GriPhyN/iVDGL 2004-46

Incentive Mechanisms for Large Collaborative Resource Sharing

Kavitha Ranganathan[1] , Matei Ripeanu[1] , Ankur Sarin[2] and Ian Foster[1,3]

[1] Department of Computer Science, The University of Chicago, Chicago, IL 60637, USA
[2] Irving.B.Harris School of Public Policy, The University of Chicago, IL 60637, USA
[3] Math and Computer Science Division, Argonne National Laboratory, Argonne, IL, 60439, USA

Abstract

We study the nature of sharing resources in distributed collaborations such as Grids and peer-to-peer systems. By applying the theoretical framework of the multi-person prisoner's dilemma to this resource sharing problem, we show that in the absence of incentive schemes, individual users are apt to hold back resources, leading to decreased system utility. Using both the theoretical framework as well as simulations, we compare and contrast three different incentive schemes aimed at encouraging users to contribute resources. Our results show that soft-incentive schemes are effective in incentivizing autonomous entities to collaborate, leading to increased gains for all participants in the system.

1. Introduction

Two recent directions in harnessing large amounts of distributed resources are 'Grid' computing [1, 2] and 'Peer-to-Peer' computing [3] . Both involve projects [4] [5, 6] developing infrastructure for the pooling of resources like data, storage or computation. A common assumption in these projects is that resource owners have committed their resources to the system and the chief task is to integrate and use them efficiently. Since a Grid deployment would typically be composed of many different administrative domains, and hence a varied set of interests, the following fundamental question arises: *Are individual resource owners (in the case of the Grid, individual administrative domains) willing to share their personal resources for the overall good of the community?* Experiences [7-10] with peer-to-peer (P2P) file sharing systems like Gnutella, Napster and Kazaa confirm our concerns that by and large, resource owners are not altruistic. In Gnutella for example 70% of all users do not share files, and 50% of all requests are satisfied by the top 1% sharing hosts.

Likewise, in the Grid community, efforts to bring together resources from different administrative domains have recognized a need to introduce some kind of negotiations. One approach is to formulate sharing policies offline and try and enforce them as and when resources need to be allocated. Another approach points to a market economy [11, 12] where resources are priced and could be bought and sold using one of the traditional market models. Both these approaches reiterate the fact that motivating users to contribute their resources may be critical to eventual success of such systems.

In this paper we take a step towards understanding the performance of incentive schemes that can be incorporated in a distributed collaboration. As a first step, we analyze sharing one kind of resource, namely data files, with the aim of extending our findings to other resources like storage and compute cycles.

Multiple users may want to share files for various reasons. For example a group of physicists solving a common problem might want access to the same raw-data for analysis. If one user creates a local copy of some of the data, it is beneficial for the whole system to have access to this replica.

Our approach involves so-called "soft incentive schemes" or non-pricing schemes [13] to facilitate sharing in such scenarios. We study three schemes: two soft-incentive, reputation-based schemes: Peer-Approved and Service-Quality and a more rigid scheme, Token-Exchange, for the Grid scenario.

We describe and apply (in Sections 2 and 3) the general analytical framework of Schelling's Multi-Person Prisoner's Dilemma (MPD) [14] to the specific case of distributed file-sharing systems. The analysis helps explain the rational behavior of individuals in a file sharing community, without an incentive mechanism in place. We then introduce Peer-Approved as the incentive mechanism and analyze the behavior of users, with the help of the MPD model. The study shows that a soft-incentive scheme like Peer-Approved is effective in incentivizing rational users to share more files (Section 4).

We then use simulations (as this allows us to look beyond the assumptions of the MPD model) to measure the effectiveness of the reputation based soft-incentive mechanisms like Peer-Approved, in comparison to a pricing scheme like Token Exchange (Section 5). Our results show that even simple soft-incentive schemes can motivate users of a resource sharing collaboration to increase contributions in a way that benefits all users, including themselves. We then discuss work in a few different related fields (Section 6) and conclude with directions for future work (Section 7).

