A Reciprocation-Based Economy for Multiple Services in Peer-to-Peer Grids

Miranda Mowbray[†], Francisco Brasileiro[‡], Nazareno Andrade[‡], Jaindson Santana[‡], Walfredo Cirne[‡]

[†]*Hewlett-Packard Laboratories Bristol*
Filton Rd, Stoke Gifford
Bristol BS34 8QZ, UK
miranda.mowbray@hp.com

[‡]*Universidade Federal de Campina Grande*
Av. Aprígio Veloso, s/n - 58.109-970
Campina Grande, Brazil
{*fubica,nazareno,jaindson,walfredo*}*@dsc.ufcg.edu.br*

Abstract

In this paper we study reciprocation-based mechanisms to encourage donation in peer-to-peer grids in which multiple services, such as processing power and data transfers, are shared explicitly. We have modeled such a system and established how peers should assess whether it is profitable to exchange services with another peer, an issue that is not present in the single service case. Unfortunately, this assessment relies on information provided by untrustworthy peers. As an alternative, we have extended, to the case of multiple services, a reciprocation-based mechanism which uses only reliable information gathered locally. We have assessed this mechanism by simulating scenarios in which services are exchanged that are combinations of two different basic services. In the explored scenarios the mechanism performs very well, and can marginalize free riders even when the cost to peers of donating a service is nearly as large as the utility gained by receiving it.

1. Introduction

A computational grid is a federation of sites across different administrative domains which shares computational services. A peer-to-peer grid is a large-scale, free-to-join computational grid in which participants do not necessarily know or trust each other. In a peer-to-peer grid, each peer represents a site; whenever the site has idle capacity it provides services to other peers, and whenever the site needs more computational services, it requests them from all other peers. Examples of projects building such grids are OurGrid [9, 17], Cluster Computing On the Fly [16] and the Self-organizing Flock of Condors proposed by Butt et al. [7].

Peers gain utility from using the services of other peers, and there is a cost for providing a service to the grid. This cost is incurred to maintain the hardware and software needed by a site to provide services to the grid, and through the security risk posed to this site when providing such services to non-trusted parties. These costs must be low in comparison with the benefit obtained by joining the grid, and they can be lowered by reducing the effort needed to install and maintain the grid middleware, as well as by improvements to security mechanisms. However, it is naive to assume that all costs can be eliminated.

In contrast to most grids currently in production, a peer-to-peer grid cannot rely on off-line negotiations or trust chains comprising all participants in the grid to enforce co-operative behavior. In this setting, an incentive mechanism plays a key role in promoting service provision to the grid. If there is a non-zero cost for donating services and a peer can obtain the same amount of service no matter how much it serves other peers, then peers have an economic incentive to contribute nothing and free ride. This behavior is indeed found in peer-to-peer file sharing systems [2, 14]. Free riding reduces the amount of service available in the grid, and diminishes the utility of the system for users of resource-intensive applications; most users of computational grids run such applications. This motivates the use of an incentive mechanism to promote collaboration in peer-to-peer grids.

In previous work, we proposed the Network of Favors, a reciprocation-based mechanism for peer-to-peer grids in which a single service is shared [3, 4]. In the Network of Favors, peers exchange donations of services. Peers always donate their spare service, and decide whom to serve based solely on the record of their past bilateral interactions with peers requesting the service.

We have shown that in a peer-to-peer grid sharing a single service – access to processing power – this autonomous behavior provides an incentive for peers to contribute as much as possible [3, 4]. This happens because peers who contribute more get more in return when they make requests. Since the balance of past interactions is very little information, a peer can keep track of its interactions with a very large number of other peers. The Network of Favors is currently implemented in a peer-to-peer grid named OurGrid [9, 17], which is in production since December 2004 (see status.ourgrid.org).

The Network of Favors is particularly suited for peer-to-peer grids because, by keeping the behavior of each peer completely autonomous, it can be implemented without depending on any centralized mechanism or trust infrastructure. In this paper we consider how to create a reciprocation-based economy grounded on the Network of Favors in which peers provide multiple services to the grid. Our main motivation comes from practical experience with the deployment of OurGrid. Users of this grid have manifested the need for the system to consider incentives not only for the provision of processing power, but also for data transfers and storage. This happens because it is not possible to abstract all services as a single one. For instance, some users run data-intensive applications and compete for the capacity of other peers to receive data, while other users run applications with only light data requirements.

Introducing multiple services in an economy complicates matters because, in such a setting, peers may value services differently. Thus, it is in the interest of peers to choose trading partners not only based on their behavior (i.e. the likelihood of reciprocation), but also on the way these peers value the services they provide and consume (i.e. how profitable it is for a peer to maintain a long-term reciprocation relationship with another peer).

