Free-Riding and Whitewashing in Peer-to-Peer Systems

Michal Feldman† Christos Papadimitriou‡ John Chuang† Ion Stoica‡

†School of Information Management and Systems, UC Berkeley
‡Computer Science Division, UC Berkeley

3rd Annual Workshop on Economics and Information Security (WEIS2004)

May 6, 2004

Abstract – We devise a simple model to study the phenomenon of *free-riding* and the effect of *free identities* on user behavior in peer-to-peer systems. At the heart of our model is a strategic user of a certain *type*, an intrinsic and private parameter that reflects the user's generosity. The user decides whether to contribute or free-ride based on how the current burden of contributing in the system compares to her type. We derive the emerging cooperation level in equilibrium and quantify the effect of providing free-riders with degraded service on the emerging cooperation. We find that this *penalty mechanism* is beneficial mostly when the "generosity level" of the society (i.e., the average type) is low. To quantify the social cost of free identities, we extend the model to account for dynamic scenarios with turnover (users joining and leaving) and with *whitewashers:* users who strategically leave the system and re-join with a new identity. We find that the imposition of penalty on all legitimate newcomers incurs a significant social loss only under high turnover rates in conjunction with intermediate societal generosity levels.

1 Introduction

Why is free-riding so widespread among users of peer-to-peer (P2P) systems? How costly is it in terms of the performance of the system? What mechanisms discourage free-riding? What is the social cost of cheap pseudonyms? What is the tradeoff between tolerating whitewashers and penalizing newcomers in the presence of cheap pseudonyms? These are the questions that motivate us.

P2P systems rely on voluntary contribution of resources from the individual participants. However, individual rationality results in free-riding behavior among peers, at the expense of collective welfare. Empirical studies have shown prevalent free-riding in P2P file sharing systems [1, 16]. While it is possible that free-riding can be sustainable in equilibrium and may even occur as part of the socially optimal outcome [12], there has been significant interest in the design of incentive mechanisms to encourage cooperation in P2P systems [4, 6, 10, 13, 17].

In many of the proposed incentive schemes, rewards and/or punishments are handed out to peers according to their contribution level. However, imposing penalties on free-riders require some means of identifying free-riders and distinguishing them from contributors. Reputation systems [11, 14] may help, but these systems may be vulnerable to the whitewashing attack, where a free-rider repeatedly re-joins the network under new identities to avoid the penalty imposed on free-riders. The whitewashing attack is made feasible by the availability of low cost identities or *cheap pseudonyms* [9]. There are two ways to counter whitewashing attacks. The first is to require the use of free but irreplaceable pseudonyms, e.g., through the assignment of strong identities by a central trusted authority [5]. In the absence of such mechanisms, it may be necessary to impose a penalty on all newcomers, including both legitimate newcomers and whitewashers. This results in a social cost due to cheap pseudonyms, as demonstrated by Friedman and Resnick [9].

We develop a simple modeling framework that helps to predict the level of free-riding in P2P systems. We use this model to quantify the effect of a *penalty mechanism*, which gives free-riders degraded service, on the emerging cooperation level. We find that the penalty mechanism is beneficial mostly when the generosity level is low, and is effective in discouraging free-riding behavior when the threat of penalty is sufficiently high relative to the cost of contribution. To study the tradeoff between tolerating whitewashers and penalizing newcomers in the presence of cheap pseudonyms, we extend the model to dynamic scenarios with turnover (users joining and leaving). We find that a significant social cost (in the form of system performance loss) is incurred only under high turnover rates and intermediate generosity levels.

2 Model

At the heart of our model is a user as a rational agent with a private and intrinsic characteristic called her *type*, a single parameter reflecting the willingness of the agent to contribute (the agent's type can be intuitively thought as a qualitative measure of her decency or generosity).

Each user decides whether to contribute or free-ride based on how the current *burden* of contributing in the system compares to her type. We assume that the cost of contributing is the inverse of the total percentage of contributors, i.e., $R = 1/x$, because when many people free-ride, the load on the contributors increases. Thus, if at present a fraction x of the users contribute, a user with type t_i will contribute if $1/x < t_i$, and free-ride otherwise.