2. Introduction to the Prisoner's Dilemma

The sharing problem can be modeled as a Prisoner's Dilemma [15]. In the classical Prisoner's Dilemma (CPD), two players choose simultaneously whether or not to co-operate. Each is rewarded if both cooperate, but at a lower rate than the penalty they receive if one cooperates and the other does not. Hence the dilemma: the rational choice of not cooperating leaves both worse off than if they had co-operated.

Not all assumptions underlying the CPD apply to our environment. First, we must extend the framework to more than two users. Second, players (different resource providers) in a typical Grid setting can observe actions of others and make choices influenced by others. Hence, the assumption that players act simultaneously needs to be relaxed. Third, the incentives to contribute/co-operate may depend on how many other users are contributing; this number needs to be incorporated into the theoretical model. The resulting Multi-Person Prisoner's Dilemma (MPD) framework provides a more realistic model (albeit with other limitations) of a resource sharing environment.

The following four conditions define a MPD [14]:

1. There are ***n*** people in the system, each with the same binary choice and payoffs.
2. Each person has the same preferred choice, which does not change, no matter what other people do.
3. A person is always better off if more among the others choose the un-preferred alternative.
4. For a certain $k > 1$, if k or more players choose the un-preferred alternative, they are better off than if all players had chosen the preferred alternative.

Figure 1(a) illustrates the situation graphically, showing the payoff curves for a player that chooses the preferred (P, upper line) and un-preferred (U, lower line) alternative, as a function of the number of other players, from 0 to n, that choose the un-preferred alternative (to share files in our case). We assume that there are $n+1$ players in total, and hence n "others." E.g., at $x = {}^{n}/_{3}$, a third of the other players choose the un-preferred alternative and two thirds choose the preferred alternative; P_x and U_x are the pay-off to the player that chooses the preferred or un-preferred alternatives, respectively.

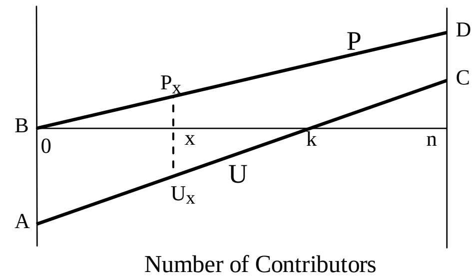

Figure 1: MPD model. Payoff curves for a player that chooses the preferred (P, upper line) and un-preferred (U, lower line) alternative, as a function of the number of other players that choose the un-preferred alternative.

While the payoff functions can be linear or curved, depending on the specific problem, the MPD definition implies that the vertical order of the four end-points of the payoff curves remain the same. That is, if everyone else chooses the preferred alternative (point A), the payoffs are the least, and if everyone else chooses the un-preferred alternative, the payoff is the highest (point D). Point C should be higher than point B, so that after some critical value ***k***, it is more profitable to join the un-preferred group than when everyone was in the preferred group (point B).

3. Applying Multi-Person Prisoner's Dilemma to File Sharing

Collaborative file-sharing systems can be modeled as MPDs. We make the following simplifying assumptions: there are ***n*** users (equivalent to ***n*** administrative domains) in the system at all times, and each user has one (unique) file that she can either decide to share with others (the un-preferred option) or keep only to self (the preferred option). All files are the same size, are equally popular and require unit bandwidth to download/upload. The benefit to a user participating in the system is the access gained to a wide range of files made available by other users in the system. Note that, in the absence of an incentive mechanism, contributors and non-contributors derive the same benefit. (In our analysis, we exclude altruism as a possible benefit for contributors.) On the other hand, a contributor incurs a cost, namely the bandwidth consumed by others when they download a file from her.

Given the above assumptions, suppose there are ***c*** contributors and ***r*** free-riders/non-contributors ($c+r=n$). Also suppose ***f*** files are requested in the system during a time unit. Thus, assuming files have similar popularity, the expected number of file requests (Π) reaching a contributor **x** is:

$$\Pi = \frac{Num.\ of\ \ files\ at\ x \times Num.\ of\ \ requests}{Total\ num.\ of\ \ files\ in\ the\ system} = \frac{f}{c}$$

Hence, the expected *cost* for a contributor is a function of $^f/_c$, while for a free rider it is always 0, as she does not share files.

