

When the experiment is over: Deploying an incentive system to all the users

Rosta Farzan
School of Information Science
University of Pittsburgh
210 S. Bouquet St.
Pittsburgh, PA 15260
rosta@cs.pitt.edu

**Joan M. DiMicco, David R. Millen, Beth Brownholtz,
Werner Geyer, Casey Dugan**
IBM Research
1 Rogers Street, Cambridge, MA 02142
{joan.dimicco; david_r_millen; beth_brownholtz;
werner.geyer; cadugan} @us.ibm.com

Abstract. Motivating user participation is an important issue for the survival of social web and social software applications. In our previous work demonstrated that a point-based incentive encourages contribution to a social networking site. This paper presents a follow-up analysis after a full deployment of the incentive mechanism to the entire user community. We address an issue uncovered in our previous experiment, we measure the long-term impact of the incentive mechanism on site content generation, and we replicate our previous result with a larger number of users. Our results will demonstrate that the incentive mechanism had a long-term effect on contribution levels and generated a second boost in contribution levels when released to a new set of users. The paper concludes with a discussion of our community's reactions to the incentive mechanism collected through the site itself, company-internal blogs, podcasts, and forums.

1 INTRODUCTION

Growth of the social web and of the popularity of social software means that, increasingly, the success of websites and software applications is dependent on user contributions. This raises the important issue of how to persuade users to participate. Prior work, such as our own [6], found through controlled experiments that different incentive mechanisms work on varied conditions. In this paper, we present a follow-up to our previous experiment that presents findings on how a full deployment of an incentive system impacts an on-line community involving thousands of users within the workplace setting.

Our previous research found that, when a points-based system was put into place for half of the users of our social networking web site, those users that received points increased their contributions to the site. There was no corresponding increase in contributions for the users who did not receive points.

After this experiment, there were several remaining questions. First, in our original analysis we observed that many users did not navigate to their personal points page, indicating they did not notice the presence of system, perhaps because there was no explicit announcement of the new site feature. In this research we are looking at the effect of personalized email notification in addition to the points system.

Second, while we found that introducing a points incentive to our original community increased the contributions of the group, we wanted to know if we could replicate that jump in contributions several months later, when the site had several hundred more users, by releasing to the points to the other half of the site's users. We also wanted to know if just before releasing the points we could see a substantial difference between the group who had points for several months and the group that was about to have points shown to them.

Our third and final research question was to ask how is a points-based incentive system received by a community of over 4000 users, as the site measured at the end of our analysis. After all of our site's users could see the points system, we observed that members of the community responded explicitly by gaming the system, discussing and sharing their thoughts on the system through the communication mechanisms in the site, and commenting on it outside of the system, within company-internal blogs, podcasts and forums. The enthusiastic, and sometimes heated, debates about incentive systems highlight some of the pros and cons of the system deployed. The response also highlights how a community of employees working in the software business respond to such systems. Others' work on incentive mechanisms has been deployed to educational communities and onto the open Internet – our study of a large community of employees and their reactions to this system within the workplace is a unique perspective that can inform the design of other planned deployments of incentive systems.

2 BACKGROUND

The appropriate way to motivate user participation always depends on the task, the application, and the users' characteristics. Researchers have looked at different incentive systems addressing different type of tasks and users. Their techniques can be classified into the following approaches:

By rewards: awarding users for their contribution [1,3,4].

By explaining community benefit: highlighting the importance of users' contributions for the community [2,7,8].

By goal-setting: setting a challenging and short-term goal for the users [2].

By reputation: enhancing users' reputation in the community through their contributions [8].

By providing self-benefit: turning users' participation into an important activity for themselves [5].

In our previous work we deployed a point-based incentive system on a social networking site inside of an enterprise, to combine ideas from incenting by rewards and reputation.

The site we deployed the incentive mechanism to was a social networking site designed for IBM employees to network with each other. On the site, users create profile pages, share photos and lists, and comment on each others' content. Persuading users to contribute content to the site is an important piece of our effort to build and sustain a lively community.

