

Using Community Visualization to Stimulate Participation in Online Communities¹

Julita Vassileva

University of Saskatchewan

Lingling Sun

Solutions AB TELUS Business Transformation

ABSTRACT

A motivational community visualization was designed to encourage users to participate more actively and to bring more contributions in an online community. The visualization is inspired by the theory of social comparison in social psychology. It evolved through two designs: a fixed and a customizable one for two different communities of students sharing papers – one of graduate students in a research lab and another one for students in an undergraduate class. The article discusses the features of the two communities, each of the two visualization designs, and their advantages and disadvantages.

Keywords: participation, online communities, social visualization

INTRODUCTION

The “cold-start” problem is well-known in most online communities, e.g., discussion forums, IRC systems, peer-to-peer (P2P) file-sharing networks, social networking, or blogging systems. While some Web-based online communities manage to attract users and grow enormously, others never reach the “critical mass” of active users needed to ensure enough new materials and attract users to revisit the community. Often, applications for expertise or document sharing created to serve a specific role in a certain geographic or organizational context remain unused. It seems like a “chicken and egg” problem: a community is useful and interesting only if many people are participating and contributing; but to get users to contribute, a useful and interesting community has to exist first.

1. An earlier version of this paper was published in the proceedings of the 12th International Workshop on Groupware: Design, Implementation, and Use, CRIWG 2006, held at Medina del Campo, Spain.

We propose to motivate users to participate by visualizing the community and the levels of participation of all community members, hoping that the social visibility will enact social norms, stimulating users to engage in responsible and reciprocal behavior. By participation, we mean activities which benefit the community and demonstrate involvement in the community, like contributing materials, rating and commenting materials contributed by others, logging into the system, and reading materials contributed by others.

The article is organized as follows: in the next section, we give an overview of some of the work on motivation from the area of social psychology as well as other approaches using visualization in online communities. Then we describe our first design of the motivational visualization and the lessons learned from its brief deployment. The next section presents the improved design of the visualization and the results of its evaluation. A discussion of the results and their implications follows. The paper ends with a conclusion and directions for future work.

RELATED RESEARCH

Asch's (1951) conformity study proved that people generally want to "fit in groups." The fitting behavior at an interpersonal level happens, for example, when someone sees a friend doing or believing in something and starts believing or doing the same thing. Fitting, at a collective level happens, for example, when one sees trends in the behavior of others, e.g., the style in dressing, and changes his/her own style of dressing to fit in, even if only for a particular occasion.

Researchers in social psychology have studied human motivation in real (geographically, socially, or organizationally based) communities. Leon Festinger (1954) created the social comparison theory showing that people tend to compare their achievements and actions with people who they think are similar to them in some way. For example, when a student wants to know if she is good at math, typically, she compares herself with the other students taking the same math class, rather than with her professor. When there is no suitable peer group, people will compare themselves with almost anyone. On the other side, when one knows that others will compare with him/her, one acts more responsibly. People normally want to be positively recognized in their community and are willing to make an effort to gain social reputation, providing the effort is affordable and worthwhile compared with the potential benefit of the reputation.

Competition is a form of upward social comparison in which one compares and tries to "fit in" with the elite, top-performing sub-group (Garcia and Tor, 2005). It seems that social comparison and competition can motivate participation in users in online communities too. Shepherd et al. (2003) showed that social comparison decreases social loafing and increases productivity in groups that are brainstorming and sharing ideas in e-collaboration. Inducing social comparison with a graphical feedback tool and increasing the salience of that social comparison with facilitation techniques, the authors in-

creased the productivity of electronic brainstorming groups by 63 percent. This shows that the designers of software infrastructures for online communities can exploit social comparison to influence user behavior.

For social comparison to take place, however, users have to be made aware of the behavior of other users as well as of their own behavior. Visualization has been used in online communities to create awareness about the other users and the things happening in the community, but it has not been used to stimulate social comparison or competition so far.

For example, VIUM (Uther and Kay, 2003) is designed to stimulate user reflection on their learning and displays learning concepts as a graph of texts in different fonts, colors and brightness to represent how much the user knows about each concept. Social visualization approaches using different metaphors have been proposed to create “social translucency” and activate social norms in online communities. For example, the Babble System (Erickson and Kellogg, 2004) visualizes a chat room as a pie with moving dots representing users who are actively participating in the conversation (those whose dots are close to the center) and users who are mostly listening (on the periphery). The Chat Circles (Viegas and Donath, 1999) uses circles filled with different colors and texts representing the conversations. The Task Proxy (Erickson et al., 2003) visualizes task groups in a company as differently-colored cells in a honeycomb creating conditions for social comparison and embarrassment for those groups who are lagging behind. This embarrassment would create social pressure and motivation to perform on par with others.

Erickson proposed a number of guidelines for designing social visualizations for online communities (Erickson, 2003). He makes an important distinction between “translucence” and “transparency,” emphasizing that the information shown in the visualization does not necessarily have to be very detailed and exact. In most of the cases, it is better just to give the user a general idea, and even in some cases a certain level of misrepresentation may be beneficial. Also, customization should be avoided; all users should see the same thing so that they feel responsible for their actions, since they know that others see the same things as they and are aware of what they do.

General design guidelines have been developed in the visualization community, which use certain physiological properties of human vision and are aimed at reducing cognitive overload, increasing usability. The choice of metaphor is very important: an appropriate metaphor is intuitive to use and does not require a complex legend for interpretation. Applying a hierarchical structure (Eick and Wills, 1993; Sprenger et al., 2000) and using composable layout and visual sets (Pattison et al., 2001) are helpful when designing information-compact visualizations of large networks. A proper use of location and color contraction of visual components will successfully attract attention (Lamme, 2003). “Richly expressive information visualization is difficult to design and rarely found” (Humphrey, 2000), so it is always beneficial to make the visualization easily reusable in similar situations.