We study the problem posed by this requirement to identify what is necessary to extend the prioritization model of the Network of Favors to a system sharing an arbitrary number of services, so that we can implement such an extension in our peer-to-peer grid. We believe our results might also be useful when extending other systems that use similar exchange-based mechanisms, such as BitTorrent [10].

The structure of this paper is as follows. We situate our work among related efforts in Section 2. In Section 3 we discuss how the Network of Favors should be extended to deal with multiple services. In particular, we discuss an issue that is not present in the single service case; that it may be unprofitable for two peers who are not free riders to exchange services. In Section 4, we model the problem and discuss how peers can assess whether it is profitable to interact with another peer, under the assumption that the other peer is not a free rider. It turns out that to do this, a peer must rely on information provided by untrustworthy peers. Since this is not desirable, we consider an alternative, in which interactions with unprofitable peers may occur, but a peer donates services preferentially to peers with whom it expects to have beneficial interactions in the future, based on its direct past experience. In Section 5 we present simulation results under several scenarios for the exchange of services that are combinations of two different basic services – for instance, the basic services might be processing power and storage. In our simulations the mechanism using only local information performs very well, and can marginalize free riders even when the costs to peers of donating services are high. We present our conclusion in Section 6.

2. Related Work

There is a body of research on using market-based mechanisms to regulate service provision in resource sharing systems [1, 5, 8, 15, 20]. Markets can provide incentives for provision in a flexible and robust way. However, we argue that market-based mechanisms are not suitable for the scenario of peer-to-peer grids that we have described. Participants in peer-to-peer grids do not trust each other, and there is no centralized or widely trusted entity in the system. Market-based resource allocation mechanisms rely on contracts, auditing, banking and electronic cash payment systems, which are very difficult to deploy with such constraints on the system. Shneidman et al. [18] discuss other open issues in getting market-based resource allocation into production.

We originally proposed the Network of Favors [3, 4] as an alternative to market-based mechanisms. The Network of Favors is designed not to depend on the existence of banking, trust or negotiation among peers. This mechanism is similar to the tit-for-tat mechanism used in BitTorrent, where peers exchange pieces of a file based on their past bilateral interactions [10]. However, neither of these works address the case in which multiple services are exchanged.

It has been suggested, as part of the vision of utility computing, that all computing services should be described in terms of a single unit, the *computon*, just as electricity is sold by the kilowatt-hour. The price for a computon would vary according to supply and demand [19]. This would reduce the economic problem of the provision of multiple computing services to the provision of a single service, measured in computons. However, attempts to produce an agreement between the few largest computing suppliers of exactly how the computon should be defined have not been successful [13]. Agreement between suppliers on how to define the computon is even more problematic in a peer-to-peer grid.

3. The Network of Favors for Multiple Services

The Network of Favors has been previously explored in a system where a single service is shared [3, 4]. In this paper we discuss how it can be extended to a system in which peers may provide more than one service and may differ in how they value services.

The basic idea of the Network of Favors is that peers prioritize the requests they receive based solely on the record of their past interactions with the requesters. As a result, there is no need to trust other peers or a central entity in order to assess the global reputation of each requester.

The extension presented here is designed for a system of peers, in which each peer owns a set of resources and can provide multiple services with them. All peers alternate independently between periods where they have spare resources, and periods where they have demand for services that cannot be immediately met by their resources. We call a peer that currently has a spare resource a *provider*, and the work done by a provider for another peer is called a *favor*, which in the general case may be any combination of the services available in the system.

When there are multiple services, a favor of one type of service may be repaid in another type of service, and peers may value services differently. The main design challenge this imposes for the mechanism is the need for peers to choose with which other peers they should interact. Providers need to decide both with which other peers it is worthwhile to interact in the long run, and which of the current requests for their services they should prioritize. They do so using a long-term and a short-term policy, respectively.

The long-term policy allows peers to protect their overall utility, by not donating to another peer if they assess that the expected effect on their overall utility resulting from a long-term interaction with this peer is unsatisfactory.

The short-term policy governs to whom a provider decides to donate a favor, when there are several other peers which are not excluded by the long-term policy and which are currently requesting favors. This decision is based on information from interactions with these peers in the past. Moreover, whenever there is no contention for the services of a provider, it donates these services to any peer requesting them that is not excluded by the long-term policy. This serves as a bootstrap for the exchange of favors.

4. Incentive Mechanisms

In this section we describe two incentive mechanisms, by giving the details for each of a long-term and a short-term policy governing the allocation of favors. As we will show, the effect of these policies is that there is an incentive to donate.

