Designing Mechanisms to Stimulate Contributions in Collaborative Systems for Sharing Course-Related Materials

Julita Vassileva, Ran Cheng, Lingling Sun, Weidong Han
Computer Science Department, University of Saskatchewan,
57 Campus Drive, Saskatoon, S7N 5A9 Canada
{jiv, rac740, lis102} @ cs.usask.ca

Abstract
This paper presents several mechanisms for stimulating contributions from students in a peer-to-peer environment. The evaluation of a status and visualization based mechanism showed a significant increase in student participation and contributions, but a decline in the quality of contributions.

Keywords: motivation, participation, persuasion, collaboration, peer-to-peer, evaluation

1. Introduction

Peer-to-peer (P2P) systems have a lot of potential uses in classroom and academic environments. The concept of P2P assumes decentralization and equality of participants: there is no central provider or services or material, but the users who consume the services or materials are also the providers of services and materials. Before the emergence of specialized P2P protocols for file-sharing, like NAPSTER, Gnutella, Chord, etc., the concept of P2P has been applied in collaborative learning environments, like Phelps (Greer et al, 1997), I-Help (Vassileva et al., 1999), where learners can ask for help their peers and receive help from their peers. Phelps was implemented with a client-server architecture, while I-Help was implemented using a multi-agent architecture and was initially fully decentralized. Later versions were centralized (all agents ran on a powerful server with an Oracle database).

More recently, P2P file-sharing protocols have been used to design shared but fully distributed learning object repositories. For example, Edutella (Neidl, et al. 2002) is based on the JXTA platform. Comtella (Vassileva, 2002) uses the Gnutella protocol and allows graduate students to share research papers that they have found on the web and stored locally on their disks.

P2P systems are fairly easy to build; they do not require any sophisticated models of collaboration, since they are mostly based on loose cooperation by the users. Users typically share resources in an asynchronous way and do not engage in dialogues or collaborative problem solving; these systems provide only the infrastructure, search and match-making, they do not facilitate explicitly the knowledge building process. There is no need for creating a pool of resources, or knowledge-base bootstrapping, as in other systems; the users provide the resources themselves. Yet, these systems are collaborative by nature, i.e. they can not exist without user cooperation, i.e. good will to share resources, to participate in the system, be on-line and answer requests. This is a much broader definition of collaboration than the one typically used in the area of CSCL, but even ensuring this level of collaboration is not straightforward. The lack of user participation has been the main problem in most such systems. It has been a problem also in online communities where statistics show that most users are lurkers or free-riders. Designing a successful on-line community is hard and most attempts fail. Only few evolve into sustained productive groups. In the next three sections we present our experience in designing mechanisms encouraging participation on different levels in a P2P system.

2. Ensuring presence

P2P systems depend very much on who is on-line at the moment, i.e. how many users have their servents running. A servent that is not running can not respond to queries, and the resources / services shared by the user are unavailable to the other users in the system. In addition, in some protocols, like Gnutella, servents participate in forwarding queries of other users, i.e. they are part of the infrastructure. If they are unavailable, parts of the infrastructure are unavailable too, and messages can not be transferred. Therefore, ensuring that users keep their servents running is important.

However, this turned out to be one of the main blocks on the way to deploying I-Help and Comtella. In the deployment of both systems, we observed that students activate their servent only when they needed help or searched for some paper. There are good reasons for this behaviour: having an extra application running all the time takes computing resources and space on the screen (even when the application icon is shrinked down on the task bar or running in the background). Unfortunately, this selfish behaviour makes it impossible for the community of simultaneous users to reach a critical mass to allow for everyone to find what they are looking for. As a result, students who try the system and don't find anything won't log in again and they are lost (together with their resources) for the community. There is a simple technical solution to avoid such negative feedback loop – to move all the user servents (or agents) on a server, where they can be running all the time and let users log-in to their own agents or servents when they need to search for resources or services themselves. This solution worked well in the case of I-Help. We applied it also in an adaptation of the Comtella system to support the "Ethics and IT" class (Vassileva, to appear) and even though it is a compromise with the architectural purity of the system (since it imposes a centralization), we argued with ourselves that it is only an implementation issue, since on a higher level the system remains distributed. All the servents communicate using the Gnutella protocol, i.e. any subset of servents can be moved to different servers, or back to the user computers and they will still work in the same way. Hybrid solutions, where some of the peers reside on user computers, if the users are willing to keep them running, and some reside on servers (where users can not ensure that they will keep their servents running) are also possible. In this way, by placing some or all the servents on a server, the basic level of participation (keeping the servent running) was ensured, therefore, access to all shared resources of all peers was ensured all the time.