Even within this minimalistic framework we can already see some interesting implications. In this "free market" environment, the percentage x of contributors is determined as the intersection of the type distribution with the curve $x = 1/t$ (see Figure 1). Of the two intersection points, the higher one is the attractor of the natural fixpoint dynamics, i.e., starting at some initial x, the agents arrive at their individual decisions, their aggregate decisions define a new x, and so on. As long as the initial x is above the lower intersection point, the process will converge to the upper one. If there is no intersection, i.e., when there are too many selfish rascals around, then x becomes 0 (the other attractor, which always exists) and the system collapses.

So far we have been interested only in costs. What is a user's benefit when the level of contribution is x? We assume that the benefit a user receives from participation in the system (whether or not she contributes), denoted by Q, is a function of the form $Q = \alpha x^{\beta}$, where $\beta \leq 1$ and $\alpha > 0$ are positive constants. Hence user benefit is an increasing function of the number of contributors, but with diminishing returns—a form widely accepted in this context (see, e.g., [2], [3], [15]). Thus, the performance of the system, denoted by W_{system}, is defined as the difference between the average benefit received by all users (including both contributors and free-riders) and the average burden experienced by all users, which effectively include only the contributors (as free-riders incur no burden):

$$W_{system} = \alpha x^{\beta} - 1 \tag{1}$$

3 Contribution Level

A user's type is a random variable with unknown distribution. For simplicity we take this distribution to be the uniform distribution between 0 and t_m. t_m is thus an important parameter of the system, as it reflects the society's "generosity" (it is twice the expected type).

To derive the contribution level x (the percentage of users who contribute) we solve this fixpoint equation:

$$x = \text{Prob}(t_i \geq 1/x) = 1 - \frac{1}{x t_m}$$

which yields: $x_{1,2} = \frac{t_m \pm \sqrt{t_m^2 - 4t_m}}{2t_m}$. The larger root x_1 is the stable equilibrium (attractor, see Figure 1) while x_2 is unstable. For $t_m < 4$ there is no intersection between the curves, thus the contribution level becomes 0 and the system collapses.

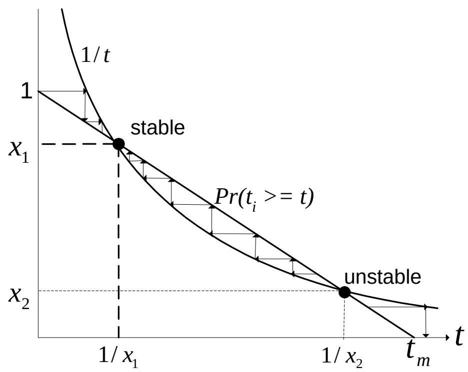

Figure 1: The intersection points of the two curves represent the two equilibria of the system. The curve $1/t$ represents the contribution cost, and $Pr(t_i >= t)$ represents the generosity CDF, where $t_i \sim U(0, t_m)$. The higher equilibrium (contribution level x_1) is stable.

General symbols	
t_i	user $i's$ type
t_m	maximal type in population
α	system benefit coefficient
β	diminishing returns coefficient
W	realized performance
Q	individual benefit
R	burden (individual contribution cost)
T	threat level
z	exclusion fraction
p	penalty level
p_m	maximal possible penalty
Static system	
x	contribution level
Dynamic system	
x_s	contribution level of stayers
x_l	contribution level of leavers
x_a	average contribution level
d	turnover rate

Table 1: This table presents the symbol notations in our model.

4 Penalty Mechanism

With the free market as baseline, we can consider different forms of intervention to discourage free-riding behavior. Assuming that user behavior is observable even when user type is not, we can institute a penalty mechanism where the users that are identified as free-riders are subject to a penalty. Since we are primarily interested in the incentive effect of the mechanism rather than its implementation, we will consider, at an abstract level, a simple penalty *p* that can be applied to the free-riders. We can interpret the value *p* as the probability that a free-rider can be caught and excluded from the system. Alternatively, we can adopt a service differentiation interpretation [4, 11, 8], where the free-riders receive a reduced benefit of (1 - *p*)*Q*. Downgrading the benefit of the free-riders increases user contribution in two ways. First, the reduction in system load reduces the burden *R* imposed on the contributors. Second, it introduces a *threat*, denoted by *T*; users know that they will received reduced service if they decide to free-ride.