For the first incentive mechanism, which we name *PosInt* for *positive interactions*, the long-term policy forbids interactions between pairs of peers unless their long-term interaction should be profitable, under the assumption that neither are free riders; and the short-term policy makes a peer donate services preferentially to peers with whom it expects to have beneficial interactions in the future, based on its past interactions and on knowledge of how the other peers value services. It turns out that knowledge of how other peers value services is also necessary for the long-term policy for *PosInt*.

The second incentive mechanism, which we name *ExtNetFav* for *extended Network of Favors*, does not require this knowledge; it relies only on peers' direct experiences and on how they themselves value services, and thus is closer in spirit to the Network of Favors. Its long-term policy is the trivial policy allowing interaction with all peers, so for reciprocation it relies on its short-term policy, under which a peer donate services preferentially to peers with whom it expects to have beneficial interactions in the future, based on its past interactions and its own (but not others') valuation of services.

Our experience is that in practice it can be difficult for peers in a peer-to-peer grid to predict the utility that they, or others, would gain by being donated a particular favor at a point in the future. We therefore do not assume that peers can estimate these utilities. However, costs are easier for peers to predict. We assume that peers can easily determine the costs that they would incur when donating a favor. These costs will in general vary between different peers, and for different types of favors.

4.1. Notation and assumptions

Here are the general assumptions that we make about the system.

G1. The system is a peer-to-peer system in which peers independently decide whether or not to do favors requested by other peers. Peers are content to participate in the system if their expected future net utility gain as a result of being in the system (that is, utility gain from being donated favors minus utility loss from the cost of donating favors to others) is positive.

G2. Each peer A can accurately estimate the utility cost $v_A(f)$ that it would incur if it provided favor f for another peer. Costs are additive, i.e. the cost of providing favor $f1$ and then favor $f2$ is $v_A(f1) + v_A(f2)$. For all non-zero favors f, the utility cost $v_A(f)$ is positive.

G3. The utility to a peer A of receiving a favor f, written $u_A(f)$, may vary over time, but always satisfies $u_A(f) > v_A(f)$. That is, the utility to A of receiving a favor that it requests is greater than the cost to A of donating the same favor to another peer.

G4. If a peer A tries to pay back a favor to peer B, it will eventually succeed in doing so. This implies that eventually A will be able to provide a service at a time that it is requested by B, and that the granularity of requests for services by B can be made small enough that a request is eventually not too large for A to satisfy.

G5. Some peers in the system are *collaborative*, and follow the algorithm specified. However, some are *non-collaborative*, and choose alternative behavioral strategies in order to maximize their expected net utility gain. In particular, one strategy that they may consider is *free riding*, that is, requesting and consuming favors from the system but donating no favors to the system. They may spread false information about other peers, or about themselves, and may conspire with other non-collaborative peers to make their falsehoods more plausible. Moreover, a non-collaborative peer can *whitewash* its identity, i.e. it can use a different identity for every interaction with other peers. Thus, there may be a very high churn of non-collaborative peers.

G6. There is low churn of collaborative peers in the system. (Note that without this assumption, our mechanisms would not work.)

4.2. The long-term policy for $PosInt$

Let $\overline{f}_A$ be the average favor that a peer A requests to the other peers. This represents a probability distribution of the types of favors requested by A over the long term. For instance, if in a typical time interval A requests on average one unit of basic service s_1 and three units of basic service s_2, then $\overline{f}_A$ will be $(s_1 + 3s_2)/4$.

Now, suppose peer A is deciding whether or not to interact with B. For a long-term interaction in which A donates $n.\overline{f}_B$ to B and B donates $m.\overline{f}_A$ to A for some m, n to be beneficial to both A and B, it needs to be the case that both $m \cdot u_A(\overline{f}_A) - n \cdot v_A(\overline{f}_B)$ and $n \cdot u_B(\overline{f}_B) - m \cdot v_B(\overline{f}_A)$ are positive. Such $m, n > 0$ exist if and only if

$$u_A(\overline{f}_A) \cdot u_B(\overline{f}_B) > v_A(\overline{f}_B) \cdot v_B(\overline{f}_A) \qquad (1)$$

It is difficult for A to estimate the functions u_A or u_B, but by assumption G2, A knows the function v_A; let us assume for now that A also knows v_B. By assumption G3, $u_A(f_A) > v_A(f_A)$ and $u_B(f_B) > v_B(f_B)$. Therefore A can check whether the following more stringent inequality holds:

$$v_A(\overline{f}_A) \cdot v_B(\overline{f}_B) \geq v_A(\overline{f}_B) \cdot v_B(\overline{f}_A) \qquad (2)$$

Roughly speaking, this inequality means that the cost to the pair of peers A and B of producing the average favors that they request themselves is greater than that of producing the requested favors for each other. Note that if $\overline{f}_A = \overline{f}_B$ or $v_A = v_B$, the inequality is automatically satisfied. In particular, in the special case that all requested favors are multiples of a single service, then all pairs of peers can gain by interacting.