In the absence of an incentives mechanism the *benefit* for all users is the same, i.e., the total number of files available in the system, which is a function of *c*, the number of contributors. We make this function logarithmic to model the intuition that the incremental benefit for gaining access to a new file decreases as the number of files available in the system increase. (Note that this could be any increasing function without changing our model). Thus, the net payoff for a contributor $payoff_c = log(c)-(f/c)$ and the net payoff for a free-rider $payoff_r = log(c)$.

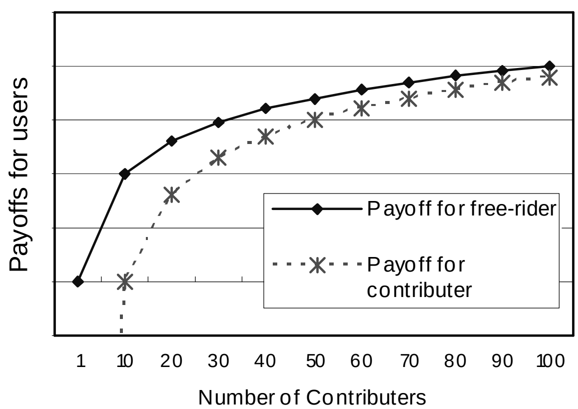

Figure 2: Payoffs for contributers and free-riders.

Figure 2 plots the net payoff curves for free riders and contributors for n = 100 and f = 10. The Y axis represents the net payoff. Note that the vertical order of the four end-points of the payoff curves in Figure 2 is consistent with the MPD definition (Figure 1). $Payoff_r$ dominates $payoff_c$, , that is at any given state of the system, a user receives higher payoff if she does not contribute. System equilibrium [16] (the state where no one has an incentive to deviate from their action given the choices made by others) is thus on the extreme left, where nobody contributes. Note that this equilibrium is *inefficient* since users could have obtained higher payoffs had they all made the opposite choice and chosen to share their files. Thus in the absence of incentives, the rational choice of not contributing, leaves users worse off than if they had contributed.

4. Incentive Mechanisms

File sharing can be 'incentivized' using either *pricing policies* that involve an explicit payment for every file transferred or *non-pricing policies* (also called soft-incentive schemes) that encourage sharing in other ways.

We describe three schemes: Token-Exchange which is similar to a pricing scheme and two non-pricing schemes: Peer-Approved, and Service-Quality.

- ***Token-Exchange***: In this scheme, a consumer of a file must transfer a token to the supplier prior to download. To enable newcomers to use the system, each first-time user might be allotted a fixed number of tokens, but once these run out, the user has to serve files to earn tokens.

This scheme is similar to a pricing scheme with fixed prices, as a user must decide for each potential download whether the file in question is worth a token. Moreover, each file exchange incurs the cost of transferring and validating the token. One of the systems implemented along these lines is Mojo Nation [17].

- ***Peer-Approved***: In this scheme, a reputation system is used to maintain ratings for users, and users are only allowed to download files from others with a lower or equal rating. This strategy motivates users to increase their rating in order to gain access to more files. User ratings can be based on different metrics: e.g., the number of files advertised by a user or the number of file-requests served by a user.

First time users without files to share should be allowed to download a small number of files so that they can enter the system and gradually build their rating.

This scheme is more flexible than Token-Exchange in that a user need not take a decision every time she wants a file. Moreover, it has been suggested that non-pricing schemes may be more practical to implement in certain kinds of collaborative networks [13], than direct payments between users. Past work also suggest that users may prefer (and thus accept more quickly) schemes that do not require payments or decisions for each transaction [18]. However, Peer-Approved needs a secure and reliable mechanism for maintaining user reputations [19, 20]. In the following, we assume the existence of such a mechanism and focus on the incentive policies that may be layered on it.

- ***Service-Quality****:* This third scheme also uses a reputation mechanism. In contrast to Peer-Approved, users advertise all their files and send download requests to any other user, but these requests are assigned to service classes and served accordingly.

Combinations of these schemes are also possible. For example, in the Paris Metro Pricing scheme [21], suggested initially for providing differentiated services in packet networks, a number of service quality classes are defined and users are assigned to one class based on how much they are willing to pay for the service.