Our incentive system rewarded points to the site's users for contributing content to in the form of profile information, photos, lists, and comments on any of the site's content. The number of points for each user was calculated based on the following formula.

$$p = 100 \times (\text{if any item in profile}) \\ + 15 \times \text{number of comments} \\ + 10 \times \text{number of lists} \\ + 5 \times \text{number of photos}$$

To extend the points incentive from reward to also reputation, we defined four different status levels based on the number of points as shown in Table 1. (The social networking site is called Beehive, hence the bee status names.)

Table 1. Points-based status levels

Number of points	Status level
< 110	New bee
< 500	Worker bee
< 2000	Busy bee
>= 2000	Super bee

To make users aware of their points and status level, and to enable them to compare themselves to others, we implemented the points system in the following ways on the web site:

- The top 10 point-earners were listed on the home page.
- Each user's points and status level appeared next to their name throughout the site.
- There was a page listing the users in each class ordered by number of points.
- And there was a personalized page for users to compare themselves with other people in their network by number of points and status.

We studied the effect of this point-based incentive system in a controlled study, assigning half of the population to an experimental group which had access to points system and the other half to a control group which didn't see any information about points. The study found that the points system motivated the experimental users to contribute all types of content; however the effect was not sustainable and the higher levels of contribution began to drop after one week [6].

At the conclusion of this study, we had three main questions. The first question was what if the users were told more explicitly about the points system? Would a personalized email notification about points make even more users contribute to the site? The second question was can we replicate these results with a different and larger population of users or was this a special group of enthusiastic early adopters? The third question was more general: what will happen when the whole site has the

points system visible? Will there be gaming of the system? (We did not observe gaming in our original experiment.) Will users refer to points in their contributions? The next three sections of this paper attempt to answer these questions.

3 EFFECT OF EMAIL NOTIFICATION

Our previous study found that most of the users in our experimental condition (72%) did not visit any of the points pages during our experiment, indicating that they probably did not notice the addition of points system. If users are unaware of an incentive system, the persuasive impact of the system is quite limited. To address this problem, we designed personal notification emails to be sent users about the points system. The site already sent email updates to users about their social network's activity either daily or weekly, depending on users' preferences. To increase the likelihood that users would discover the points system, we added the following information to the top of these personalized email notifications:

You are a **new bee** with **25** points:

You only need **85** more points to become a worker bee.

- Each time you share a photo, you earn 5 points.
- Each time you create a hive5, you earn 10 points.
- Each time you comment on a profile, a photo, or a hive5, you earn 15 points.
- The first time you put content into the "about you" section of your profile page, you earn 100 points.

Click [here](#) to find out the details about the beehive points and compare yourself with your network.

The highlighted portions of the above text were personalized for each user, showing status, total number of points, and number of points needed to jump to next class. The message about adding information into the profile was included if the user had no profile information.

To evaluate the effect of this addition to the email notifications, we looked for changes in contribution levels within the experimental group (41 users), as compared with the control group who did not see any changes to the site or their emails (40 users). Looking at users' total point values, in the week following the email notification, 12% of users in the experimental group moved out of zero points status and 7% of them jumped from New bee to Worker bee status, while only 5% of the control group moved out of zero-point status and no one in the control group earned enough points to jump from New bee to Worker bee.

Figure 1 compares the control versus experimental groups in terms of the average number of photos, lists, and comments added the week before and the week after receiving the first enhanced email notification. As the data shows in Figure 1, within the experimental group, there was a 67% increase in the number of photos added, a 75% increase in lists, and a 92% increase in comments. Figure 2 shows the percentage of users in each group who added content to their profile section for the first time, the week before and the week after the email was sent.

These results suggest that the email including notification about the points caused an increase in the amount of content added by the experimental group. However, the difference is not statistically significant for each content type separately or for the overall amount of points earned by each group (We ran a repeated-measures ANOVA, considering time as the repeated

measure, group as the independent variable and total amount of new content added to the site as the dependent variable).