Comtella is an online community sharing resources (papers and links) developed

in the Multi-Agent, Mobile and Ubiquitous Computing (MADMUC) Lab at the University of Saskatchewan. The first version of the system (Vassileva, 2002) was a peer-to-peer (P2P) file-sharing system, used to support sharing materials among university students. The first P2P version of the system was used to support graduate students and faculty in sharing research papers. It was deployed in the MADMUC lab and across the Computer Science Department at the University of Saskatchewan. A new version of the system was developed in the fall of 2004 and it was used in the spring of 2004 to support students in a 4th year undergraduate class to share class-related papers downloaded from the Web. This system has subsequently evolved in Web-based, server-centric applications, which were used in the offerings of the same class in 2005 and 2006. An overview of all versions of the Comtella system and a downloadable version of the P2P client from 2004 can be found at <http://bistrica.usask.ca/madmuc/comtella.htm>.

A motivational community visualization was included from the very first Comtella implementation (Bretzke and Vassileva, 2003). The visualization itself also evolved through two versions, a fixed and customizable one. The community in both cases was closed and relatively small (under 40 users). The remainder of the article presents the design of these two versions, case studies with these designs in the context of their respective communities, and a discussion of the results and their implications for design of motivational visualizations for online communities.

In the rest of the article, we describe two versions of the community visualization for Comtella. The visualization was designed according to the guidelines mentioned above with the goal of stimulating social comparison and competition and increasing the user participation and contribution of resources. Each version was deployed in a case study not so much with the goal of formal evaluation, but more with the goal of an open-ended discovery of the effects, problems, and issues that need to be considered. As seen from the brief review of previous social visualization systems above, this is a novel use of visualization in online communities addressing the problem of participation, a problem which is gaining importance with the advance of Web 2.0.

FIXED VISUALIZATION FOR A P2P PAPER-SHARING COMMUNITY

The first version of the visualization was designed for the decentralized (Gnutella-based) version of Comtella. The system was intended to serve a small community of graduate students and faculty members to share academic articles as files in PDF format. We describe first the functionality of the system and our hypothesis about the role of the community visualization in encouraging participation. Then we present the fixed design of the visualization and, finally, the results of the case study.

Functionality

The goal of the system was to support the distribution and search of relevant papers within the group by keeping local copies of the articles. Usually, these copies were stored in some folder on the user's personal computer and, depending on the personal file management habits of the user, could be hard to find again. Using a P2P file-sharing system with a built-in search function allowed users to search conveniently for copies of articles that they had saved on their own hard disks, and then for articles saved on the computers of their colleagues, thus allowing sharing of relevant articles without the need of sending them as e-mail attachments.

To enable search in Comtella, which did not have full-text search functionality, the user had to classify every new article s/he shared according to a pre-defined hierarchical list of categories (a sub-set of the ACM subject categories with 3 levels of depth). In addition, they could add keywords to allow narrowing down the search. To remind users to share papers while they browse online, a small application in C++ was embedded in Internet Explorer, the main browser used in the lab, which prompted her to share in Comtella whenever the user opened a PDF file and viewed it for more than 30 seconds.

Since the system was based on the Gnutella P2P protocol, which is fully decentralized, there was no central repository containing the shared files, but the server residing on each user's computer held a local repository containing both the files uploaded by the user (new files) and the copies of files downloaded by the user from others. To ensure availability of files that can be searched, it was essential that many users were online at any time and that they re-shared the files they had downloaded from others. Therefore, two types of participation were needed for ensuring a sustainable community:

- sharing files (new or downloaded from others), thus ensuring redundancy and availability of resources necessary for the P2P search to work, and
- contributing new files, thus increasing the diversity of resources.

To enable user awareness of their community and encourage social comparison, we designed the social visualization to emphasize user participation of the two types listed above.

Hypothesis and Visualization Design

Our general hypothesis was that an appropriately designed visualization of the community will stimulate social comparison among the users and will result in increased user participation, measured by appropriate metrics.

The visualization shown in Figure 1 used the metaphor of a mid-summer night sky, first proposed in Bretzke and Vassileva (2003). This metaphor was chosen because it is rich (allows naturally different colors and sizes of stars and different ways of clustering

the stars), fits intuitively both with the popular expressions “being a star” and “cyber-space,” and is aesthetically pleasing.

The visualization showed the two desirable types of participation – the number of original contributions and the total number of files shared by each user. The graphical language of the visualization used two dimensions – the size of a star (large, medium, small) and the color of the star (red or yellow). The size of a star was determined by the user’s number of contributions (files shared in the Comtella community, both original and downloaded by others). Users who shared more than the average number of files per user in the community were represented by larger stars, those who shared around the average were represented by medium-sized stars, and those who shared less were represented by small stars.