3. Ensuring regular new contributions

A P2P file-sharing system without new materials shared regularly by at least some of the users quickly reaches equilibrium: all users have already found and downloaded the files they are interested in. If they don't find anything new of interest when they search several times, they won't try again and will loose interest in the system. Maintaining a constant stream of new contributions (or shared materials) is very important for the success of the system. However, this problem can not be solved so easily as the problem of ensuring presence. To contribute new resources, the users need to invest some effort, including finding the new material, annotating it, performing the needed actions to share it with other peers. We applied several different methods of motivating users to make contributions.

3.1. Convenient interface for sharing new resources – "zero-effort" sharing

For example, a pop-up window can appear in the browser reminding the user to share the file, if the user spends a certain amount of time on a web-page or if she prints a page, or if she is looking at a PDF document. The window can automatically suggest annotation for the resource using tools like semantic classification of the document based on the text, URL etc. We applied such proactive reminder in a previous version of the Comtella system for sharing academic papers, however some of the users found it somewhat too intrusive. If activated only for a certain type of files (PDF and PS, in which typically research papers are available on the internet), it is more acceptable. However, in our "Ethics and IT" class the resources to be shared were articles appearing on web-journals, like www.wired.com, www.itbusiness.ca, the web-edition of New York Times, Atlantic Monthly, or on discussion

forums, on-line communities like Slashdot.org, etc. The format of files from such sources is too heterogeneous and automatic prompting (without using heavy semantic classification tools) is unlikely to be successful: it will either omit too many files or it would be annoying.

3.2. Rewarding contributions

Reciprocation is a basic norm in human society [4]. Reciprocation, through rewarding users for sharing files has been used in existing P2P filesharing systems. Some approaches introduce the notion of artificial currency and micro-payments, i.e. create a virtual market. We attempted this approach earlier (Vassileva et al., 1999, Kostiuk & Vassileva, 1999) in I-Help, but it didn't stimulate much participation (Vassileva, Deters, 2001). Partially this was due to the inappropriate way our e-currency was "cashed" in the experiment. Instead of cashing it in grades (e.g. participation marks), we distributed souvenirs to the top users, which didn't thrill the students. However, before attempting a second time, in a class that was in our control to award participation marks based on accumulated currency, the fate of Mojo-Nation and Clay Shirky's article on micro-payments (Shirky, 2003) convinced us that any payment-based approach will not work because the act of "buying" a resource, even at a very small price, creates mental transaction costs, that is, energy required to decide if the resource is worth buying or not.

There are other approaches to rewarding participation. KaZaA rewards active users with a better quality of services. The system records the actions of users and maintains a numeric participation level for each user. The speed of downloads the user can get is based on this value. We applied in a previous version of Comtella (Sun et al., 2003) a similar idea, but it turned out that a better download speed was not seen by users as a particular reward, since the speed was very good anyway, the files were relatively small and a couple of seconds didn't matter. In the "Ethics and IT" version of Comtella, we rewarded user participation with access to more powerful search options (remove duplicate results, sort results by different criteria, show only new resources, shared after the user's last logout, etc.). While the data from the usage of these options as well as the user questionnaire showed that these extra-functions were valued by the users, they could have acted as a double-edged sword, since they empowered the users who were already actively using the system. Perhaps allowing these functions for users who didn't participate actively would have made the system more attractive for them.

3.3. Social Visibility

One fundamental way how people decide what to do in a situation is to look at what other people are doing or have done. If many individuals have decided in favour of doing something, more people will tend to follow their way. Moreover, a group of people sharing some similarity can influence each others' behaviour more effectively (Cialdini, 2001). According to the social validation theory (Cialdini, 2001), it would be possible to persuade people to make contributions to a P2P community by demonstrating that many people just like them have contributed a lot to the community and benefited from their contributions. Visualizing the community and the level of contribution of each member in some appropriate way can serve this purpose (Bretzke & Vassileva, 2003). Again, however, the visualization if not well-thought can be a double edged sword. If only a small portion of all users are active, it wouldn't be a good idea to provide a whole picture of the community, because the behaviour of the inactive users can discourage the user from contributing, if she identifies with them. It is possible, and even in some cases it is recommended (Erickson, 2003) to misrepresent the actual level of contribution.