There is further evidence to support that the email encouraged users to contribute content found in the percentage of users contributing to the site each week. We found that the percentage of users contributing photos and lists nearly doubled in the week following the email notification (7% to 12% of users for photos, 5% to 10% of users for lists); the percentage of users adding comments dropped by 2% from 24% to 22%.

Our conclusion from this analysis is that notifying users by email about the points was an effective way to extend the awareness of the points system and in turn generate more user content.

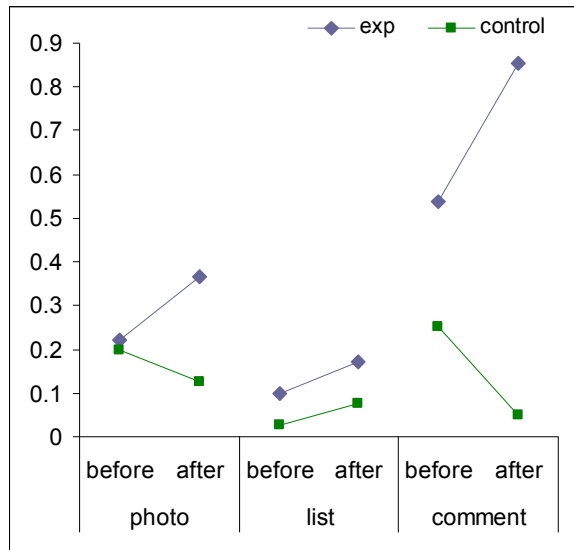


Figure 1. The average number of photos, lists, and comments added by each group the week before and the week after the email notification

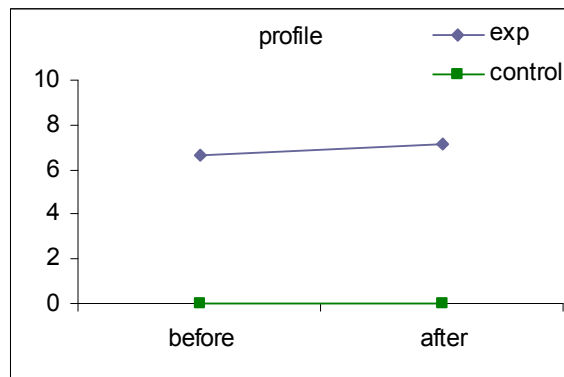


Figure 2. Percentage of users in each group who added content to their profile “about me” section for the first time

4 RELEASING POINTS SYSTEM TO ALL USERS

At the conclusion of our experiment, we left the points system visible to just half of the users for another three months. At that point, after the site’s community had more than doubled in size, we determined it was time to let all users see the points system, particularly since the points and the email had generated boosts in site contributions.

By releasing the points to the entire site, we can also assess two aspects of the points system. First, we can assess the long-term effect of one group having points and another not. Second, we can see if our original findings, that points generated a boost in contributions, are repeatable, with a new, larger community, when we release the points to the control group.

In our previous experiment [6], the experimental group had 63 users, as did the control group. Now the group that had the points had grown to 207 users and the group that would be newly introduced the points had 214 users. For clarity, we will refer to these two groups in our analysis as **PtsGroup1**, the group of users who had access to the points system for several months and in most cases since the day they joined the site, and **PtsGroup2**, the group who the points were released to all on a single day.

Our two evaluations were done through log analysis considering six weeks of usage logs – the three weeks before and the three weeks after releasing points system to PtsGroup2. For consistency we limited the users of the study to those who used the system consistently over the six-week period, which means they used the system at least once every week.

4.1 Differences between PtsGroup1 and PtsGroup2

Before releasing the points, we assessed the long-term impact of the points by looking for differences between PtsGroup1 and PtsGroup2. Table 2 presents the amount of content added by the users in each of these groups, both before releasing the points and three weeks after the points were released. From the numbers, we can see that PtsGroup1 added much more content to the site than PtsGroup2, especially in the case of comments, where users added three times the number of comments to the site. Fitting negative binomial regression model shows a significant effect of group for lists and comments (lists: $df=418, 1, \chi^2=6.38, p=0.01$, comments: $df=418, 1, \chi^2=25.14, p<0.0001$). These data provide strong evidence that the points system generates more content on the site over the long term.