For pairs of peers A, B for which inequality (2) holds, both peers know that they can benefit from a long-term exchange of favors, and so both have an incentive to initiate such an exchange of favors. When A has spare resources it will look for peers B requesting favors for which this inequality is satisfied, and use the spare resources to grant favors to one of these peers.

Pairs of peers for which (2) does not hold cannot tell whether or not it is possible for them both to benefit from a long-term exchange of favors. Conservatively, they do not interact, so as to avoid being drawn into an interaction that decreases their utility.

4.3. The short-term policy for $PosInt$

Now suppose that there is more than one peer B requesting favors with which the long-term decision function of a provider A does not prevent interactions. How does A decide which of these peers to donate to? The answer is that A donates to whichever of these peers it expects to gain most from interacting with; A calculates this using its short-term policy.

The short-term policy uses the *cost balance* from the previous interactions between peers. Each peer keeps a record, for each other peer with which it interacts, of this number, which is calculated as follows. Before the two peers have ever interacted the cost balance is equal to zero (although this value is not explicitly recorded). If peer A donates a favor f to peer B, then A decreases its cost balance for interactions with B by the cost $v_A(f)$ of producing this favor – or, if the original cost balance was less than $v_A(f)$, peer A sets its cost balance for B to zero. Meanwhile, B increases its cost balance for A by $v_A(f)$. Note that different peers will in general record different cost balances for the same peer A.

When choosing which peer to donate a favor to, provider A donates the favor to the candidate peer with the highest cost balance, where the candidate peers are the peers requesting the favor for whom inequality (2) holds. If all candidates have cost balances equal to zero, the peer chooses one of the candidates randomly.

4.4. The policies for $ExtNetFav$

An obvious caveat of the way the policies for $PosInt$ are calculated is that they require knowledge of the average favors requested by other peers ($\overline{f}_A$), and also the costs that other peers would incur when donating a favor ($v_A(f)$). Other peers are not necessarily trustworthy, and may be able to increase their expected utilities by giving false information about these. In the absence of reliable information on other peers' costs, it is not feasible to use a long-term policy which distinguishes profitable collaborative peers from unprofitable ones. The alternative used by the incentive mechanism $ExtNetFav$ is to remove this check, or equivalently to use a trivial long-term decision function that never prevents two peers from interacting, and rely simply on the short-term function to marginalize both free riders and unprofitable collaborative peers.

In addition, the way the cost balance accounts for favors received should also be changed. The approach that we use for $ExtNetFav$ is instead of a peer using the unknown cost function of a favor's provider to calculate the cost balance, it uses its own cost function. Thus, the short-term policy for $ExtNetFav$ differs from that of $PosInt$ in that when B receives favor f from A, B increases its cost balance for A by $v_B(f)$ rather than by $v_A(f)$.

4.5. Consequences of the policies

Whether the cost balance is calculated in the way specified by the short-term policy for *PosInt* or by the short-term policy for *ExtNetFav*, the cost balance that A records for B decreases when A donates favors to B (provided that it was not zero to start with) and increases when B donates favors to A. It can therefore be thought of as an indication of the net benefit that A has gained so far by interacting with B. By donating the favor to the candidate peer with the highest value of the cost balance, provider A is choosing to interact with a peer with whom it expects to have beneficial interactions in the future, based on its past experience. A similar way of selecting peers to interact with is described (for a different context) in [6].

The cost balance is greater than or equal to zero for all peers. If it is zero for all candidate peers, A still donates the service to one of the candidate peers. This allows newcomers to the system to have a chance of donating and receiving favors. It also serves as a bootstrap mechanism.

Free riders will sometimes be donated favors, and A will lose utility as a result of any donations that it makes to a free rider. However, A will not donate a favor to a free rider unless all the candidate peers have zero cost balance. On the other hand, the way that the cost balance is calculated has the effect that if collaborative peer A ever donates a favor to collaborative peer B, then throughout the subsequent history of the system either A's cost balance for B will be positive, or B's cost balance for A will be positive, or both will be positive. Thus the system promotes continuing interactions between collaborative peers.

Since donating a favor to a collaborative peer increases the expected amount of favors received from that peer in the future, over the long run the more a peer donates to the system the more it can expect to receive back.

Free riders receive services with low priority. Thus the expected long-term utility gain for free riders should be lower than that for peers that do not free ride, and so non-collaborative peers will choose not to free ride.

Since we do not assume that there is a limit on the number of different IDs that a peer can use in the system, free riders can whitewash their identities, and thus become indistinguishable from collaborative newcomers. However, they cannot increase their chances of obtaining services from any peer A by doing so, because A's cost balance for their new identity will be zero, the minimum value. Peers can only increase their priority for donations from others by making donations to others themselves. As soon as a peer makes a donation it becomes distinguishable from a free rider, so collaborative peers should not remain isolated indefinitely.