5. Performance of Peer-Approved Policy

We want to compare the effectiveness and fairness of soft-incentive policies such as Peer-Approved and Service-Quality with those of pricing policies such as Token-Exchange.

As a first step, we analyze the performance of *Peer-Approved* using the MPD model (We will study *Service Quality* and other soft-incentive schemes in the future.) We use our theoretical model to determine if such an incentive scheme would indeed motivate rational users to share more.

However, since our model has its limitations (binary choices for users, lack of heterogeneity in user characteristics, and is unable to capture dynamics) we also perform simulations to test the incentive scheme under a wider range of possible resource distributions and user behaviors. We also use these simulations to compare the performance of Peer-Approved and Token-Exchange.

5.1. Theoretical Analysis

Let all users have $D=h*d$ files. In order to conform to the binary choice model, we assume two choices for users. Users can either advertise all their files or a fixed fraction ($^1/_h$) of their files, that is d files. A user's rating is the number of files that she advertises. Note that the *Peer-Approved* based on the number of advertised files policy would immediately exclude true free riders from the system as users are only allowed to download files from others with a lower or equal rating. Thus we assume that all users advertise at least part of the files available locally.

The analysis follows as in Section 3 (Table 1 describes the variables used).

Table 1: Variables used in analysis

$D=h*d$	Num. of files advertised by a contributor
d	Num. of files advertised by a partial contributor
f	Total num. of **f**iles requested at unit time
u	Num. of contributors (**u**n-preferred alternative)
p	Num. of partial contributors (**p**referred alternative)
n	Total num. of users ($n=p+u$)

Recall that ***u*** users choose the unpreferred alternative (to share all $d*h$ files) and ***p*** choose the preferred alternative (to share only d files). If ***f*** file requests are made per unit time, the expected number of file requests (Π_x) reaching a user ***x*** is:

$$\Pi_x = \frac{Num.\,of\ \ files\ advertised \times Num.\,of\ \ requests}{Total\ num.\,of\ \ files\ advertized\ in\ the\ system}$$

Thus, the expected cost for a contributor is:

$$c_c = \frac{hd*f}{hdu+dp} = \frac{hf}{hu+p}.$$

Similarly, the expected cost for a partial contributor is $c_p = \dfrac{f}{hu+p}$ and the expected benefit for both is *log(hdu+dp)*. Figure 3-left plots the payoff curves for both kind of users, and by the same logic used earlier, users tend to conglomerate at an inefficient equilibrium on the extreme left.

We now introduce the Peer-Approved incentive policy according to which contributors will only serve other contributors where as partial contributors serve both kinds of users (and thus all users in our scenario). The expected number of requests originating from a contributor is

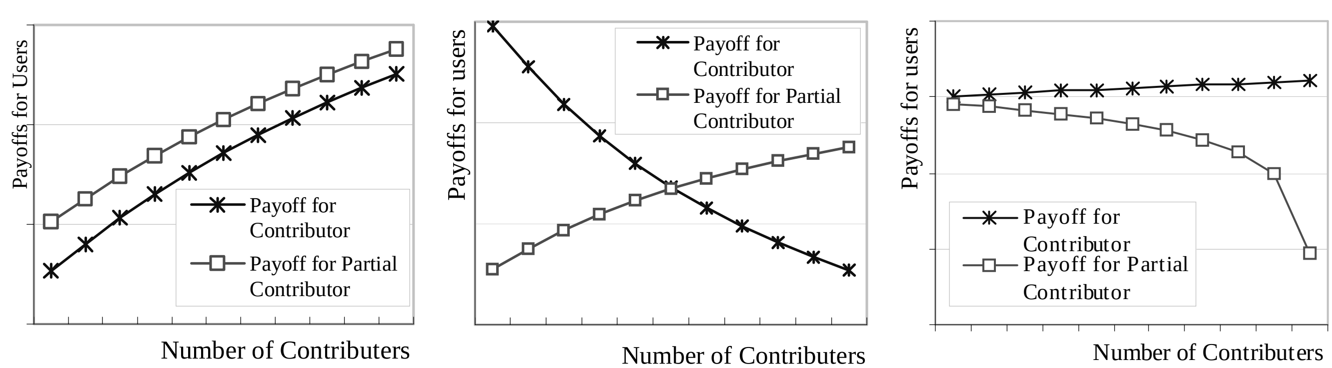

Figure 3. The net payoff curves for both camps (n=100, d=1, f=10, h=2) when (left) there is no incentive policy in place, (middle) peer-approved incentive with the benefit function increases very slowly with the number of files a user can access (right) the benefit function is log(number of files accessible by a user).