Under the penalty mechanism, the realized performance of contributors and free-riders is:

$$W_{contributors} = Q - R = \alpha x^{\beta} - \frac{x+(1-x)(1-p)}{x}$$

$$W_{free-riders} = Q - T = \alpha x^{\beta} - p\alpha x^{\beta}$$

Consequently, the contribution level, x, is derived according to the following expression:

$$x = \text{Prob}(t_i \geq R - T)$$

$$x = \text{Prob}(t_i \geq \frac{x + (1-x)(1-p)}{x} - p\alpha x^{\beta}) \qquad (2)$$

In what follows, we set $\beta = 1$ for tractability and presentation clarity. If the threat is high enough such that $R - T < 0$, we get $x = 1$. If $R - T > t_m$, no contribution emerges, and for intermediate values, $0 \leq R - T \leq t_m$, the contribution level in equilibrium is:

$$x = \frac{p - t_m + \sqrt{p^2 + 2t_m p + t_m^2 - 4t_m + 4p\alpha - 4p^2\alpha}}{2(-t_m + p\alpha)}$$

System performance now becomes:

$$W_{system}^{penalty} = (\alpha x^{\beta} - 1)(x + (1-x)(1-p))$$

and the optimal penalty level is: $p^* = \text{argmax}_p \quad W_{system}^{penalty}$

While p yields a social benefit due to the higher contribution level it achieves, it also incurs a social cost in the form of reduced benefit to free-riders. However, we identify the interesting case in which p is set high enough such that it achieves full cooperation ($x = 1$). If so, the penalty mechanism introduces no social cost since it only serves to threaten users but no penalty is effectively imposed. In this case, we achieve the maximal benefit (denoted by Q_m) and the maximal system performance as a result.

$$Q_m = Q(x = 1) = \alpha$$

Based on equation 2, if $p = \frac{1}{\alpha}$, x becomes 1, and the system performance cannot be further improved. For example, if $Q_m = 10$, we only need a mechanism that can catch and exclude a free-rider with 10% probability. Observe that the threshold value equals $\frac{1}{Q_m}$, thus decreases in Q_m.

The above results seem very optimistic; if we impose a high enough penalty, or are able to identify and exclude free-riders with high probability, we achieve optimal system performance. However, these conclusions should be viewed with caution since they assume that the system designer has the freedom to set p as high as desired. In reality, it may be difficult or costly to exclude free-riders with very high probability, so in many cases p will be restricted by a maximal feasible value, denoted by p_m.

Figure 2 presents the percentile of optimal performance that can be achieved by the penalty mechanism for different p_m and α values. We observe that optimal system performance can be achieved, regardless of the value of t_m, as long as the penalty is set to be at least $1/\alpha$. This is true for both the $p = .91$, $\alpha = 1.1$ and $p = 0.1$, $\alpha = 10$ curves. On the other hand, if the penalty is set too low (e.g., $p = 0.1$ when $\alpha = 1.1$), then the outcome is not significantly better than the free-market ($p = 0$) outcome.

An additional issue is the stability of the equilibrium. It is true that if $p = 1/\alpha$, $x = 1$ is an equilibrium for all t_m values. However, the basin of attraction, $[1 - \epsilon, 1]$, varies depending on t_m, and a low t_m value may lead to an extremely small ϵ or even $\epsilon = 0$, which means that the system will never converge to $x = 1$ unless the initial x is 1. The threshold value above which the system converges to $x = 1$, denoted by $t_m^{threshold}$, is a function of α and ϵ as follows:

$$t_m^{threshold} = \frac{1 - 2\alpha + \epsilon\alpha}{\alpha(\epsilon - 1)}$$

$t_m^{threshold}$ increases in ϵ as expected, and also increases in α. For example, if $\alpha = 10$, $t_m^{threshold}(\epsilon = 0) = 1.9$, which means that for $t_m < 1.9$, $x = 1$ is not an attractor, and the wider the desired basin of attraction is, the higher $t_m^{threshold}$ becomes.

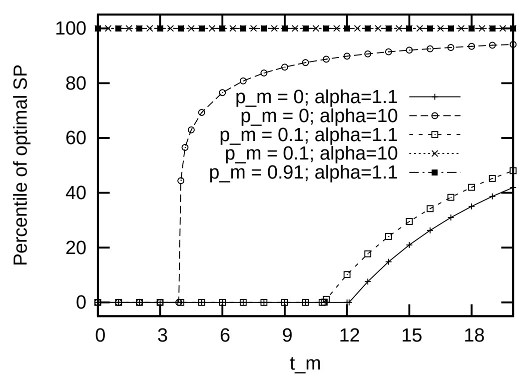

Figure 2: Percentile of optimal performance for different values of α and p_m. Note that $p = 0$ corresponds to the free-market scenario.