We implemented visualization of the user community in the previous version of Comtella for sharing academic papers among graduate students. It showed the currently online members as stars with size dependent on the number of their shared files. The size was relative to the total number of files shared at the moment. This lead to some confusion among the users, since it was possible to share very few files and one day appear as a tiny star but the next day as the biggest star, since there was currently nobody online sharing more files. Also, when very few stars were present, the visualization actually worked discouraging, making users feel alone. The conclusion was that we needed to represent always all the members and to have absolute rather than relative classification of the contribution level, so that users see consistently their size grow with their contribution and not in relation to who is currently on line. The good news was the fact that users actually cared about why their star that was so big and bright the previous day has shrunk to a dwarf. In the Comtella version for the Ethics and IT class we had a permanent visualization of all

users, independently whether on- or off-line. The online users were visualized with filled circle, while the offline users were visualized with empty circle. There were 4 possible sizes and they were used to show different aspects of contribution: number of new shared papers, number of papers downloaded from others, frequency of logging into the system, etc. The visualization was well used by the students, and the questionnaire showed that students cared about how they appeared – the majority replied that if they see that they are little stars and there are many large stars, they will try to contribute more to improve their standing in the community.

3.4. Status: Combining Reward and Fear

Discrete emotions have unique appraisal patterns, motivational functions and behavioural associations. According to the theory of fear, people feel fear when they perceive some threat to themselves or their properties. This fear makes incoming messages, especially those containing reassuring information or information about how to avoid the threat more persuasive (Nabi, 2002). One possible persuasion strategy based on fear will arouse fear in the user to loose something, for example, some privilege. At that point, information how to avoid the problem is provided, e.g. by contributing more resources. This information will be more persuasive than just a general advice by the system.

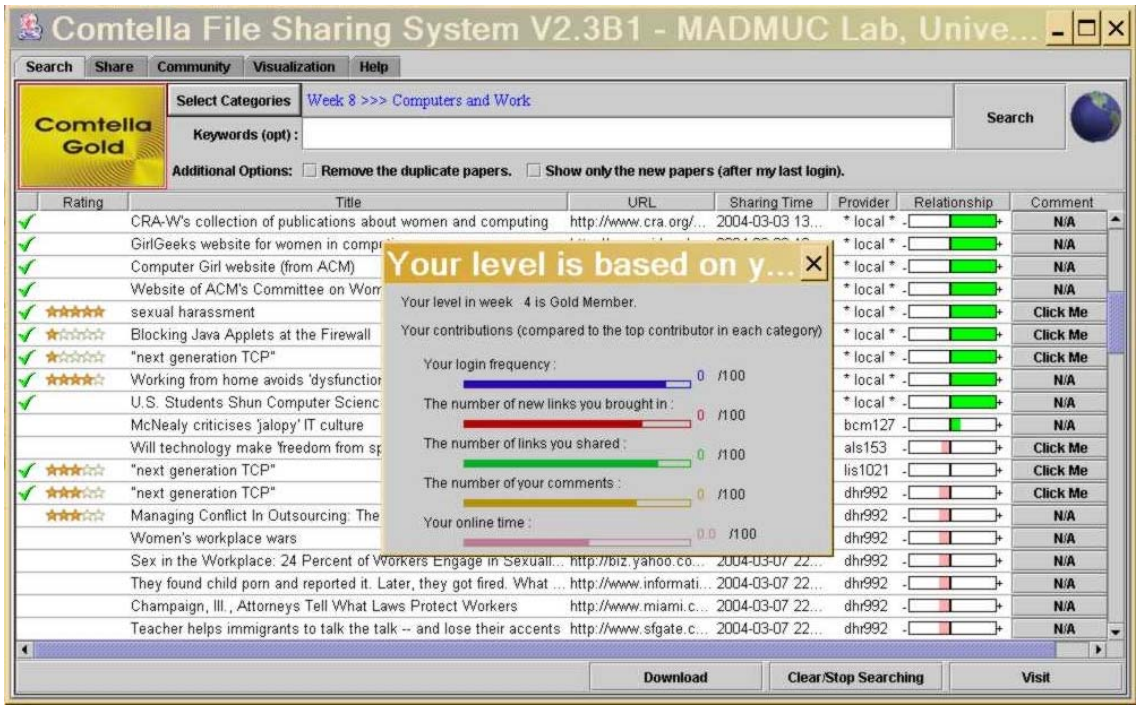

Figure 1: The Comtella Interface showing the membership card and explanation of the level.