$\frac{uf}{n}$ and the expected number of these requests reaching a certain contributor is $\frac{uf}{n} \times \frac{h}{hu+p}$ which is, in fact, the expected cost for a contributor. The cost for a partial contributor is unchanged: $\frac{f}{hu+p}$, since requests originating from anywhere can access a partial contributor. The number of files accessible to a contributor is *hdu+dp*, which is the total number of files advertised in the system, while the number of files available to a partial contributor is *dp* as the only files accessible to a partial contributor are those accessible by other partial contributors.

One extreme case (presented in Figure 3, center) emerges when the benefit users perceive from the files available in the system increases very slowly with the number of files. In this case introducing the incentive scheme moves the equilibrium to the right: a number of users find it is in their advantage to contribute. (Remember that without an incentive scheme, at equilibrium, users did not contribute to the system regardless of perceived benefits. In this case the equilibrium was at the leftmost point in our graphs – Figure 3, left).

The relative 'strength' of the benefit function will determine how much more to the right the equilibrium shifts. That is, how many more users are motivated to contribute. In the case where the benefit function is the logarithm of the number of files accessible (Figure 3, right), when a user is faced with a choice, at any point on the X-axis, she will choose to contribute, thus shifting the system more and more to the right and to its *efficient* equilibrium (an efficient equilibrium is defined in the Pareto [22] sense: no one can be made better off without making someone else worse off). At this point everyone shares all of their files, resulting in higher payoffs for all users.

5.2. Simulations

We have shown that an incentive scheme can significantly improve the efficiency of equilibrium state of a distributed file sharing system. However, to overcome the limitations in our analytical model (binary choices for users, lack of heterogeneity in user characteristics and inability to capture dynamic situations) we use simulations to study a more general case where a heterogeneous set of users can (incrementally and dynamically) change the number of files they share, depending on perceived benefits. Simulations also help us compare *Peer-Approved* to (the pricing like scheme) *Token-Exchange,* in this dynamic and more realistic scenario.

We assume a fixed number of users, who have limited storage and bandwidth, and an initial state in which files are placed at users according to a distribution function. Files are assumed to be equally popular. Each user initially advertises only a percentage of the files she actually has (again according to a distribution). At each iteration ***f*** users request one file each. No individual files are modeled, and requests get assigned to a peer probabilistically. Hence, a peer advertising more files will receive proportionally more requests than a peer advertising fewer files.

A request is satisfied if and only if the requesting user meets the criteria for the incentive scheme in use: i.e., if the user has a token to spend in the case of *Token-Exchange*, or a rating that is not less than the server rating in case of *Peer-Approved.* Note that in these two schemes, unlike in *Service-Quality,* the peer advertising a file, cannot block an eligible user from downloading that file. Hence, even though there are no immediate costs to advertise a file, an advertised file would attract requests that cannot be denied. Hence, advertising a file has a potential cost associated with it.

In the case of *Peer-Approved*, the rating of a user is the number of files currently advertised by that user. For the *Token-Exchange* case, each user is initially assigned a small number of tokens.

Given the above scenario, we model 'rational' user behavior in two ways. First, a user is motivated to advertise one more file when she is denied access. For example, when *Peer-Approved* is the incentive scheme in place, the user is denied a file if she has lower ratings than the server peer. Since the goal of the user is to gain access to a wide variety of files, she is motivated to increase her rating. To do so, she will advertise more files. In the case of *Token Exchange,* a user is denied a file if she ran out of tokens. The only way to gain additional files is to gain more tokens, which can only be achieved if someone downloads a file from her. Thus, the user is once again motivated to advertise more files, to attract more requests.

Second, a user reduces the number of shared files if too much of her own bandwidth is consumed by others. At each iteration users keep track of how much local bandwidth was used (how many downloads were served). If this exceeds a threshold, the user reduces the number of advertised files, which will lead to fewer download requests in the future.