5 The Social Cost of Free Identities

In Section 4, we show that a penalty mechanism can be effective in discouraging free-riding behavior. However, the effectiveness of penalties can be undermined by the availability of cheap pseudonyms. In particular, a free-rider might choose to *whitewash*, i.e., leave and re-join the network with a new identity on a repeated basis, to avoid the penalty imposed on a free-rider. The lower the cost of acquiring new identities, the more likely a free-rider will engage in whitewashing. Since whitewashers are indistinguishable from legitimate newcomers, it is not possible to single them out for the imposition of a penalty. Of course, it is possible to counter the whitewashing

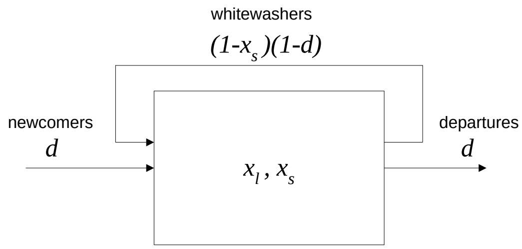

Figure 3: Dynamic system with arrivals, departures, and whitewashers. A fraction d of users depart and are replaced by the same number of newcomers. At the same time, a fraction $(1 - d)(1 - x_s)$ of users whitewash under FI.

strategy by imposing the penalty on all newcomers. However, this results in a social cost, as shown by Friedman and Resnick [9].

In this section, we are interested in quantifying the social cost of cheap pseudonyms in terms of system performance loss. We do so by extending our model from section 4 into a dynamic model with user joins and leaves. To quantify the social loss due to cheap pseudonyms we consider two dynamic scenaria, namely, *permanent identities* (PI) and *free identities* (FI).

Under PI, identity costs are taken to be infinity, while under FI, they are costless [1]. Therefore, these two cases represent two extremes. In actuality, identity cost can take any positive finite value, and users decide whether to whitewash or not depending on how the identity cost compares to the penalty imposed on free-riders and newcomers. In this paper, we focus on the two extreme cases, as we believe they provide important insights while still preserving some level of simplicity.

5.1 System Dynamics and Population Mixture

We model a system with joins and leaves, with a turnover rate of d (Figure 3). We assume that arrivals and departures are type-neutral and therefore do not alter the type distribution [2].

The population at each point in time is composed of the following four groups:

- existing contributors (EC)
- existing free-riders / whitewashers (EF/WW)
- new contributors (NC)
- new free-riders (NF)

The difference between the permanent and free identities scenaria is signified by the members of the second group. While free-riders stay in the system if identities are permanent, they will adopt whitewashing behavior under free identities. However, if penalty is imposed also on newcomers, free-riders are indifferent between staying or whitewashing.

5.2 Burden, Threat and Contribution Levels

An important property of the dynamic scenaria is that not all users care about the threat. The users who leave the system at the end of each period are not affected by the penalty they would have paid had they stayed in the system. Consequently, we get two separate contribution levels:

[1]Identity cost refers to the cost of acquiring any additional identity after the first one, which is considered to be a sunk cost.
[2]The model can be extended in future work by considering more sophisticated dynamics, as discussed in Section 6.

	NC not penalized	NC penalized
% penalized	$(1-d)(1-x)$	$d+(1-d)(1-x)$
% not penalized	$(1-d)x+d$	$(1-d)x$

Table 2: The respective fraction of users who get full and reduced benefit with and without penalties to newcomers.

x_l: the contribution level of the "leavers"

x_s: the contribution level of the "stayers"

The values of x_s and x_l in equilibrium satisfy the following equations:

$$x_l = \text{Prob}(t_i \geq R) \tag{3}$$

$$x_s = \text{Prob}(t_i \geq R - T) \tag{4}$$

The average contribution level in the system, denoted by x_a, is:

$$x_a = dx_l + (1-d)x_s$$

The contribution level of the stayers is always greater than or equal to that of the leavers. Unlike the static system, where $x = 1$ can be achieved for a sufficiently high p, in dynamic scenaria we cannot achieve $x_a = 1$ due to the leavers.