We implemented in the Ethics and IT version of Comtella this idea by introducing a set of hierarchical memberships in the community depending on the level of contribution of the user. The levels are currently three: "bronze", "silver" and "gold" (see Figure 1). The users with higher level membership are rewarded with better search functionality as explained in section 3.2. and better visibility (there is a special view "By Status" in the visualization, where all the gold, silver and bronze members can be seen). Therefore, according to the reciprocation theory (Cialdini, 2001), there are rewards for higher participation and higher status. However, when a user has reached a gold membership level, there is little incentive for her to contribute more. Therefore, the status is only temporary, in our case, on a weekly basis. According to the theory of fear (Nabi, 2002) if a silver- or gold-status user, who has already enjoyed some privileges, stops contributing to the community and the system shows a warning, this will arouse fear that her membership will be degraded in the next week. At this time, a message related to the actions that the user can take to avoid demotion will provide effective persuasion. This was achieved through a window showing

the level of contribution of the user with respect to several factors in comparison with the top-contributor for the week for each factor.

Essential for this strategy is how the participation level is calculated. In our experiment we used a weighted sum of the following factors (in decreasing order of the weight): number of new resources brought, number of ratings or comments on resources, number of downloaded resources from others, relative frequency of logging in, total time spent on line. The students were very curious to know the formula and immediately attempted a variety of ways to "cheat" the system, especially the ones that didn't involve significant effort: keeping logged in for long time, logging in and out frequently, downloading many files from other peers. We also saw users who contributed new links that were not related to the topic of the week, and unfortunately, they were successful, since our participation formula did not take into account the quality of the resources shared.

3.5. Results

In the first 6 weeks of the class we used Comtella without status, visualization and reward for participation. We introduced the version implementing the motivational mechanism in the middle of the term (after week 7). There were 35 students registered in the class, 29 of which used the system actively. Figure 4 shows the number of new links shared by students each week. The bar for week 6 represents two weeks, since the two weeks discussed the same theme.

It is evident that the number of new articles shared over the last four weeks is about twice bigger than that in the first six weeks (641 versus 332). Although there is a decrement in the last two weeks (due to the fact that the students had to focus more efforts on their course projects, scheduled for presentation in the last week of the class), the number of new resources shared increased significantly. Besides, other kinds of contributions, such as ratings and comments of articles, increased as well after the motivation mechanism was introduced in the 7th week.
However, the quality of shared resources somewhat decreased. Several (4) users shared a lot of articles not related directly to the topics of the class to gain a higher membership level or maintain their gold level. One extreme case was the top contributor of the class, who shared 131 of the 973 links (13 %), of which 40 (~30%) were irrelevant. Since the hierarchical memberships were introduced into the community, the quantity of the users' comments increased but the length of the comments became shorter on average. Obviously, in the next version of the system, the participation metric has to include a measure of the quality of shared resources and comments.

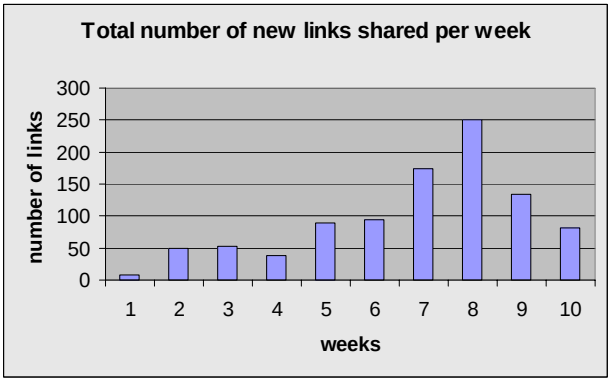

Figure 4. The total number of new links shared by all students over the 10 weeks.
(The motivational version was introduced after week 6.)

4. Ensuring high-quality contributions (current work)

Obviously, integrating persuasion mechanisms can effectively encourage participation and contributions in a P2P community. However, whenever there is some reward (independently whether it is tangible or simply social visibility), people will be tempted to cheat the system. Therefore, adequate measures should be taken to discourage

cheating behaviours. Our current mechanism allows users to cheat by submitting low quality resources. It is important to include controls of quality, and it would be the best if the community, i.e. the users are empowered to exercise this control.

