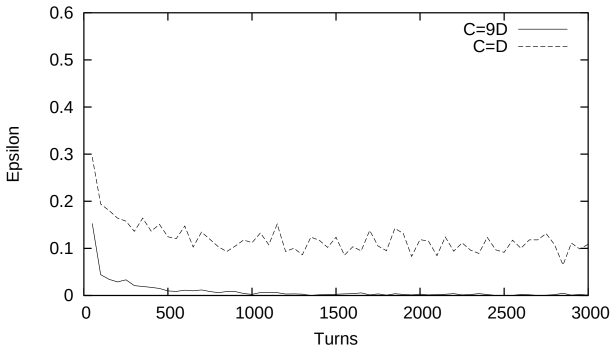

Figure 3. The proportion of resources donated to free-riders using our autonomous reputation scheme for $f = 0.5$, $\rho = 0.5$ and different values of C

of free-riders. As a result, before the system reaches this state the reputation scheme should have less effect than the perfect reputation system in discouraging free-riding. We now investigate whether the eagerness level C has an impact on the time needed for the system to reach the steady state.

Figure 3 shows the proportion of the available resources that were donated to free-riders in the last 50 turns, which we denote ϵ. When the system is in famine, ϵ expresses how well the community has identified the free-riders. We found that C does not impact on the time needed for reaching the steady state, but on the actual value of ϵ that the system shows when in steady state. We found that, except for the scenarios where $\rho = 0.9$, ϵ is approximately inversely proportional to C. When $\rho = 0.9$, although ϵ is greater for $C = D/10$ than for $C = D$ and $C = 9D$, it is very similar in scenarios where peers have one of these two last eagerness levels. We suspect that in practice, when they have one of these two eagerness levels, the consumers act as eager consumers for the system. Thus, our observations show that eagerness makes it easier for our autonomous reputation scheme to identify free-riders, although we have already found that a high level of eagerness is not a necessary condition for the system to work.

6 Analysis of dynamic system behaviour

Now we consider the effect of allowing peers to change their strategies. The value of f_t will vary over time according to strategy choices, whereas the values of all other system parameters are fixed over time. We assume that peers change their strategy in their own best interest, choosing to be either a free-rider or a collaborator so as to maximize their expected utility. So the gradient of f_t is positive at t if the advantage to collaborators is negative at time t, and negative if the advantage to collaborators is positive at time t.

We do not need to offer peers a third option of leaving the system, because free-riders always have an expected utility at least as great as they would have outside the system. We assume that the choice of strategy is binary, that is, peers either choose to be a collaborator and offer all their spare resources to the community, or to be a free-rider and offer none: we do not consider the option of peers offering some but not all of their available resources.

We give a general analysis under the assumption that the values of C and D do not vary widely between peers. If this is not the case, then the system is harder to analyse in general, but we describe one example with heterogeneous peers for which it is possible to determine the system's dynamic behaviour.

6.1 Analysis of dynamic behaviour

For this subsection we assume that the values of C and D do not vary widely between peers. More specifically, we assume (as in Subsection 4.2) that we can determine whether the system is in famine, glut or the middle case for the system at time t, and whether the advantage to consumers is positive or negative, by calculating the case and the sign of the advantage to consumers for the expected values of x_d, x_c and x_f at that time. In practice statistical fluctuations may temporarily move the system into another case, but we assume that these excursions are sufficiently rare and short-lived that they can be ignored when we are considering the large-scale long-term dynamic behaviour of the system.

Figure 4 illustrates the dynamics of the system given these assumptions. The ratio of the mean values for x_d and x_c is independent of f_t, so is fixed over time. Therefore if famine holds for these mean values for the initial state of the system, it continues to hold for the subsequent evolution of the system. As free-riders become collaborators in a system with a famine of donations, the average amount of resources donated to the system at a given time increases, but the average amount consumed at that time also increases, because the new collaborators are donated (and hence consume) more resources than when they were free-riders. This increase in consumption is large enough to keep the system in famine. So if the system is initially in famine it should remain in famine as f_t decreases, (except for rare excursions given by statistical fluctuations), and eventually all the free-riders become collaborators.