The individual contribution cost (burden) in each period is determined by the ratio between the fraction of users who get the full benefit and those who get the reduced benefit. If only the existing free-riders are penalized (feasible only under PI), all groups except for the EF get the full service. However, if all newcomers are penalized, all groups except for the EC get the reduced service. Table 5.2 presents the respective fraction of users who get full and reduced benefit under the two scenaria in steady state:

Based on this table, the burden under PI, when newcomers are not penalized is:

$$R_{PI} = \frac{(1-d)x + d + (1-d)(1-x)(1-p)}{x}$$

and the burden under FI, when newcomers are penalized is:

$$R_{FI} = \frac{(1-d)x + d(1-p) + (1-d)(1-x)(1-p)}{x}$$

Observe that the burden is lower under FI because a larger fraction of users are penalized, therefore the demand that is placed on the system is lower. Nevertheless, the benefits of all users, except for the EC, is also reduced. Observe that under PI, if we set p sufficiently high, we can get into the scenario where p is used simply as a method to threaten users but no penalty is effectively imposed (similar to the static system, see Section 4). In contrast, under FI, imposing a penalty always results in a social loss because newcomers are effectively penalized, independent of their behavior.

5.3 System Performance

The fraction in the population and realized performance level of each group under the two scenaria are presented in Table5.3. System performance is $W_{system} = \sum_j (f_j * W_j)$, and the best strategy is to impose a penalty p^* that satisfies:

Group (j)	Group Size (f_j)	Realized Performance (W_j)	
		Permanent identities	Free identities
EC	$(1-d)x$	$Q - R_{PI}$	$Q - R_{FI}$
EF / WW	$(1-d)(1-x)$	$Q(1-p)$	$Q(1-p)$
NC	dx	$Q - R_{PI}$	$Q(1-p) - R_{FI}$
NF	$d(1-x)$	Q	$Q(1-p)$

Table 3: The size and realized performance level of the different groups under the PI and the FI scenaria.

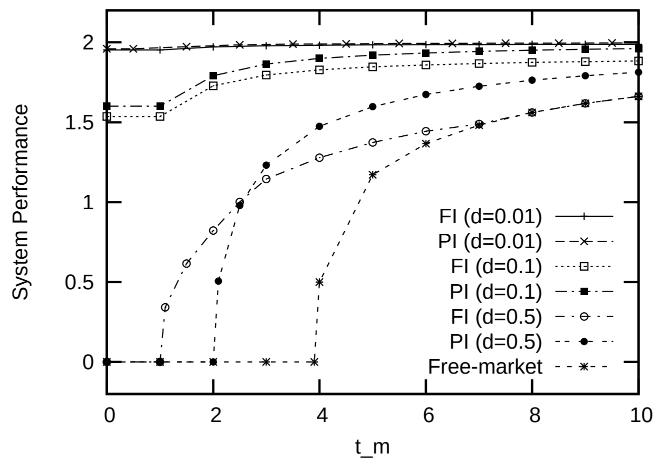

Figure 4: System performance subject to $p = p^*$ under free-identities, permanent identities and free market; $\beta = 1$, $\alpha = 3$.

$$p^* = \text{argmax}_p \quad W_{system}$$

Figure 4 compares the system performance, W_{system}, subject to a penalty $p = p^*$, under PI and FI as a function of t_m for different turnover rates (d). We make the following observations:

- For very small turnover rates ($d = 0.01$), the system performs close to its optimal level since the threat is imposed on the majority of the population, thus a small penalty level is suffi cient to achieve a high cooperation level. If so, no notable performance gap exists between PI and FI. As the turnover increases, a higher penalty is required, which incurs a social cost and reduces the system performance.
- As t_m increases, system performance converges to its optimal performance level under both scenaria. Thus, the performance gap between the two scenaria shrinks.
- Under a high turnover rate ($d = 0.5$) and low t_m values, the system performs better if penalty is imposed on newcomers even under PI. In cases where the generosity level is low and the turnover rate is high, it is hard to obtain satisfactory cooperation levels, and penalties to newcomers, although incurring a social loss, may help improve the cooperation level by reducing the load placed on the system.