Table 2. Differences in contribution amounts between PtsGroup1 (207 users) and PtsGroup2 (214 users)

	Total Photos	Total Lists	Total Comments
Before releasing points:			
PtsGroup1	823 (48% more)	443 (74% more)	2703 (299% more)
PtsGroup2	556 (32% fewer)	254 (43% fewer)	678 (75% fewer)
Three weeks after:			
PtsGroup1	991 (37% more)	484 (49% more)	3630 (234% more)
PtsGroup2	722 (27% fewer)	324 (33% fewer)	1086 (70% fewer)

Table 3. The percentage of users in each group that had contributed content before the release of the points to PtsGroup2

	Photos	Lists	Comments	Profile "About Me"
PtsGroup1	44.4%	43.0%	42.1%	43.9%
PtsGroup2	45.6%	38.3%	43.7%	36.4%

Table 3 presents the percentage of users who contributed content to the different parts of the site. These percentages are all roughly the same: between 36 and 46% of users are contributing photos, lists, comments, and profile information to the site. These data indicate that the points system over the long term did not encourage more users to contribute to the site. So at this point in our analysis of the system, it appears that the long term impact of the points system is that users contribute *more content* to the site but the same number of users contribute as would without the points system.

4.2 Impact of releasing points

In releasing the points to the entire site so that PtsGroup2 would see the system, we hypothesized that we would observe a jump in the amount of content added by PtsGroup2, replicating result observed in our previous study. Comparing the contribution of the groups the week before and after releasing the points finds significant interaction between time and group in terms of number of points earned ($df=1, 418, F=4.19, p=0.02$), number of added lists ($df=1, 418, F=5.35, p=0.01$), and number of added comments ($df=1, 418, F=3.71, p=0.02$), but not significant for photos ($df=1, 418, F=0, p=.98$) (These tests were repeated-measures ANOVAs, considering time as the repeated measure, group as the independent variable and total amount of new content added to the site as the dependent variable). The results, in Figure 3, show that the group new to points system started adding more content, earning more points. These results charts are very similar to those from our previous experiment, so it appears that the release of the points had the same effect as the introduction of the points to the initial, smaller group.

Focusing on just the content added by PtsGroup2, we compared different content types added by them over time. As mentioned before, users received the highest number of points for adding comments, followed by lists. As a result, we expected to observe a more significant increase of contribution in the form of comments, and lists. The result is shown in Figure 4, supporting our expectation that the rate of growth of comments and lists are higher than photos.

Our final analysis of the effect of releasing points to PtsGroup2 was looking at the percentage of users in each group that contributed the different content types. As shown previously in Table 3, there was little difference between the two groups before the release of the points: between 36% and 46% of users were contributing in both groups to all categories of content. This may indicate that the points system encourages quantity of contributions, rather than encouraging new users to contribute to the site. But there is now evidence to dispute that, at least in the short term. Three weeks after releasing the points to PtsGroup2, the number of users that added new content to each content type is much higher than in PtsGroup1. Between 3.4% and 6.4% of the PtsGroup2 users added content to site for the first time, compared to 0.9 to 1.9% of PtsGroup1. This indicates that in the

short term, the points system encouraged users to add content for the first time. It may be that this finding reveals a distinction between short-term and long-term effects of incentive systems: that short-term they encourage new users to contribute, but over the long-term those users would have contributed anyway. Our consistent finding is that this incentive system encourages users to share more content than they would have otherwise.