If the system is initially in glut then f_t will increase until eventually the system is no longer in glut - at this point it will be in the middle case. At the border between the glut and middle cases free-riders have a higher expected utility than collaborators.

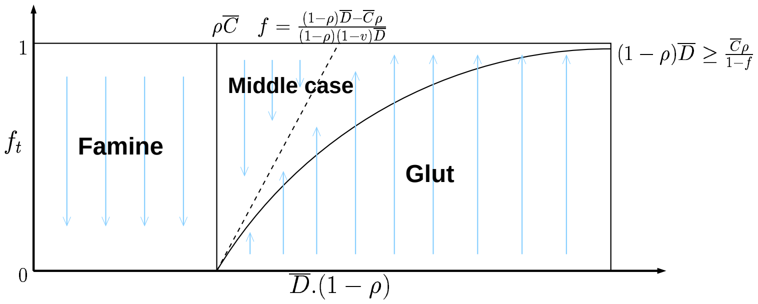

Figure 4. Dynamics of the system, varying f_t and $\overline{D}.(1-\rho)$.

For the middle case, either the advantage to collaborators is negative for all $f_t \in (0,1)$, in which case collaborators will eventually die out, at which point the system will not work, or else it is a monotonically increasing function of f_t for $f_t \in (0,1)$ which tends to $-\infty$ as $f_t \to 0$ and to a positive number as $f_t \to 1$. For this second alternative the system will evolve to a stable equilibrium at which $f_t > 0$ and free-riders and collaborators have equal expected utilities. At the stable equilibrium the system does not work, because the advantage to collaborators is zero. In practice the system may oscillate around the stable equilibrium rather than reaching it precisely, because of the statistical fluctuations in the value of resources donated and requested; but the average over time of the advantage to collaborators for the oscillating system will still be zero.

It follows that if there is not a famine of donations, the proportion of free-riders will evolve over time to a value at which the system does not work. It is possible that the system initially is in a state in which it works, if the conditions for the middle case hold and the advantage to collaborators is positive, but it will eventually evolve to an equilibrium state in which the system does not work. This happens even though in this analysis we are assuming perfect identification of free-riders. On the other hand, if there is a famine of donations, then the system will work and free-riders will eventually die out.

This gives a relatively simple heuristic for checking if the system is eager enough for a reputation system to have a chance of acting to drive out free-riding by prioritizing donations to peers with high reputation: assuming that all peers choose to free-ride or not so as to maximize their expected utility, free-riding can only be driven out if there is a famine of donations, ie.

$$\overline{D}.(1-\rho) \leq \overline{C}.\rho \tag{2}$$

In our simulations, peers did not change strategies. However for most of the scenarios we simulated, after some time interval the effect of the autonomous reputation scheme was similar to that of a scheme with perfect prioritization, and this held for a range of values of f. It is therefore reasonable to expect the eventual behaviour of a community under our autonomous reputation scheme to be similar to that predicted by this analysis, provided that the community has enough time to identify free-riders when peers change strategy.

6.2 An example with heterogeneous peers

In this paper we have assumed that the values of C and D do not vary very widely between different peers. In particular we have assumed that the standard distributions of these values are small and that it is sensible to discuss their means. For some potential application systems this assumption is unrealistic. If the values do vary widely our previous analysis will not apply, although it may still be possible to predict the system behaviour.

For example, begin with a network in which the values of C and D do not vary widely between peers, and the average amount of resources donated in a given timeslot is less than half than the amount requested. Clearly, there is a famine of donations. Now make the set of peers heterogeneous by adding a new peer that is in donating state at least half the time, and has such a high value of D that when it is in this state it is able to meet all the unfulfilled requests in the system. The new peer's expected utility is greater outside the system than as a contributor within the system, even if all the other peers donate all their free resources to the new peer when it is in consuming state. So the new peer will become a free-rider. The remaining peers will eventually detect that this has happened, and after that the system will evolve as in our analysis, except that there is one additional free-rider (the new peer): eventually, all other peers will become contributors.

Essentially, our reputation scheme encourages the donation of resources by a peer only if there is a reasonable chance of the donation being repaid.