The shade of the color of a star was either in red or in yellow. Reddish stars represented “givers” – users who brought to the community more new files than the number of files they downloaded from other users, while yellow stars represented “takers” – users who downloaded more files from other users than the number of new files they brought to the community. In this way, people who enriched the shared resources pool with new files were recognized by having “warmer” stars. The “sun,” a big yellow star, represented the “best user” among those who were online at the time of the viewing. The “best user”

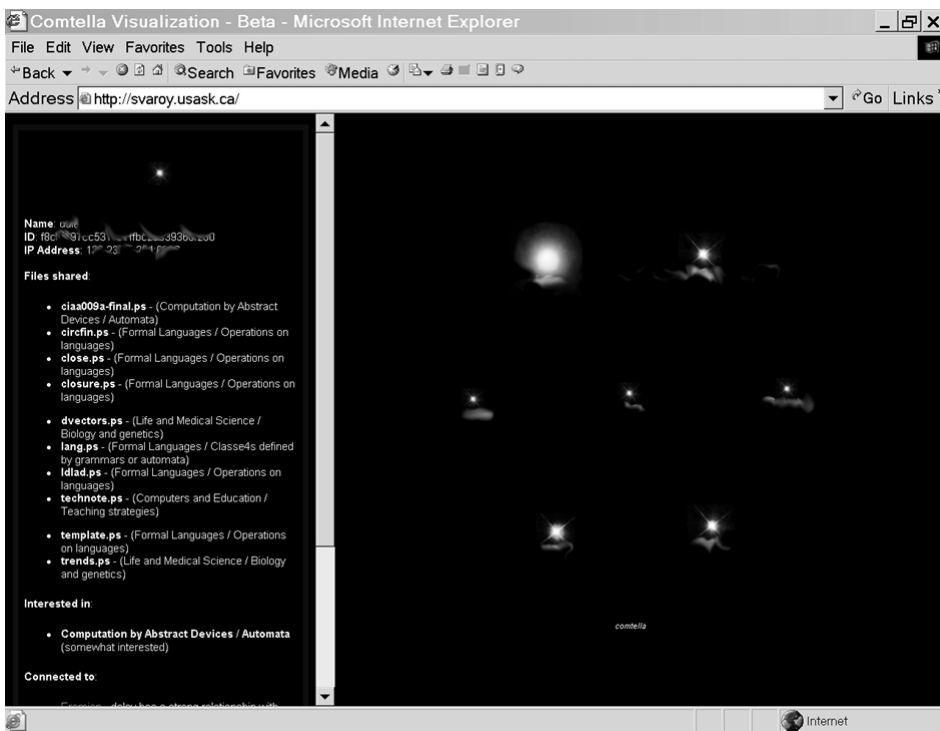


Figure 1. Motivational visualization - fixed version (user names and IP address are blacked out to preserve the privacy of the participants).

was the one who shared more than everyone else and contributed more new things to the community, rather than downloading from others.

When the user clicked on any of the stars, a list of the files shared by the represented user was displayed in the left side of the window, along with the categories of these files (indicating in this way the interests of the user), the number of new files shared by the user and the number of files downloaded from others.

The visualization was a fixed Web page displaying all users who were online at the moment of viewing. It was generated by a specialized visualization node in the Comtella peer network, which received information periodically about the number of shared files from all servents that were on line at the time. In addition to the functionality of the original Gnutella servent, the Comtella servents were modified to keep track of the number of files downloaded or uploaded by each user and to report them periodically (every 30 seconds) to the visualization node. The visualization was generated using off-the-shelf graph-generation software from the data received in the last 10 minutes before the moment a request for a visualization Web page was made by one of the servents. Thus the same view was shown to every user requesting the community visualization at approximately the same time. This design is in line with one of Erickson's (2003) recommendations, stating that customization or adaptive views are undesirable since they prevent a shared social view that is conducive for social comparison.

Case study

The system was advertised for use in the computer science department (about 90 graduate students and 25 faculty) in the fall of 2003. Sixteen users (11 students and 5 faculty) volunteered to use it and participate in the case study, representing several distinct areas of interest: networks, graphics, HCI, AI, software engineering, and distributed systems.

Approximately ten of the participants used Comtella for more than one month. Most of the users accessed the visualization Web page at least once in each session. About 200 unique papers were shared in the Comtella community, but mostly by one user. Another user shared about 20 articles and the rest shared 2 to 5 articles.

There were two main usability problems with the system. First, the hierarchical list of categories (implemented as 3 nested menus) was too complicated and discouraged users from sharing papers (for comparison, it required nearly the same effort as selecting the right subject index category for a paper needed for an ACM publication). Since we had provided an option of "Other" for categories that were not on the list, most users selected this category when sharing papers; very few users took the time to classify the papers appropriately. Since search was organized according to categories (the user has to select the category first to see a list of all articles shared under this category), usually searches for a specific category yielded "No results," which was discouraging.

The second problem was that at any given time there were very few users on line. The graduate students typically activated their Comtella servers only when they wanted to search for something, and switched them off afterwards. Therefore, there were no other peer-nodes available most of the time, and it was hard or impossible to search.

The data collected was so limited that we could not confirm or disprove the hypothesis. However, as a formative study, it was successful. Through interviews with the active users, we obtained feedback about the visualization that revealed the following:

1. All users said that in principle the visualization is a good idea, since it allows to see “who has what” and “who is contributing and who is free-riding.” Yet, there was no consensus on the question, if such visualization could have a strong motivational effect on their own participation.
2. The users typically had a specialized interest in one or two categories, e.g., “distributed systems,” but because of the lack of other users interested in their category, were forced to compare themselves with users with different interests, e.g., in “human computer interaction,” who they did not necessarily care about or want to compare with.
3. The visualization was misleading since it showed only the users who were online at the moment and the sizes of stars were calibrated accordingly. In this way the best user at the moment, even with relatively few contributions, appeared as the sun. This led to inconsistency in time, since for example, a user who saw herself as a large star one day, found that her star had shrunk the next time she logged in, even though she had made several contributions in the meantime, simply because a more active contributor happened to be logged in at this time. This inconsistency was discouraging for some users.
4. Because of visualizing only the users who were online at the moment, often users saw themselves as the only star on the night sky, which amplified a feeling of isolation and lack of community, instead of creating a feeling of co-presence.
5. The visualization was not self-explanatory. While the size was considered by everyone an intuitively clear indicator of the amount of contribution, it was not clear if it shows the number of original contributions, or total shared papers (original papers and papers downloaded from others). It was not easy to see a difference in the sizes of stars with close, but not identical numbers of contributions. The meaning of the color nuance (reddish or yellowish) of a star remained unclear to most participants. The users felt that a legend was needed to explain the meaning of each size and color of a star.
6. The visualization did not provide ways to create customized views, which was considered a weakness by the users. Perhaps due to the fact that the users were computer scientists, they wanted to be in control and select the criteria for viewing rather than having just one view.