Another consequence of peers' potential ability to use multiple identities is that it is particularly difficult to design a reliable global reputation system, because a non-collaborative peer has the possibility of creating very many clones of itself which propagate false reputations [12]. Our incentive for donation does not rely on a global reputation system.

5. Evaluating the Network of Favors for Multiple Services

In this section we evaluate the performance of the incentive mechanisms proposed in the previous section. Our evaluation is based on simulations of the two mechanisms discussed. Before analyzing the results attained from the simulations, we present the model implemented by our simulator, the metric measured and the scenarios simulated.

5.1. System Model

We consider a grid comprised of N peers which offer and consume services whose resource requirements are combinations of two different basic services, s_1 and s_2. For instance, the two basic services might be processing power and storage. Peers can either be collaborative or free riders. In our simulations, the timeline is in turns, and at each turn a peer can be either in consuming or in non-consuming state. When in non-consuming state, collaborative peers donate the use of their spare resources, while free riders go idle. In addition to not donating any services, we assume that free riders change their identities at each turn. The design parameters that we consider for the system are:

- *Frequency of consumption*. We assume that at a given turn each peer has an independent probability ρ of being in consuming state.
- *Service availability*. $D = (d_1, d_2)$, where d_1 and d_2 are, respectively, the maximum amount of services s_1 and s_2 that a peer is able to donate in a given turn.

- *Relative favor consumption profile*. For a peer A, $\pi_A = (\pi_1^A, \pi_2^A)$ is such that $\overline{f_A}$ is a multiple of $\pi_1^A \cdot s_1 + \pi_2^A \cdot s_2$.
- *Relative cost of donation*. For a peer A, $\kappa_A = (\kappa_1^A, \kappa_2^A)$ measures how costly it is for A to provide a unit of basic services s_1 and s_2. A favor f provided by A that corresponds to x units of service s_1 and y units of service s_2 has a cost to A equal to $v_A(f) = x \cdot \kappa_1^A + y \cdot \kappa_2^A$.
- *Prevalence of free riding*. ϕ is the number of peers that free ride and change their identities, divided by the total number of peers. The peers who do not free ride, collaborate.
- *Incentive mechanism*. The incentive mechanism is either $PosInt$ or $ExtNetFav$.

We assume that peers in consuming mode are able to consume as much service as the providers are able to offer them, restricted only by the consumption profiles of the consuming peers. Our practical experience has shown us that this is generally the case in a computational grid. In each turn all providers are selected in a random order and donate as much as possible of their available services. Each provider performs the following steps: (i) uses the long-term policy to select among all peers in consuming state which are the peers with whom it may interact; (ii) uses the short-term policy to select a consuming peer to donate its services; (iii) donates as much service as it can to the selected peer; and, (iv) updates the corresponding cost balance. Collaborative consuming peers also update their corresponding cost balance.

5.2. Performance metric

The metric M that we use to evaluate the performance of system is

$$M = \frac{\sum_{C \in \mathcal{C}} v_C(f_{prov}(C))}{\sum_{C \in \mathcal{C}} v_C(f_{rec}(C)) - \frac{|\mathcal{C}|}{|\mathcal{F}|} \sum_{F \in \mathcal{F}} v_F(f_{rec}(F))}$$

where $f_{rec}(A)$ and $f_{prov}(A)$ are the accumulated amount of favors that a peer A has received and provided, respectively, $\mathcal{C}$ is the set of collaborative peers and $\mathcal{F}$ the set of free riders. The metric M was designed by considering the special cases where, for some Ω, utilities are given by:

$$u_A(f) = \Omega \cdot v_A(f) \text{ for all peers } A \text{ and favors } f \tag{3}$$

M is the minimum value of Ω such that on average free riders gain no more from the system than collaborators. This is because if (3) holds, then the average benefit to a collaborator from being in the system, minus the average benefit to a free rider, is $\frac{1}{|\mathcal{C}|} \sum_{C \in \mathcal{C}} [\Omega \cdot v_C(f_{rec}(C)) - v_C(f_{prov}(C))]$ minus $\frac{1}{|\mathcal{F}|} \sum_{F \in \mathcal{F}} [\Omega \cdot v_F(f_{rec}(F))]$. It can be easily checked that this difference is positive if and only if $\Omega > M$.

By assumption G3 in Subsection 4.1, $\Omega > 1$. It follows that if M is close to 1, the system is very effective at marginalizing free riders.