We conclude that a notable social cost due to free identities is incurred only under specifi c conditions, in which a penalty on all newcomers is unnecessarily imposed. In particular, it is incurred only under high turnover rates (d) and only in conjunction with intermediate generosity levels (t_m) and low system benefi ts (α). In contrast, in cases where the system can tolerate the newcomers, the imposition of penalty on all newcomers incurs a social loss. In what follows, we provide some observations that help explain these fi ndings:

- If the turnover rate is low, the fraction of newcomers in the population is small. Therefore, penalizing newcomers does not have a big impact on the system performance. In addition, because the population is fairly permanent, a low p imposes a sufficient threat to obtain high cooperation.

- If the turnover rate is high and the societal generosity level is low, system collapse can only be avoided by reducing the demand placed on the system. Assessing a penalty on all newcomers is one method to limit the demand. Therefore, in these situations, not only does penalizing newcomers not incur a net social loss, but it helps to sustain the system by reducing the load enough to avoid system overloading.

- If the societal generosity level is high, a high cooperation level is obtained even in the absence of intervention. Therefore, the best policy under both scenaria is to impose an extremely small penalty or no penalty at all. Hence, no notable social loss is incurred due to free identities.

- If the benefits of the system (α) are high, even a small p results in a high threat that is imposed on free-riders. Once again, the optimal p will be very small, thus no notable gap will occur.

6 Discussion and Future Work

We have presented an economic model of user behavior in peer-to-peer systems, and derived some useful observations. In particular, a mechanism that penalizes free-riders can improve system performance by reducing the burden placed on the contributors. This mechanism is especially effective when the societal generosity level is low, in which case the system exhibits low or zero performance level in the absence of intervention. Additionally, penalizing all newcomers may be effective in discouraging whitewashing behavior and will incur a social cost (in the form of reduced system performance) only for high levels of turnover rates and in conjunction with low system benefits and intermediate societal generosity levels.

Our model is flexible enough to account for a diverse set of characteristics. For example, we extend our model to account for heterogeneity with respect to resources. To do so, we split each *user* to a number of *virtual users* that is proportional to the amount of resources he has. We find that users with many resources bear a higher burden, therefore exhibit lower contribution levels. Because contribution from high-resources users is more valuable in terms of system performance, a heterogeneous system results in a lower system performance than a homogeneous system. However, if the amount of resources and the generosity level of users correlate, a heterogeneous system may result in better performance than a homogeneous one. Several research questions arise in this context as discussed below.

Unlike many works in this area, we are not proposing a new protocol or incentive scheme, neither do we specify any implementation details. Instead, the objective of this work is to develop a game theoretic framework that helps to gain insights into the effect of incentives schemes on user behavior and system performance, and to obtain a better understanding of the impact of the different factors and system parameters on the need and effectiveness of these schemes. For this purpose, we have simplified the model with a set of somewhat restrictive assumptions. In future work, we plan to relax or modify some of the assumptions and possibly extend the model in several directions:

- Additional incentive schemes. We plan to analyze the effect of system partitioning on user behavior and system performance. In particular, if the system is partitioned into two or more sub-systems that impose different penalties on their free-riders, how would it effect the results?

- Additional penalty forms. Consider other forms of penalty to newcomers. One candidate is entry fee that can be used as a pure transfer to the system. Some examples are monetary payments which can be distributed among the participants or entry fees in the form of contribution of resources. These mechanisms entail no direct loss in efficiency, but introduce a different set of issues. First, this type of mechanism essentially *forces* contribution at the entering stage, and may therefore prevent some users from participating. Second, contribution of resources from newcomers prior to their participation may be limited because in many cases the resources are gathered while being members. Third, redistribution of monetary payments may be difficult due to the highly dynamic membership in these systems.

- Type distribution. It would be interesting to derive general results for various user type distributions. One particular distribution to consider would be a bimodal distribution which is uniform between 0 and t_m, but has two spikes at the extremes.

- System dynamics (Section 5). In the dynamic scenario, the model can be extended by assuming (1) departure rates that depend on performance, (2) arrival rates affected by p, and (3) dynamics that affect the distribution by postulating type-dependent departures and arrivals. In particular, one could imagine that p would affect arrival rate in different directions. On the one hand, imposing penalty on newcomers may discourage them from joining the system, which will reduce arrivals. On the other hand, users who join a system that penalizes its users may have expectations for higher performance levels. Hence, it may attract users even more. Depending on the effect of p on arrival rate, the social loss due to free identities may increase or decrease relative to our results.