7. The graphical location of each star on the screen was random and meaningless, but the users were trying to attach meaning to it, e.g., the proximity of two stars representing the degree of closeness in their interests.

In conclusion, even though there was no evidence that the fixed visualization encouraged social comparison or participation in any way, the users generally liked the idea and indicated in their feedback that an enhanced version of this visualization would be useful in quickly discovering what their colleagues were working on, easily finding out what the hot topics were, and raising the users' awareness of the online community.

Another lesson we learned from this experience was the principal impossibility of sustaining such a small decentralized P2P community covering a wide spectrum of interest categories. When there are only a few people interested in a given topic, the redundancy of the resources across the network is very low. When very few users are online at any given time (which, as mentioned before, was the case in our study), the likelihood that someone else sharing the same interest is online at a given moment is low, and the chance that this user has brought a new paper of interest is even lower.

To allow focusing the study on the effect of the visualization, it would have been better to start with a community which is functioning, i.e., participating already at a certain level. The functionality of sharing academic papers among faculty and graduate students in a department, though deemed useful, was not certain to generate the level of participation required to evaluate the effect of visualization, especially considering the factors listed above (different interests, schedules, and work habits). It would have been better if the users had a common schedule, or were expected to participate in accord, to avoid the problem of too large an asynchronicity, leading to no noticeable effects in the visualization and reducing its potential effect on social comparison.

Therefore, we decided to change the target community to one with a more cohesive scope of interests, where some level of participation is generated in any case, to compare the level of participation of the same group of users under two conditions – a test condition in which the users have access to the visualization, and a control condition where the same users have no access to the visualization. These conditions could be ensured in a community of students taking a class, where participation is expected and rewarded with a certain percentage of the overall grade. An undergraduate class has a fixed and relatively small set of themes or topics, so that sharing resources does not require laborious navigation through a complex subject category index. The students in the class have a common schedule of assignments, projects, exams, and so on, so they are likely to be active at approximately the same time. Yet, to use Comtella as a support tool for a class, it had to be completely reliable, to ensure access to all shared resources at any time to all students. We modified the architecture of the system. We also re-designed the visualization, taking into consideration the user feedback received from the case study with the first version,

even though it was generated by a different community. We considered the feedback useful and general enough to be able to improve the visualization in this new community. The resulting, dynamic design is presented in the next section.

CUSTOMIZABLE VISUALIZATION FOR A CLASS-SUPPORT COMMUNITY

The new customizable visualization was applied in a new version of Comtella supporting a 4th year undergraduate class at the Department of Computer Science. We present first the functionality of the modified system; then we present the modified hypothesis regarding the visualization and the new visualization design, and finally, the results of the new case study.

Community and System Functionality

The new version of the Comtella system supported the students in a senior undergraduate class to share Web resources – articles published in popular Web-magazines, news networks and software – relevant to the weekly topics discussed in the class. “Ethics and Information Technology” is a 4th year undergraduate class required for students taking the B.Sc. Software Engineering Honours Degree in the Computer Science Department at the University of Saskatchewan. It involves discussion of various societal, ethical, and legal issues surrounding information technology. The list of topics includes privacy, freedom of speech, intellectual property, computer crime, wiretapping and encryption, workplace issues (including discrimination, harassment, whistle-blowing, outsourcing), and general social issues (de-skilling, telework, e-commerce and the Wal-Mart effect), as well as professionalism and ethical decision making. A number of the issues, cases, and developments related to these topics are discussed actively in the popular media, and most of the technology-related laws (e.g., those related to patenting software, intellectual property, privacy, freedom of speech, and outsourcing) are still in flux, so there is a need for students to read fresh stories from the media and use them as context and examples for discussing the issues in the class. In previous offerings of the class, as a part of their coursework, the students had been required to create personal Web pages dedicated to the class and to post on them each week links to two articles related to the weekly topic. This activity had been rewarded with 5% participation marks. However, typically the students posted all the links after the classes were over, just before the final exam, and it was not possible to discuss in class any of the stories or facts that were found.

Comtella offered itself as the natural tool to support the convenient and timely submission of links by students. We had to make two modifications to the system to allow for this. First, we had to ensure that all the shared resources were constantly available on the server, independently from whether the users had their clients running or not, to avoid the main problem encountered in the previous version. We solved this problem by artificially centralizing the resources shared by students. This was done by separating

the user interface part of each server from the back-end part of the server (the component that keeps the shared files, forwards and responds to search queries), and placing all the back-ends on a central server, where they were constantly available. When the user “started” her client, in fact, she started just a user interface, which logged the user to the corresponding back-end on the server and allowed her to generate searches herself. Second, to save space on the server and prevent problems with copyright of shared materials, we disabled the functionality of uploading files. Instead, the users uploaded URLs or links to Web-resources (bookmarks), thus creating one of the first handful of tools for social bookmarking, one of the first Web 2.0 applications. Finally, we modified the subject category index menus, to reflect the topics of the class. Thus, we obtained a very simple menu with 10 topics corresponding roughly to the weeks in the class. To maintain a record of the participation of each user, we used a database residing on the same server as the back-ends. The instructions for using the system, as well as the Comtella client, can be found at <http://bistrica.usask.ca/madmuc/Projects/Exp04Jan/comtella.html>.