Thus, depending on the perceived benefits and costs of file-sharing, at each iteration the user is motivated to either stay in the current status or to increase or decrease (by one), the number of files advertised.

Since the overall goal of an incentive scheme is to motivate users to share, or in our case to advertise more files, we measure the success of a scheme by the total number of files advertised.

We present our simulation results in Figure 4, which shows the performance (expressed in terms of the total number files shared) of *Peer-Approved*, *Peer-Approved-Tier* (a variation of Peer-Approved in which there are only a limited number of user rating categories), and *Token-Exchange*, under two different initial file-sharing distributions.

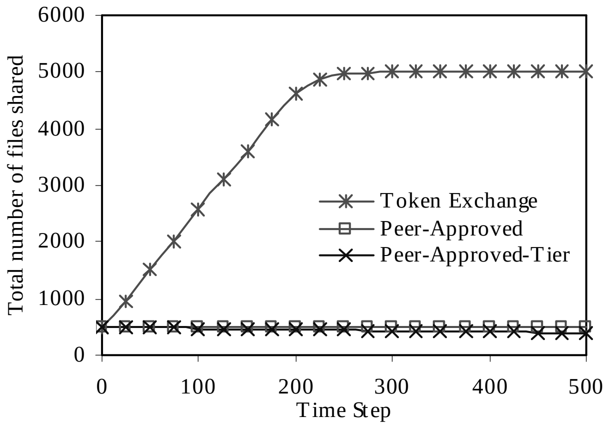

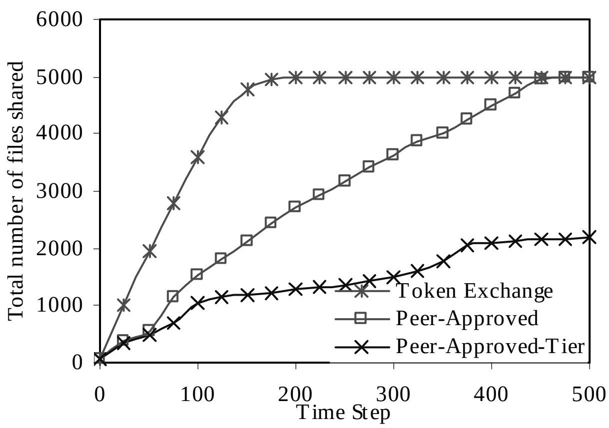

Figure 4: Simulation results for initial (top plot) Uniform file distribution and (bottom plot) Zipf file distribution.

In the uniform file distribution scenario (top) every user initially has 50 files and shares/advertises 5 files each. In the non-uniform scenario (bottom) each user has 50 files and advertises according to a Zipf distribution [23] (N= 50, α=2). In both cases, users have the same bandwidth and storage space.

We see that in the case of a uniform initial distribution of files shared, all users start with the same rating and hence can access files from all other users. Thus the rating schemes do not motivate users to advertise more files.

However, in the non-uniform case user ratings vary from the beginning, which in the case of *Peer-Approved* motivates lower-ranked users to advertise more files, albeit more slowly than in the case of *Token-Exchange*. *Peer-Approved-Tier* is even slower to converge to an equilibrium because users are distributed into a smaller number of rating slots and thus have access to more files than in *Peer-Approved*, and so are slower to advertise more of their files.

We believe that the non-uniform case is a more realistic scenario. In this case, a Grid collaboration is made up of heterogeneous users with varying resource capabilities. Thus, *Peer-Approved* could be a useful incentive scheme in such scenarios, since without involving direct payments, its performance is comparable to a pricing scheme like *Token-Exchange*. However, *Token-Exchange* converges faster in most settings considered.

6. Related Work

We discuss related work in three fields. First we describe sharing in social collaborations, and venture how studies in this traditional field can be applied to virtual communities like Grids. Second, our soft incentive schemes rely on an infrastructure for reputation management. Hence we detail related work in this arena. Finally we mention projects that aim to solve the same problem we contend with: incentivising users to contribute resources.