5.3. Scenarios

We chose scenarios in which the parameter values satisfy $D = (10, 10)$; $\rho \in \{0.1, 0.5, 0.9\}$; and, $\phi \in \{0.25, 0.50, 0.75\}$. We chose these parameter values to cover a variety of scenarios, to include both low and high realistic values. Moreover, we chose the values κ_A and π_A (with A ranging over all the peers) so to create three settings, each with different profitable relationships among collaborative peers: (i) the interactions between any two collaborative peers are mutually profitable (*single-set*); (ii) there are two disjoint sets of collaborative peers with the same cardinality, such that interactions between collaborative peers of the same set are mutually profitable, while interactions of collaborative peers in different sets are mutually unprofitable (*full-mutex*); and, (iii) there are two disjoint sets of collaborative peers with the same cardinality, such that interactions between collaborative peers of the same set are mutually profitable, while interactions of collaborative peers in the first and second sets are profitable only for the peers on the first set (*half-mutex*). For each of these settings, the free riders are uniformly distributed over all the sets. Also, in all settings the total number of peers in the system is 600. (This relatively small value, compared to the potentially many thousands of peers that a peer-to-peer grid may encompass, is a conservative choice for the evaluation. For both mechanisms, free riders only get served when no resources are being requested by a consumer with positive cost balance. Thus, the larger the number of peers, the larger the number of consumers, and the easier it is to marginalize free riders.) We simulated all these scenarios both for $PosInt$ and for $ExtNetFav$.

Recall that inequality (2) is always satisfied if either $\pi_A = \pi_B$ or $\kappa_A = \kappa_B$, for any two peers A and B. Thus, a trivial way to generate a single set of peers whose pairwise interactions are always mutually profitable is to have all peers either with the same π or with the same κ. However, to have more diversity among peers in this setting, we have divided them in two subsets - say $\mathcal{P}$ and $\mathcal{V}$, such that in subset $\mathcal{P}$ all pairs of peers A and B, $A, B \in \mathcal{P}$, have $\pi_A = \pi_B = (1, 1)$, while in subset $\mathcal{V}$, all pairs of peers A and B, $A, B \in \mathcal{V}$, have $\kappa_A = \kappa_B = (2, 2)$. The values of κ_1^A and κ_2^A for a peer A, $A \in \mathcal{P}$ and π_1^B and π_2^B for a peer B, $B \in \mathcal{V}$ are drawn from uniform distributions. To make the interactions of any peer $A \in \mathcal{P}$ and any peer $B \in \mathcal{V}$ mutually profitable, we have set $\kappa_1^A < \kappa_2^A$ and $\pi_1^B > \pi_2^B$. Moreover, $\kappa_1^A + \kappa_2^A = 4$ and $\pi_1^B + \pi_2^B = 2$.

For the *full-mutex* and *half-mutex* settings, we have divided the peers in two subsets - say $\mathcal{A}$ and $\mathcal{B}$, such that all peers A, $A \in \mathcal{A}$, have $\pi_A = (2, 1)$ and all peers B, $B \in \mathcal{B}$, have $\pi_B = (1, 2)$. Then, we have drawn κ_A and κ_B from uniform distributions such that: (i) for the *full-mutex* setting $1 < \kappa_1^A < 2 < \kappa_2^A < 3$ and $1 < \kappa_2^B < 2 < \kappa_1^B < 3$; and, (ii) for the *half-mutex* setting $1 < \kappa_2^A < 5$, $\kappa_1^A = (\kappa_2^A - 1)/2$, $5 < \kappa_2^B < 10$ and $\kappa_1^B = \kappa_2^B/2$.

5.4. Analysis

We ran enough simulations to reduce the error to 1% with a 95% confidence interval. Table 1 summarizes the results for M in all scenarios simulated.

As can be seen, for most cases the performances of $PosInt$ and $ExtNetFav$ are equivalent, considering the 1% margin of error. In particular, for the *single-set* setting, $PosInt$ and $ExtNetFav$ perform equally well in all scenarios simulated. Since in this case their long-term decision functions have the same output (all peers have mutually profitable interactions), this shows that the different ways of calculating their cost balances can lead to equivalent results.

In the other two settings (*full-mutex* and *half-mutex*) there are pairs of peers that have unprofitable interactions. These peers never interact under $PosInt$, while under $ExtNetFav$ they occasionally do.

For the *full-mutex* setting, there are cases where $PosInt$ is better than $ExtNetFav$ by 3% or more (see the entries in bold in Table 1). In these cases there is either a high probability of peers being in consuming state ($\rho = 0.9$), or a large number of free riders ($\phi = 0.75$). For larger ρ, there are more consumers, and hence free riders are more easily marginalized. Although peers under $ExtNetFav$ are as able as peers under $PosInt$ to marginalize free riders, the former occasionally provide favors to unprofitable collaborators, while the latter do not. This explains why, in these cases, peers under $ExtNetFav$ perform a little worse than those under $PosInt$.