In the new community, the desirable user actions were defined by us in this order:

- Bringing new resources in the community.
- Downloading new resources from others. While, due to the centralization of the resources, it was not needed anymore to have a big redundancy achieved through downloading and re-sharing, downloading from others was still required to read the papers shared by others (one of the pedagogical goals in using the system).
- Rating and commenting resources. These two actions correspond to the pedagogical goals of encouraging students to read the shared resources and easing the navigation in the shared pool of resources.

Hypothesis and Customizable Visualization Design

When designing a system for real use, the space of possible design decisions is huge. It is impossible to find a theoretical motivation or to evaluate in isolation every possible option. We used as guidance through the huge space of possibilities the feedback from the previous study, even though the community was different and the functionality was modified. The main idea and hypothesis remained the same – that a customizable visualization would stimulate social comparison among users and result in increased participation, which would be measurable quantitatively through appropriate metrics and qualitatively through a user questionnaire.

We expected also that allowing users to customize the visualization by selecting which types of participation they wish to view would allow them to define a dimension of comparison and competition of their liking, and we would see more diversified participation.

The customizable visualization allowed the user to select a view of interest for him/her. The view was generated on the client upon request with the latest data about the community from the server database (to maintain this database, each Comtella server

reported regular participation data to a central server). Although this conflicted with Erickson's (2003) guideline against customization, it was a design decision made in response to the user feedback from the experiment on the previous first design of the visualization. Some users had only one or very few areas of interests, and some users only wanted to compare in one particular aspect such as sharing more files, giving more comments, or rating other users' shared files, but not in all of these activities. Another reason for ignoring Erickson's recommendation was that a customizable visualization would have allowed us to see if users engage in social comparison and competition in the type of activities they care about. Correlating the number of contributions in a given type of participation with the number of accesses to the corresponding view in the visualization could provide a measure for finding out if the respective view in the visualization was able to stimulate social comparison with respect to this type of participation.

The same metaphor (night sky with stars) was used, but highly stylized – instead of stars, the visualization showed rows of nodes (circles) on a black background (see Figure 2). The reason for this was that in the customizable design the user client had to draw every component of the visualization or user request, depending on the selected view. It was easier and faster to draw primitive geometry elements such as a circle and a dot. Generating in real time more complex images would have required more resources and would have taken a longer time, which would have been a deterrent from using the visualization.

All users were shown, regardless whether they were online or off-line at the moment. The goal was to alleviate the problem with the inconsistent size of stars encountered by the users of the first version and to create a stronger feeling of “co-presence” in the system. Each user was represented by a node which was either filled or empty. A filled node represented a user who was online at the moment and an empty node represented an offline user.

In the new design, the users were provided with a simpler graphical language. Color was not used. The new version of the visualization did not represent if the user had taken more from others, or given more to others, since from the user feedback in the first study this turned out to be too subtle a difference, too hard for users to comprehend, to compare themselves with others, and take action on. Instead of using color and size, only the size of stars had a meaning, which depended on the viewing criterion selected by the user when customizing his/her view.

As in the first version, the number of contributions by each user was visualized by the size of the star, since it was found to be an intuitive way of showing the amount of contributions. However, the size represented different types of contributions depending on the viewing criterion chosen by the user.

There were two viewing criteria: the *topic in the class* (selected by the user by choosing from the vertical menu on the left of the screen shown in Figure 2) for which the contribution would be shown and the *type of contribution* that would be shown (by selecting the radio buttons along the upper horizontal part of the screen in Figure 2). Allow-

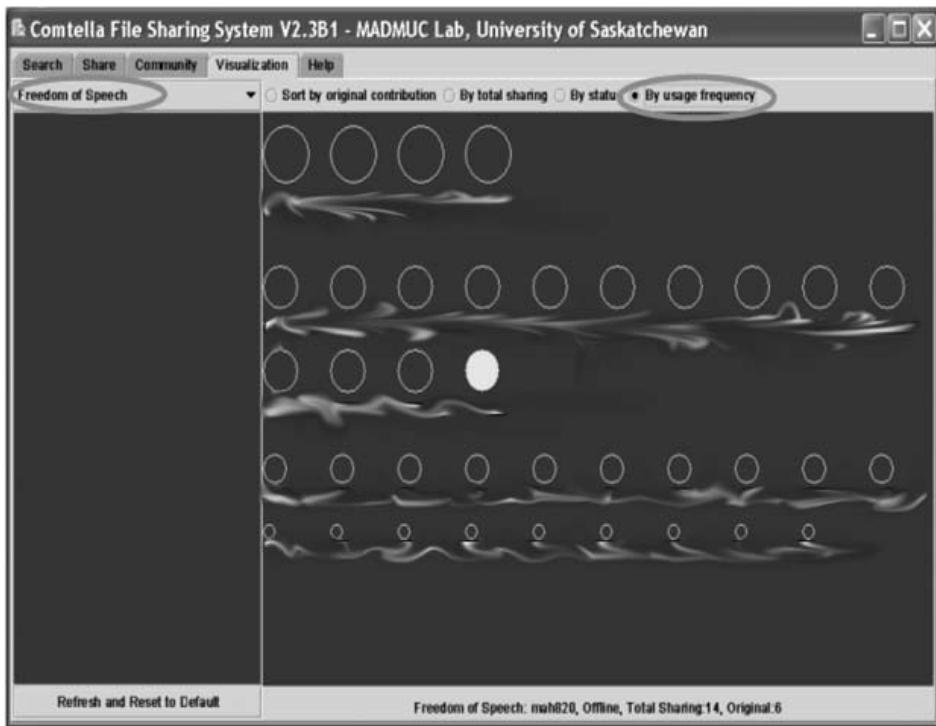


Figure 2. The dynamic design of the visualization (user names blacked out).

ing the users to compare their contributions with those of their peers selectively for different topics in the class was intended to avoid the problem that appeared in the fixed visualization by allowing users who have particular interest in a given topic to compare themselves with similar users. Yet the difference in the communities reduced the severity of the problem. While the members of the community of graduate students and faculty were affiliated in different research groups with different areas of interest, thus forming naturally sub-communities based on areas of interest, the class community did not have subgroups affiliated with particular topics. The users did not select the topics, but rather the selection of topics was external for the entire group, according to the weekly schedule defined by the instructor in the curriculum.