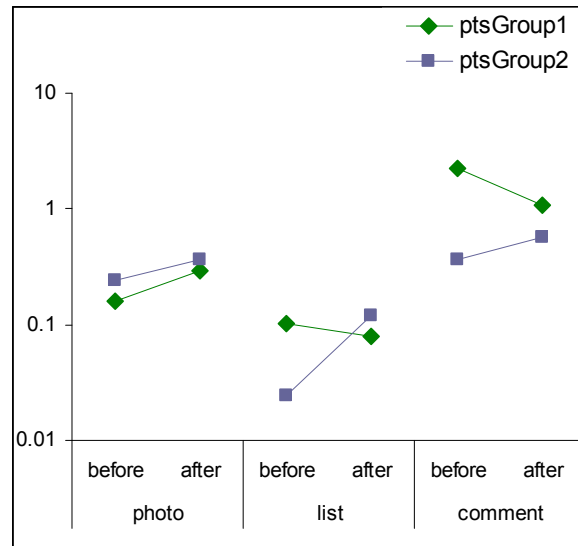


Figure 3¹. The average number of photos, lists, and comments added by each group each week

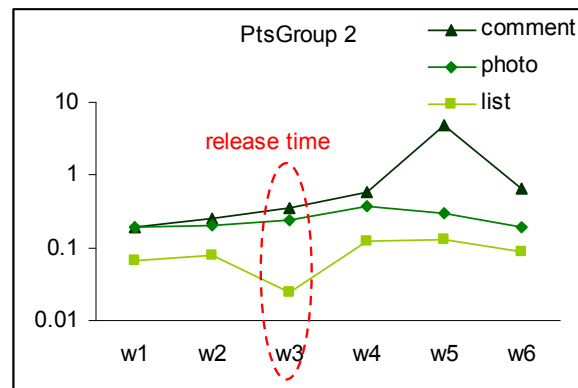


Figure 4¹. Comparing the amount of each content type added by PtsGroup2 over time

Table 4. During the three weeks after releasing points, the percentage of users in each group to add photos, lists, and comments for the first time.

	Photos	Lists	Comments	Profile
PtsGroup1	0.9%	0.9%	1.9%	1.4%
PtsGroup2	3.4%	6.4%	5.3%	5.3%

¹ For better presentation, the graphs in Figure 3 and 4 are shown on a logarithmic scale.

5 USER REACTIONS TO THE POINTS

Shortly after the points system was released to the entire community, members of the community began to mention the points and the different bee levels on their profiles, in the comments, in their status messages, and as the topic of discussion for lists. While most of the site content relating to points can be summarized as users talking about their own point level and bee status level, a vocal minority began discussing the points system in general, having a lengthy discussion of the pros and cons of the system and suggesting alternatives. Another small set of users openly gamed the system by adding low-value content to the site for the purpose of gaining points. This section describes these different explicit responses to the system. When we performed this analysis, the points system had been released to the entire site for two months and the system had over 4000 users.

5.1 Discussing merits of the points system

The most visible discussion of the merits of the points system occurred within seven different lists on the site that had 87 comments made on them. Outside of the site, there was also a blog post with 19 comments, a discussion forum thread with 20 posts that both debated the pros and cons of the system, and a company-internal “social software podcast” invited on one of the vocal members of these discussions specifically to lead a conversation during the podcast about the drawbacks of our site’s points system. All of these discussions involved 39 people, not counting the site’s project team members. (We, the team, posted minimally to the discussion, only to clarify the mechanics of the system and ask for clarifying points from the users on what they would like to see and to understand which parts of the system were particularly objectionable to users.)

The main arguments levelled against the points system were:

Users cannot opt-out. This point was raised by some of the highest point earners involved in the discussion. They felt they should have a choice as to whether or not their activity level was revealed on the site. They felt the visibility of their activity level was a violation of their privacy. Others stated that they personally would not opt-out of the system, but they agreed with this opinion and wanted this ability added to the points system.

Points do not reflect quality or meaning of the contributions. There was a lot of discussion about how quantifying contributions lessens the meaning of the contributions. An ideal system would measure the quality of the contributions in more qualitative ways, which users admitted was difficult to do. As an alternative opinion, some felt that any incentive system detracted from the real “value” of the site, which is the personal connections made between people. Those who agreed with this would prefer the site have no incentive mechanism in place.