When ϕ is large and ρ is not large, peers donate more to free riders. Additionally, peers under $ExtNetFav$ also lose utility by donating to unprofitable collaborators. Therefore, in these cases peers under $ExtNetFav$ have a higher probability of participating in interactions that will not be profitable. For instance, when $\rho = 0.5$ and $\phi = 0.75$, from the 300 peers with which a peer under $PosInt$ may interact, 225 of them are free riders and interactions with them will result in utility loss. On the other hand, peers under $ExtNetFav$ may interact with any of the 600 peers, and will lose utility in interactions with 575 of them. However, it is important to notice that even with the disadvantage of not being able to identify unprofitable peers, $ExtNetFav$ performs very close to $PosInt$. For the scenarios we evaluated, the difference is never greater than 5%.

Table 1. Summary of simulation results

Setting	ρ	ϕ	M for PI	M for ENF
single-set	0.1	0.25	1.02	1.02
single-set	0.1	0.50	1.03	1.03
single-set	0.1	0.75	1.06	1.06
single-set	0.5	0.25	1.01	1.01
single-set	0.5	0.50	1.01	1.01
single-set	0.5	0.75	1.02	1.02
single-set	0.9	0.25	1.02	1.02
single-set	0.9	0.50	1.03	1.03
single-set	0.9	0.75	1.05	1.05
full-mutex	0.1	0.25	1.02	1.04
full-mutex	0.1	0.50	1.03	1.05
full-mutex	**0.1**	**0.75**	**1.05**	**1.09**
full-mutex	0.5	0.25	1.01	1.02
full-mutex	0.5	0.50	1.01	1.02
full-mutex	**0.5**	**0.75**	**1.01**	**1.04**
full-mutex	**0.9**	**0.25**	**1.01**	**1.04**
full-mutex	**0.9**	**0.50**	**1.02**	**1.05**
full-mutex	**0.9**	**0.75**	**1.03**	**1.08**
half-mutex	0.1	0.25	1.02	1.02
half-mutex	0.1	0.50	1.03	1.03
half-mutex	0.1	0.75	1.05	1.06
half-mutex	0.5	0.25	1.01	1.01
half-mutex	0.5	0.50	1.01	1.01
half-mutex	0.5	0.75	1.02	1.02
half-mutex	0.9	0.25	1.01	1.02
half-mutex	0.9	0.50	1.02	1.03
half-mutex	0.9	0.75	1.04	1.05

The *full-mutex* case is a very constrained setting in which every pair of peers either have mutually profitable interactions or will both lose utility in a long-term reciprocation relationship. In the less constrained *half-mutex* setting, the fact that only half of the peers lose utility in the long-term interactions with half of the other peers makes the performance of $ExtNetFav$ equivalent to that of $PosInt$ again.

More importantly, in most cases the value of M for both $PosInt$ and $ExtNetFav$ is very close to 1, and in all cases it is less than 1.11. If utilities are given by (3), this implies that there is an incentive for non-collaborative peers to change their strategy from free riding to providing services, provided that the cost to a peer of donating a favor is less than nine-tenths of the utility gained by the peer if it receives the same favor. In practice, the cost of providing a service in a peer-to-peer grid is typically small compared to the utility gained. The simulation scenarios in [4], which were chosen to be realistic based on practical experience with a system running the Network of Favors, satisfy (3) and have the cost of donation equal to at most 0.4 times the utility gained.

We have also performed simulations with the 600 peers unevenly distributed among the sets. We have simulated the same scenarios presented in Table 1 for the *full-mutex* and *half-mutex* cases, considering 100 peers in one set and 500 in the other, as well as 200 in one set and 400 in the other. For the *full-mutex* case the results attained were equivalent to those when peers were evenly distributed in the two sets. For the *half-mutex* case there was only one scenario ($\rho = 0.9$ and $\phi = 0.75$) for which $ExtNetFav$ performed worse than $PosInt$. As discussed before, the scenarios with a large number of free riders or a high probability of peers being in consuming state are the less favorable for $ExtNetFav$. Nevertheless, in all cases the difference between the two mechanisms remained not larger than 5%.

In our design of $ExtNetFav$, we dealt with the problem of unreliable second-hand information by avoiding the use of such information altogether. A different approach to unreliable information in peer-to-peer systems is to use majority voting by peers to ascertain which information is correct (see e.g. [6, 11]). However, this will fail if fewer than half the peers are collaborative, because the non-collaborative peers can collude to outvote peers that provide truthful information. In contrast, note that the incentive mechanisms presented here perform well even in the cases for which $\phi = 0.75$: in these cases only a quarter of the peers are collaborative.