6.1. Resource Management in Social Groups

The resource sharing problem has been studied in great detail in various social settings. A recurring problem is the over-utilization of natural resources that are shared among a number of users, where individual rational behavior leads to the 'Tragedy of the Commons' (for example over-fishing, cattle grazing in a common pasture) [24]. In the much quoted phrase from The Logic of Collective Action [25], Olson states that "*unless the number of individuals is quite small, or unless there is coercion or some other special device to make individuals act in their common interest, rational self-interested individuals will not act to achieve their common or group interests.*"

Ostrom [26] describes mechanisms for self governance and proposes voluntary organizations rather than coercive methods to curb free-riders.

6.2. Reputation Management

Reputation mechanisms have long been studied, initially in traditional, offline settings where consumer reports tallying complaints, personal experience, and gossip have been shown to influence a consumer's decision to engage in a transaction. More recent online reputation systems are at the core of the success of many recent Internet services (e.g., eBay's auction services) [27]. These centralized mechanisms however are unsuitable for Grids that support collaborations backed by resources from multiple administrative domains. The main challenge for decentralized reputation algorithms is aggregating local reputation values while bounding the generating network traffic. Recently proposed solutions [20] that fit this requirement are based on the same random walk model as Google's PageRank algorithm [28].

6.3. Sharing in Virtual Communities

The Grid Economy project [11] propose an economic framework for resource management and scheduling of jobs. Motivated by the issue of what causes a resource owner to contribute resources to the collaboration, they propose mechanisms for participants to buy and sell individual resources.

Golle et al [12] examine the sharing problem in peer-to-peer file sharing systems They show that the free-rider problem is a real issue for peer-to-peer systems and propose different payment mechanisms to encourage file sharing. Feigenbaum and Shenker [29] examine the application of distributed algorithm mechanism design (DAMD) to distributed communities composed of autonomous entities. They identify numerous open problems to be addressed for successful deployments of such systems [26]. Fu et al. [30] propose an architecture for secure distributed resource sharing wherein resources are partitioned into small exchangeable slices. These slices can then be used for bartering or trading in a computational economy.

7. Conclusions

We have presented a model based on the Multi-Person Prisoner's Dilemma (MPD) for studying the resource contribution problem in Grid communities. Our model illuminated the paradox of “individually rational strategies leading to collectively irrational outcomes” [26] . We then used both this model and simulations to analyze the effectiveness of different incentive schemes designed to motivate increased user contributions. We compared one such scheme, the reputation-based *Peer-Approved*, with a *Token-Exchange* based scheme. Our results support the intuition that these simple incentive schemes can be used effectively to counter selfish user behavior.

We leave a number of important issues open. First, the mechanisms required to support the incentive schemes that we study can impose significant communication costs on a system. Although we perceive the costs associated with the different schemes to be relatively similar, we have not analyzed those costs in detail, nor have we investigated alternative incentive schemes that might involve lower costs. We would like to compare the communication costs of different schemes in order to quantify, ultimately, tradeoffs between cost and effectiveness. Second, our simulation study considers only two simple incentive schemes and in relatively standard settings. We plan to study additional incentive schemes, focusing on schemes that have low overhead, are easily deployable, and are acceptable by users. Third, we want to extend our analysis to other resources like storage and compute cycles that are more transient in nature.

Acknowledgements

We thank Adriana Iamnitchi for helpful discussions and feedback. We also thank the organizers of the workshop on Economics of peer-to-peer systems for the opportunity to present our initial ideas.

This work was supported in part by the National Science Foundation under contract ITR-0086044 (GriPhyN).