The other viewing criterion was what type of contribution for each topic/week to display: there was a view by number of original contributions (default), view by total shared links (both original and downloaded from others), view by login-frequency, and view by status (membership) which could be “bronze,” “silver” or “gold,” depending on a summative measure of the total participation of the user in the system (presented in Cheng and Vassileva, 2005).

Every user was able to use the same set of selection criteria to view the community, so everyone had an equal opportunity to see each possible customized view, which is somewhat coherent with Erickson’s guidelines (2003). The customization allowed social

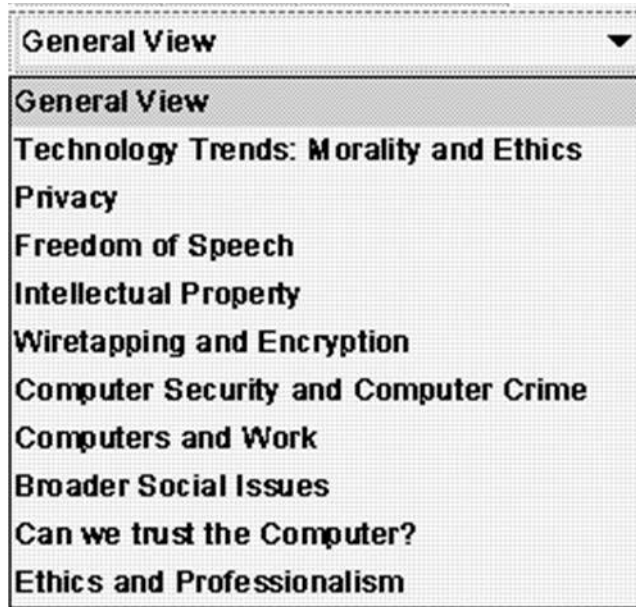


Figure 3. Topic selection box

comparison to take place in a variety of dimensions, depending on how the users wished to define their “peers.”

An important design decision was which view to be shown as the default view. Keeping in mind the requirement that we need to encourage bringing original (new) contributions more than sharing copies downloaded from other users, we decided to use the “original contribution” view as default for the second viewing criterion. For the other criterion – topic of the class (area of interest) – there was an obvious choice for a default – the topic discussed in the current week. However, if the current week was to be set as default, it was necessary to change the default topic automatically in the user interface on a weekly basis, which was very hard to do on the client. It was much easier to get the user to select the topic and then generate the corresponding view. Also it was not clear if the users of the class community will, similar to the community of graduate students in the department, want to compare themselves with others in specific areas of interest, because of the difference in the communities, as discussed above. Arguably, they could have found it more natural to compare in general, for all topics, since all students studied the same topics in the class. For all these reasons, we introduced a “General View” topic (see Figure 3) as a default for the first viewing criterion, which showed all contributions made up to the moment of viewing. The default view was set so that when the user clicked on the visualization, she first saw the “General View,” showing the nodes with sizes reflect-

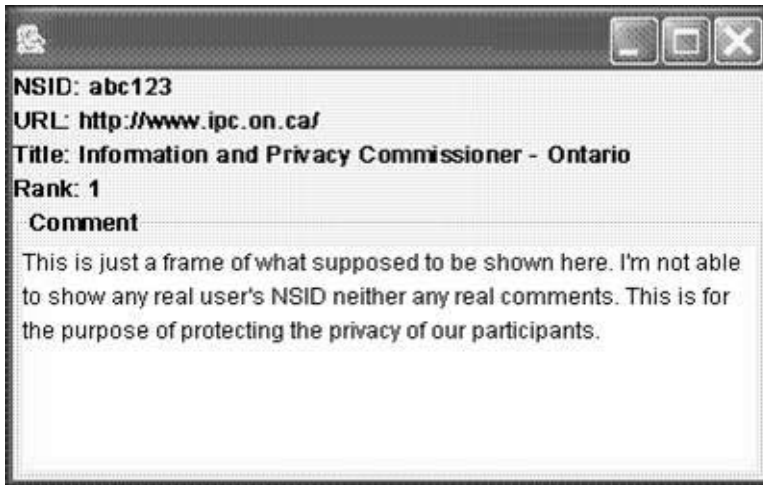


Figure 4. A comment window

ing and sorted by their original contributions for all the topics. The user could also select any of the topics from the menu, if they wanted to view the contributions for this topic.

To facilitate making a distinction between the levels of contribution, the nodes were not randomly located. Instead they were laid out hierarchically at four levels in the descending order of sizes.

The users were able to double-click on a node to see the list of papers shared by this user (the left black part of the window in Figure 2); a single click on any of these files opened up a comment window (see Figure 4). When a user moved his/her mouse over a node in the visualization, a brief summary of the contributions made by the represented user appeared in the bottom bar of the window, on the right side, below the visualization.

The new version of the visualization was implemented as a graph-generating application written in Java and embedded in the Comtella client interface. The visualization was generated on the client upon request with the latest data received from the database on the server. In the next section, we present the evaluation of the social visualization.