6. Conclusions

In this paper we have designed and evaluated mechanisms for promoting service provision in a peer-to-peer grid in which several services are available. We have extended the Network of Favors, which is a mechanism to provide incentives using a simple autonomous behavior of the peers in the system. Simulation results show that the extended Network of Favors has an excellent performance, even when the costs of providing services are high and fewer than half the peers are collaborative.

References

[1] D. Abramson, R. Buyya, and J. Giddy. A Computational Economy for Grid Computing and its Implementation in the Nimrod-G Resource Broker. *Future Generation Computer Systems J.*, 18(8):1061–1074, 2002.

[2] E. Adar and B. A. Huberman. Free Riding on Gnutella. *First Monday*, 5(10), Oct. 2000.

[3] N. Andrade, F. V. Brasileiro, W. Cirne, and M. Mowbray. Discouraging Free Riding in a Peer-to-Peer CPU-sharing Grid. In *Proc. 13th IEEE Symp. High Performance Distributed Computing, Honolulu, Hawaii*, pages 129–137, June 2004.

[4] N. Andrade, M. Mowbray, W. Cirne, and F. V. Brasileiro. When Can an Autonomous Reputation Scheme Discourage Free-riding in a Peer-to-Peer System? In *Proc. 4th Workshop on Global and Peer-to-Peer Computing (GP2PC 04), Chicago, USA*, Apr. 2004.

[5] A. AuYoung, B. N. Chun, A. C. Snoeren, and A. Vahda. Resource allocation in federated distributed computing infrastructures. In *1st Workshop on Operating System and Architectural Support for the On-demand IT InfraStructure (OASIS'04), Boston, MA*, Oct. 2004.

[6] D. Banerjee, S. Saha, S. Sen, and P. Dasgupta. Reciprocal resource sharing in P2P environments. In *4th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'05), Utrecht, The Netherlands*, pages 853–869, 2005.

[7] A. Butt, R. Zhang, and Y. Hu. A Self-Organizing Flock of Condors. In *Proc. Supercomputing 2003, Phoneix, Arizona, USA*, Nov. 2003.

[8] R. Buyya and S. Vazhkudai. Compute Power Market: towards a market-oriented grid. In *First IEEE/ACM Int'l Symp. Cluster Computing and the Grid, Brisbane, Australia*, pages 574–581. IEEE Computer Society Press, May 2001.

[9] W. Cirne, F. Brasileiro, N. Andrade, L. Costa, A. Andrade, R. Novaes, and M. Mowbray. Labs of the world, unite!!! *Journal of Grid Computing*. To appear. Available as technical report at http://www.ourgrid.org/twiki-public/pub/OG/OurPublications.

[10] B. Cohen. Incentives Build Robustness in BitTorrent. In *Proc. Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, June 2003.

[11] E. Damiani, S. D. C. di Vimercati, S. Parabosci, and P. Samaranti. Managing and Sharing Servents' Reputations in P2P Systems. *IEEE Trans. Data and Knowledge Engineering*, 15(4):840–854, July/Aug. 2003.

[12] J. R. Douceur. The Sybil Attack. In *Int'l Workshop on Peer-to-Peer Systems*, 2002.

[13] M. Hall. Utility pricing. *Computerworld*, *21 Nov.*, 2005.

[14] D. Hughes, G. Coulson, and J. Walkerdine. Freeriding on gnutella Revisited: the bell tolls. *IEEE Distributed Systems Online*, 6(6), Feb. 2005.

[15] K. Lai, B. A. Huberman, and L. Fine. Tycoon: A distributed market-based resource allocation system. Technical report, HP Labs, Apr. 2004. http://arxiv.org/abs/cs.DC/0404013.

[16] V. M. Lo, D. Zappala, D. Zhou, Y. Liu, and S. Zhao. Cluster Computing on the Fly: P2P Scheduling of Idle Cycles in the Internet. In *Peer-to-Peer Systems III: 3rd Int'l Workshop, IPTPS 2004*, volume 3279 of *Lecture Notes in Computer Science*, pages 227–236. Springer Berlin / Heidelberg, Feb. 2004.

[17] OurGrid Project. OurGrid site. http://www.ourgrid.org/, 2003–2006.

[18] J. Shneidman, C. Ng, D. Parker, A. AuYoung, A. C. Snoeren, and A. Vahda. Why markets could (but don't currently) solve resource allocation problems in systems. In *10th USENIX Workshop on Hot Topics in Operating System*, June 2005.

[19] The Economist. Who wants to buy a computon? 12 Mar., 2005.

[20] R. Wolski, J. Plank, and J. Brevik. G-commerce: Market formulations controlling resource allocation on the computational grid. Technical report UT-CS-00-450, University of Tennessee at Knoxville, Oct. 2000.