References

1. Foster, I. and C. Kesselman, eds. *The Grid: Blueprint for a New Computing Infrastructure*. 1999, Morgan Kaufmann.
2. Foster, I., C. Kesselman, and S. Tuecke, *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*. International Journal of Supercomputing Applications, 2001. **15**(3).
3. Oram, A., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. 2001: O'Reilly & Associates.
4. *GriPhyN: Grid Physics Network*.
5. Chervenak, A., et al., *The Data Grid: Towards an Architecture for the Distributed Management and Analysis of Large Scientific Data Sets*. Journal of Network and Computer Applications: Special Issue on Network-Based Storage Services, 2000. **23**(3): p. 187-200.
6. Rowstron, A. and P. Druschel. *Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility*. in *Proceedings of 18th SOSP'01*. 2001.
7. Ripeanu, M., I. Foster, and A. Iamnitchi, *Mapping the Gnutella Network: Properties of Large-Scale Peer-to-peer Systems and Implications for System Design*. IEEE Internet Computing Journal, Special issue on Peer-to-peer Networking, 2002. **6**(1).
8. Saroiu, S., P.K. Gummadi, and S.D. Gribble. *A Measurement Study of Peer-to-Peer File Sharing Systems*. in *Multimedia Computing and Networking Conference (MMCN), San Jose, CA*. 2002.
9. Adar, E. and B.A. Huberman, *Free Riding on Gnutella*. 2000, First Monday.
10. Leibowitz, N., M. Ripeanu, and A. Wierzbicki. *Deconstructing the Kazaa Network*. in *3rd IEEE Workshop on Internet Applications (WIAPP'03)*. 2003. Santa Clara, CA.
11. Rajkumar Buyya, D.A., Jonathan Giddy. *Economy Driven Resource Management Architecture for Global Computational Power Grids*. in *The 2000 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000)*. 2000. Las Vegas.

12. Golle, P., K.Leyton-Brown, and I.Mironov. *Incentives for Sharing in Peer-to-Peer Networks*. in *ACM Conference on Electronic Commerce*. 2002.
13. Krishnan, R., M.D. Smith, and R. Telang, *The Economics of Peer-To-Peer Networks*. 2002, Working Paper, Carnegie Mellon University.
14. Schelling, T.C., *Micromotives and Macrobehavior*. 1978: W.W.Norton & Company.
15. Hardin, R., *Collective Action*. 1982, Baltimore: John Hopkins University Press.
16. Osborne, M.J. and A. Rubinstein, *A Course in Game Theory*. 1994, Cambridge, Massachusetts: MIT Press.
17. Wilcox-O'Hearn, B. *Experiences deploying a large scale emergent network*. in *IPTPS'02*. 2002. Cambridge, MA.
18. Szabo, N. *Micropayments and mental transaction costs*. in *2nd Berlin Internet Economics Workshop*.
19. Ngan, T., D. Wallach, and P. Druschel. *Enforcing Fair Sharing of Peer-to-Peer Resources*. in *2nd International Workshop of Peer-To-Peer Systems (IPTPS'03)*. 2003. Berkeley, CA.
20. Kamvar, S.D., M.T. Schlosser, and H. Garcia-Molina. *EigenRep: Reputation Management in P2P Networks*. in *The 12th International World Wide Web Conference (WWW-12)*. 2003. Budapest, Hungary.
21. Odlyzko, A.M. *Paris Metro Pricing: The minimalist differentiated services solution*. in *7th International Workshop on Quality of Service (IWQoS '99)*. 1999.
22. Varian, H.R., *Microeconomic Analysis*. 1992: W.W.Norton & Company.
23. Adamic, L., *Zipf, Power-laws, and Pareto - a ranking tutorial, http://ginger.hpl.hp.com/shl/papers/ranking/*.
24. Hardin, G., *The Tragedy of the Commons*. Science, 1968. **162**(1243-8).
25. Olson, M., *The Logic of Collective Action. Public Goods and the Theory of Groups*. 1965, Cambridge, Mass.: Harvard University Press.
26. Ostrom, E., *Governing the Commons. The Evolution of Institutions for Collective Action*. 1990: Cambridge University Press.
27. Gross, B. and A. Acquisti. *Balances of Power on eBay: Peers or Unequals?* in *Workshop on Economics of Peer-to-peer systems*. 2003. Berkeley, CA.
28. Page, L., S.B.R. Motwani, and T. Winograd, *The PageRank Citation Ranking: Bringing Order to the Web*. 1998, Stanford Digital Library Technologies Project.
29. Feigenbaum, J. and S. Shenker. *Distributed Algorithmic Mechanism Design: Recent Results and Future Directions*. in *6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*. 2002. New York: ACM Press.
30. Fu, Y., et al. *SHARP: an architecture for secure resource peering*. in *Nineteenth ACM symposium on Operating systems principles*. 2003.