CASE STUDY WITH THE CUSTOMIZABLE VISUALIZATION

The case study with the new version of Comtella with the customizable visualization involved 35 students, of which the data of 32 who signed the consent form for participation in the study was used. The class covered 10 topics in 13 weeks. The students could decide if they wanted to use Comtella or not. There was an alternative way of sharing and accessing materials for each week, through personal Web pages, as it was done traditionally in the class before. All students signed the consent form and agreed to have their data used for the study.

The same user group used the system in two conditions: control and test condition,

in which, respectively, the group did not have access and had access to the community visualization. If our hypothesis held, we expected to see an increase in participation measured by the number of original contributions, total shared files (contributions downloaded from others), comments, and ratings under the test condition.

Ten class topics were covered during the time of the case study, where each topic was discussed for one week except the sixth topic, “Computer Crime and Security,” which was discussed for two weeks plus an extra week in-between (reading week); thus this topic ran over three weeks. The system in the control condition was used during the first 6 topics (8 weeks) and the system in test condition – during the remaining 4 topics (4 weeks). To implement this experimental setup, the students first downloaded a Comtella client with an interface that provided all the sharing and searching functionality, but did not have community visualization. After week 8, the subjects were asked to download a new Comtella client which included the visualization option in its interface. Quantitative data reflecting the user participation was collected for each user: numbers of original contributions, total contributions (original files and downloaded files from others), number of logins, and the number of accesses to each of the view provided by the visualization.

A qualitative study of the user experience was conducted in the end of the 13th week to throw light on the motivations of users, and to find out if the students felt that the visualization had encouraged them to compare themselves with others, and to what degree it motivated them to participate. The qualitative study was done through an online questionnaire which the students voluntarily filled after the end of the class. All students were invited to fill in a questionnaire and received a 3% bonus towards their total mark in the class as a reward for the effort. To prevent any influence on the evaluations that the students would give, the students received the same benefit regardless if they participated in the experiment (signed the consent form) or if they used the system. The three students who did not sign the consent form exited the questionnaire immediately and no data was stored in the system about them apart from the fact that they have logged in to fill it, so that they could get the 3% bonus. The class instructor had no access to the students’ data about the experiment and did not know who used the system and who did not. The instructor also did not know who answered the questionnaire and who did not. The answers were anonymized immediately, as well as all the participation data from the database, to mitigate eventual influence on the answers.

Quantitative Results. To evaluate the quantitative effect of the visualization on user participation, we used four participation metrics: (1) the total number of shared articles for each topic, (2) the number of original (new) shared articles for each topic, (3) the number of comments given on the shared articles, and (4) the number of ratings given on the shared articles.

Figure 5 shows the distribution of the total number of contributions made by all users for the 10 topics. One can see that the total number of contributions in the test condition (topics 7 to 10) is significantly higher than in the control condition (topics 1 to 6). In

particular, the contributions for topic 7 are about 3 times higher than for each of the previous topics and the contributions for topic 8 are about 6 times higher. This means that introducing the visualization had a big initial impact on participation according to this metric.

Similarly, Figure 6 shows the distribution of the original contributions made by all users on each topic throughout the experiment. Topic 7 represents the contributions made in the first week under the control condition.

The marked increase in contributions (also in ratings and comments) in the second half of the term is obvious from Table 1. The row marked “before 7” in the table refers to all the participation data of different types obtained during the control condition (topics one to six), and “after 7” refers to all the participation data of different types obtained during the test condition, i.e., the data in all topics after and including topic 7.

The z-ratio of Wilcoxon Signed-Rank test based on each user’s average number of times of logon in first six weeks and in last four weeks was equal to 2.812, which proved that users logged on the system more frequently in the last four weeks than they did in the first six weeks with statistical significance beyond the 0.005 level.

The visualization access data showed that in most of the cases the users resorted to the default sorting criteria (“General view” and “Sorted by Original Contribution”). There were a total of 698 accesses to the visualization (the default view) by 27 users. Fourteen users (44% of the users) continued to access the view “sort by total contribution” in 159 (25%) of the total 698 access cases. The same 14 users continued in 120 (17%) access cases to view the community “by status.” No users ever chose to select the “sort by usage frequency” view, which indicates that it was irrelevant.

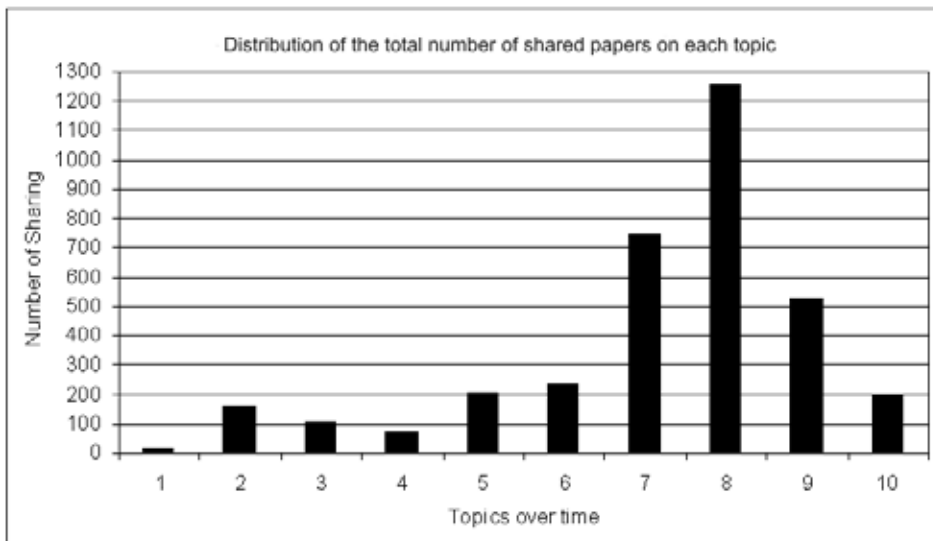


Figure 5. The total number of contributions (original and downloaded from others) made by all users in the group for each topic during the control condition (topics 1 to 6) and the test condition (topics 7 to 10).