Desire for customizable systems and points for more actions. A small number of users suggested that users be able to craft the incentive systems by assigning point values to the things they felt were most valuable. Others wished users could earn points for “friending” because that was what they perceived as the most valuable activity.

The users involved in the discussions who defended the points system also made some important points about its limitations:

When getting started, the points system can be motivating, but after you attain a high level of points, the incentive no longer works. The users who felt this way were high-point earners, contributing to many discussions elsewhere on the site, and connecting to many people. So while they stated that the points were “meaningless” to them, they also stated that they felt when they were starting out on the site, the points were a helpful way to gauge progression through the site.

The points help users distinguish experts from newbies. A couple of users specifically approved of the label “New bee” as a good label to attach to users who were new to the site and perhaps needed encouragement to start contributing. The “Super bee” label was mentioned by others as a way of finding out which users were experienced with the site and heavy users.

Improve the points with a decay function. An issue that had come up in our earlier paper’s user interviews [6] and was also brought up in this discussion was that users would like to see the points decay over time. The main reasons for wanting points to decay or reflect time in some way was that they wanted to have a way to detect recent high contributors to the site, rather than the current model which highlights the highest contributors over a long period of time.

From these valuable discussions, we learned that our users have well thought-out opinions about incentive systems and their role within an enterprise environment. They have contributed to our overall understanding of these systems. We also agree with many of their suggestions and plan to implement the opt-out feature and possibly a time-based decay function on the points.

5.2 Gaming the points system

Most of the discussion on the site about the points systems was led by a single user who felt passionately that users should be able to opt out of the points system. To demonstrate his position that the points system should not be a representation of a person’s contribution to the site, he launched a keyboard macro that submitted comments onto one of his photos with the single word “test.” This generated 1000’s of points for his profile and raised him to the second highest point position on the system. (He stopped the macro just short of the top position because he was friends with the person in the top position and said he did not want to offend him.) Once the flurry of activity around the points discussion died down, this user deleted the photo that had all of these “test” comments, and thus returned his point value back down to a lower level.

In addition to this example of gaming, we have evidence of three other instances of gaming the points system.

Two users earned 100 points each by adding content to their profile that mentioned the points system and did not provide any information about themselves. These two profile self descriptions were “What do you think about putting this here to get the 100 points I needed to be a nicer bee than a noobie?” and “Why [am I] entering a question here? I read on the newsletter that you will get 100 points and hence here I am...all part of the master plan to becoming a Queen Bee).”

The third instance of gaming was a user who posted the exact same comment on 21 different profiles. The comment said “You have zero beehive points. Maybe that's something to work on, so you don't look like such a newbie. Thanks for letting me make points on you just for posting this comment!”

Overall, these instances of gaming seem fairly benign considering there were over 4000 users on the site at the time of our analysis. In designing any incentive mechanism, it is always a challenge to design one that discourages gaming. Although the points system is technically easy to game, we believe it is *socially* difficult to game because users are communicating with their colleagues every time they post content. By gaming the system, one's coworkers become acutely aware of the gaming because of the system's network notifications.

As was brought up in the discussions about the merits of the points system though, there is an issue of low quality contributions being counted equally with high quality contributions. We did not assess the quality of the comments being left on the site because the quality of communication between two users cannot be judged by an external party. For example, a simple "hi" can convey a lot of meaning between colleagues, which an automated analysis would miss. Therefore we have not been able to assess whether the quality of the comments, photos, or lists left on the site has decreased due to the points system.

5.3 Mentioning one's own points

Outside of the discussion about the merits of the system, there is evidence throughout the site that users were aware of the system and had a desire to earn more points. Sixty-five people commented about their points in a variety of locations on the site: 21 people commented on people's profile pages about their own or their colleague's points; two people mentioned the points in their profile self descriptions (in addition to the two who gamed the points); one person used his list's comment section to test out how points were earned; two people talked about earning points while adding a comment about photos; and 37 people mentioned points in their status messages.