We are currently designing a mechanism to encourage users to rate fairly resources and comments. One possible approach is to include the quality of contributed resources (e.g. the average rating by other users) in the participation formula. However, this may encourage users to rate highly the papers of their friends, or generally inflate the ratings. To ensure fair ratings it will be necessary to measure in some way the credibility or competence of users as raters, e.g. to compute a reputation (Wang & Vassileva, 2003), similar to the "karma" used in Slashdot.org. Incorporating techniques measuring similarity of users' ratings will allow forming sub-communities of users with similar interests and criteria, which can be useful to extend the system functionality into recommending resources (similar to collaborative filtering). The more ratings a user submits, the better recommendations she can expect from the system, which will create a positive feedback loop and reward for fair and frequent ratings. The visualization will be able to show users based on their reputation, and clustered by similarity of ratings, which will elevate user participation to a higher, more knowledge-dependent level.

Acknowledgements: The authors thank Jeremy Long for implementing the interface of the Comtella version for sharing course materials. This research has been supported by the NSERC Discovery grant of the first author. Comtella is available for download from `http://bistrica.usask.ca/madmuc/comtella.htm`

References

Bretzke H., Vassileva J. (2003) Motivating Cooperation in Peer to Peer Networks, *Proceedings User Modeling UM03,* Johnstown, PA, June 22-26, Springer Verlag LNCS 2702, 2003, 218-227.

Cialdini. R. (2001) Influence: The Science of Persuasion. Scientific American, (Feb. 2001), 76-81.

EDUTELLA: a P2P Networking Infrastructure based on RDF. In Proc. of 11th World Wide Web Conference (Hawaii, USA, May 2002).

Erickson, T. (2003) Designing Visualizations of Social Activity: Six Claims, In Proc. CHI'03.

Kostuik K. and J.Vassileva (1999) Free Market Control for a Multi-Agent Based Peer Help Environment, in Proceedings of the Workshop on Agents for Electronic Commerce and Managing the Internet-Enabled Supply Chain, held in association with the 3rd International Conference on Autonomous Agents (Agents '99), Seattle, May 1-5, 1999.

Nabi. R. (2002) Discrete Emotions and Persuasion, The Persuasion Handbook: Developments in Theory and Practice, In J. Dillard & M. Pfau, Sage Publications, ISBN: 0761920064, P 289-308.

Neidl, W., Wolf, B., Qu, C., Deceker, S., Sintek, M., Naeve, A., Nilsson, M., Palmr, M., AND Risch, T. (2002)

Sun L., Upadrashta Y., Vassileva J. (2003) Ensuring Quality of Service in P2P File Sharing through User and Relationship Modelling, UM03 Workshop on User and Group models for web-based adaptive collaborative environments, Johnstown, PA, June 2003.

Shirky. C. (2003) Fame vs Fortune: Micropayments and Free Content, First published September 5, 2003 on the "Networks, Economics, and Culture" mailing list, available online at: http://shirky.com/writings/fame_vs_fortune.html

Vassileva J. , J. Greer, G. McCalla, R. Deters, D. Zapata, C. Mudgal, S. Grant (1999) A Multi-Agent Approach to the Design of Peer-Help Environments, in Proceedings of AIED'99, Le Mans, France, July, 1999, 38-45.

Vassileva J., R. Deters (2001) Lessons from Deploying I-Help, in J. Whatley & M. Beers (eds.) Proceedings of the Workshop on Agents and Internet Learning, AIL'2001 at the Autonomous Agents'2001 Conference, Montreal, May 28, 2001.

Vassileva J. (2002) Supporting Peer-to-Peer User Communities, in R. Meersman, Z. Tari et al. (Eds.) "On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE" Coordinated International Conferences Proceedings, Irvine, 29 Oct - 1 Nov. 2002, LNCS 2519, Springer Verlag: Berlin-Heidelberg, 230-247.

Vassileva, J. (to appear) Harnessing P2P Power in the Classroom, in Proc. ITS'2004, Maceio, Brazil, August 2004.

Wang Y., Vassileva J. (2003) Bayesian Network-Based Trust Model, Proc. of IEEE/WIC International Conference on Web Intelligence (WI 2003), October 13-17, 2003, Halifax, Canada.