For distributions of C, D with large deviations which are not as extreme as in this example (for example, Zipf distributions, which we have investigated), the system behaviour is hard to analyse, as it is contingent on the behaviour of peers with high values of C and D. More research needs to be done.

7 Conclusion

In this paper we have demonstrated through simulations that in a network of similar peers, our autonomous reputation scheme is sufficient to discourage free-riding when there is a famine of donations. Our analysis of a system model indicates that when there is not a famine of donations, no reputation scheme should be able to discourage free-riding by prioritizing donations to peers with high reputations.

In our simulations of the autonomous reputation scheme for sample scenarios in which there was a famine of donations, there was an incentive for collaborating as opposed to being a free-rider in almost all scenarios where an ideal reputation scheme would also provide such an incentive. The cases where the results where different showed that our autonomous reputation scheme has a slightly worse prioritization than the ideal reputation scheme and therefore requires slightly more contention for resources to keep the utility of free-riders low. For the majority of the scenarios, both schemes performed similarly. The autonomous scheme discourages free-riding by prioritizing donations almost as well as the ideal reputation scheme, despite being very lightweight and easy to implement, and requiring neither central coordination nor a cryptographic infrastructure.

References

[1] ADAR, E., AND HUBERMAN, B. A. Free riding on Gnutella. *First Monday 5*, 10 (2000). http://www.firstmonday.dk/.

[2] ANDRADE, N., CIRNE, W., BRASILEIRO, F., AND ROISENBERG, P. OurGrid: An approach to easily assemble grids with equitable resource sharing. In *Proceedings of the 9th Workshop on Job Scheduling Strategies for Parallel Processing* (June 2003).

[3] CIRNE, W., BRASILEIRO, F., SAUVÉ, J., ANDRADE, N., PARANHOS, D., SANTOS-NETO, E., MEDEIROS, R., AND SILVA, F. Grid computing for Bag-of-Tasks applications. In *Proceedings of the IFIP I3E2003* (September 2003).

[4] COHEN, B. Incentives build robustness in bitTorrent. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems* (June 2003).

[5] DELLAROCAS, C. Efficiency and robustness of binary feedback mechanisms in trading environments with moral hazard. Working Paper 4297-03, MIT - Sloan School of Management, January 2003.

[6] EBAY, INC. Reputation - eBay feedback: Overview, 1995-2003. http://pages.ebay.com/help/confidence/reputation-ov.html.

[7] EMULE-PROJECT.NET. eMule site. http://www.emule-project.net/.

[8] GROTHOFF, C. An Excess-Based Economic Model for Resource Allocation in Peer-to-Peer Networks. *Wirtschaftsinformatik* (June 2003).

[9] KUNG, H. T., AND WU, C.-H. Differentiated admission for peer-to-peer systems: incentivizing peers to contribute their resources. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems* (June 2003).

[10] LAI, K., FELDMAN, M., STOICA, I., AND CHUANG, J. Incentives for cooperation in peer-to-peer networks. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems* (June 2003).

[11] LEIBOWITZ, N., RIPEANU, M., AND WIERZBICKI, A. Deconstructing the KaZaA network. In *3rd IEEE Workshop on Internet Applications* (June 2003).

[12] OOI, B. C., LIAU, C., AND K.L.TAN. Managing trust in peer-to-peer systems using reputation-based techniques. In *International Conference on Web Age Information Management* (August 2003).

[13] RANGANATHAN, K., RIPEANU, M., SARIN, A., AND FOSTER, I. To share or not to share: An analysis of incentives to contribute in collaborative file sharing environments. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems* (June 2003).

[14] RESNICK, P., ZECKHAUSER, R., FRIEDMAN, E., AND KUWABARA, K. Reputation systems. *Communications of the ACM 43*, 12 (2000), 45–48.

[15] RIPEANU, M., AND FOSTER, I. Mapping the Gnutella network: Macroscopic properties of large-scale peer-to-peer systems. In *First International Workshop on Peer-to-Peer Systems* (2002).

[16] SAROIU, S., GUMMADI, P. K., AND GRIBBLE, S. D. A measurement study of peer-to-peer file sharing systems. In *Proceedings of Multimedia Computing and Networking 2002* (January 2002).