An interesting aspect to the status messages is that one does not earn points for changing one's status, so these comments were not motivated by a desire to earn points. Most of the status messages were about one's point level: "just a new bee :(" "a simple worker bee." "figuring out how to get points," "just over 800 (make that 900) points, narf, narf," "deciding whether to be a busy bee or not," and "wow, Ricardo has 15000+ points, and I have 80... suddenly, I feel inadequate..."

5.4 Summary of employees' responses

At the conclusion of this several week interaction within the community, we believe there are several important lessons to be learned about introducing incentive systems to an existing site. First, established community members, given the opportunity to give feedback, will have strong opinions on the introduction of a feature such as this that ranks users and establishes an evaluation criterion. Second, much of the feedback provided by the community was insightful and the suggestions for changing the system were provided from the perspective of users within a larger ecosystem, our enterprise setting. So we believe that the suggestions made for changing the points system should be taken very seriously. Based on their feedback, we plan to provide an opt-out ability so that users can elect to not have their point levels shown on the site. Our third lesson is that a small feature

does have an impact on how people communicate on the site, from the gaming to the joking references about the points. Discussion of the single feature was pervasive and was not limited to the features of the site that granted points.

6 CONCLUSION

Our qualitative and quantitative results show that overall the points system is a successful method for motivating users' contributions in an enterprise social networking: it boosts users' contributions when introduced, particularly when paired with an email notification, and over time the users who have points visible to them contribute much more content to the site, albeit not at higher levels of user participation. Furthermore, it also stimulates discussion amongst the users.

A broader implication of this finding is that if designers of a site wish to generate a burst in user activity, the introduction of a simple points-based system can accomplish that. Therefore, in addition to increasing the content on the site overall, the launching of a points system can be seen as a way of controlling the behavior of the users to generate a predictable result at a specific moment in time. This gives an opportunity to site owner's to target users and could have value for marketing efforts.

7 REFERENCES

- [1] Bretzke H., Vassileva J. (2003) Motivating Cooperation in Peer to Peer Networks, In Proc. Of UM03, Springer Verlag LNCS 2702, (2003), 218-227.
- [2] Beenen G., Ling K., Wang X., Chang K., Frankowski D., Resnick P., Kraut R.E., Using social psychology to motivate contributions to online communities, In Proc. CSCW 2004, ACM Press (2004), 212-221.
- [3] Cheng R., Vassileva J. Design and Evaluation of an Adaptive Incentive Mechanism for Sustained Educational Online Communities. User Modeling and User-Adapted Interaction, special issue on User Modeling Supporting Collaboration and Online Communities (2006), 16 (2/3), 321-348.
- [4] Cheng, R., Vassileva, J.: User motivation and persuasion strategy for peer-to-peer communities. In Proc. of the 38th Hawaii International Conference on System Sciences (Mini-track "Online Communities in the Digital Economy"), IEEE Press (2005), 193-202
- [5] Farzan, R., and Brusilovsky, P. Social navigation support in a course recommendation system. In: V. Wade, H. Ashman and B. Smyth (eds.) In Proc. AH'2006, Springer Verlag (2006), pp. 91-100
- [6] Farzan, R., DiMicco, J. M., Millen, D. R., Brownholtz, B., Geyer, W., and Dugan, C. (2008). Results from Deploying a Participation Incentive Mechanism within the Enterprise. To be appear in proceeding of CHI2008.
- [7] Rashid, A. M., Ling, K., Tassone, R. D., Resnick, P., Kraut, R., and Riedl, J. Motivating participation by displaying the value of contribution. In Proc. of CHI 2006, ACM Press (2006). 955-958.
- [8] Kollock P., The economies of online cooperation: gifts and public goods in cyberspace, In: Smith, M.A., Kollock, P. (Eds.), Communities in Cyberspace. Routledge, (1999) London.
- [9] Ling K., Beenen, G., Ludford, P., Wang, X., Chang, K., Li, X., Cosley, D., Frakowski, D., Terveen, L., Rashid, A. M., Resnick, P., and Kraut, R. Using social psychology to motivate contributions to online communities. JCMC, 10(4), 2005.