- Identity costs. In our analysis we consider the two extreme cases of infinite and zero identity costs. In future work, we intend to study cases in which the cost of identity, c, is a positive finite value. In this context, imposing different penalties on free-riders and newcomers may be beneficial (e.g., $p_{newcomers} = p_{free-riders} - c$).

- Additional performance metrics. In this paper we use the metric of *system performance*. This metric assigns equal weights to the realized performance of all users, whether they are contributors or free-riders. In the future, we plan to consider other performance metrics that might be more appropriate on grounds of fairness. One natural metric is one that assigns more weight to the performance of contributors than the performance experienced by free-riders.

- Resource heterogeneity. Several interesting directions can be examined in the context of resource heterogeneity. First, how do the results change if users can elect to contribute some amount of resources as opposed to the binary decision assumed here. Second, it will be interesting to experiment with alternative contribution cost functions that may be more reflective of the opportunity cost. For example, users with high resources may experience lower opportunity costs even when they contribute more resources [7]. If so, the system may exhibit better performance.

References

[1] ADAR, E., AND HUBERMAN, B. A. Free Riding on Gnutella. *First Monday 5*, 10 (October 2000).

[2] ASVANUND, A., CLAY, K., KRISHNAN, R., AND SMITH, M. An Empirical Analysis of Network Externalities in Peer-to-Peer Music Sharing Networks. In *Proceedings of the 23rd International Conference on Information Systems (ICIS)* (December 2003).

[3] BRESLAU, L., CAO, P., FAN, L., PHILLIPS, G., AND SHENKER, S. Web-Caching and Zipf-like Distributions: Evidence and Implications. In *Proceedings of Infocom* (1999).

[4] BURAGOHAIN, C., AGRAWAL, D., AND SURI, S. A Game-Theoretic Framework for Incentives in P2P Systems. In *International Conference on Peer-to-Peer Computing* (Sep 2003).

[5] CASTRO, M., DRUSCHEL, P., GANESH, A., ROWSTRON, A., AND WALLACH, D. S. Security for Structured Peer-to-Peer Overlay Networks. In *Proceedings of Multimedia Computing and Networking 2002 (MMCN '02)* (2002).

[6] CROWCROFT, J., GIBBENS, R., KELLY, F., AND ˘ OSTRING, S. Modeling Incentives for Collaboration in Mobile Ad Hoc Networks. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks* (2003).

[7] FELDMAN, M., LAI, K., CHUANG, J., AND STOICA, I. Quantifying Disincentives in Peer-to-Peer Networks. In *1st Workshop on Economics of Peer-to-Peer Systems* (2003).

[8] FELDMAN, M., LAI, K., STOICA, I., AND CHUANG, J. Robust Incentive Techniques for Peer-to-Peer Networks. In *ACM Conference on Electronic Commerce (EC'04)* (May 2004).

[9] FRIEDMAN, E., AND RESNICK, P. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy 10*, 2 (1998), 173–199.

[10] GOLLE, P., LEYTON-BROWN, K., MIRONOV, I., AND LILLIBRIDGE, M. Incentives For Sharing in Peer-to-Peer Networks. In *Proceedings of the 3rd ACM conference on Electronic Commerce, October 2001* (2001).

[11] KAMVAR, S. D., SCHLOSSER, M. T., AND GARCIA-MOLINA, H. The EigenTrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the Twelfth International World Wide Web Conference* (May 2003).

[12] KRISHNAN, R., SMITH, M., TANG, Z., AND TELANG, R. The Virtual Commons: why Free-Riding can be Tolerated in Peer-to-Peer Networks. In *Worshop on Information Systems and Economics* (December 2003).

[13] MICHIARDI, P., AND MOLVA, R. A Game Theoretical Approach to Evaluate Cooperation Enforcement Mechanisms in Mobile Ad Hoc Networks. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks* (2003).

[14] NOWAK, M. A., AND SIGMUND, K. Evolution of Indirect Reciprocity by Image Scoring. *Nature 393* (1998), 573–577.

[15] SAROIU, S., GUMMADI, K. P., DUNN, R. J., GRIBBLE, S. D., AND LEVY, H. M. An Analysis of Internet Content Delivery Systems. In *Proceedings of Symposium on Operating Systems Design and Implementation* (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 (MMCN '02)* (2002).

[17] URPI, A., BONUCCELLI, M., AND GIORDANO, S. Modeling Cooperation in Mobile Ad Hoc Networks: a Formal Description of Selfishness. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks* (2003).