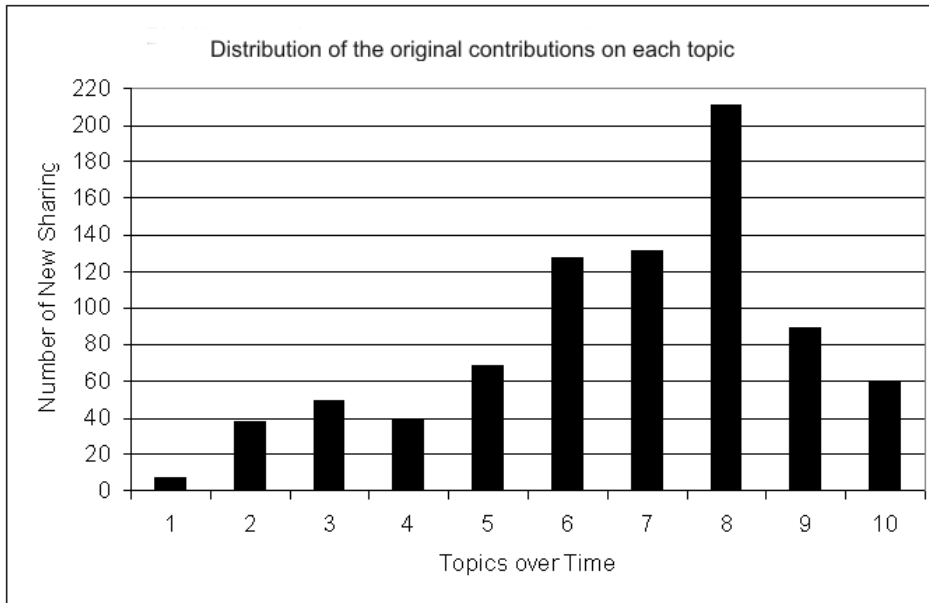


Figure 6. The original contributions made by all subjects for each topic during the control condition (topics 1 to 6) and test condition (topics 7 to 10).

To study the effect of the visualization on the individuals who used it, we used as a metric the correlation between the number of accesses of the visualization by each user and the number of contributions they made in the test condition.

Since the default view showed the users' contributions for the current week sorted by the number of original contributions, we expected to see the strongest correlation between the number of accesses to the default view and the number of original contributions of the users. Each point in Figure 7 represents a user, where the x-coordinate of the point shows the number of times she accessed the visualization and the y-coordinate shows the number of original contributions made by this user. We can see that most points (22) are scattered around the diagonal, with 5 outliers. The correlation is rather high (0.66), which shows that those users that accessed the visualization more often made more original contributions or reverse.

	Total Contribution		Original Contribution		Comments		Ratings	
	number	%	number	%	number	%	number	%
All	3526	100%	821	100%	888	100%	578	100%
before 7	803	22.77%	331	40.32%	176	19.82%	73	12.63%
topic 7	745	21.13%	131	15.96%	162	18.24%	112	19.38%
after 7	2723	77.23%	490	59.68%	712	80.18%	505	87.37%

Table 1. Participation data

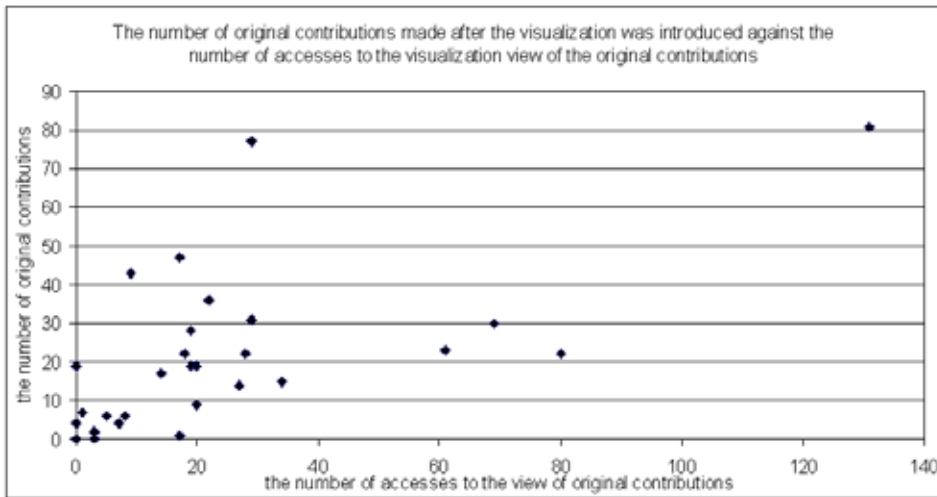


Figure 7. Original Contributions against Usage of the Visualization (correlation 0.66)

The correlation between the number of visualization accesses (the default view) by each user and his/her number of total contributions was 0.41. The correlation between the total contributions of each user against the number of times s/he accessed the view by total contributions in the visualization was lower, 0.34.

The visualization did not provide specific views by number of comments or ratings. We computed the correlation between the number of visualization accesses (to the default view) of each user with the number of comments (0.46) and the number of ratings (0.44) given by that user. The correlation between the number of times the users accessed the status view and the number of their original contributions in the test condition was close to 0 (0.03), but the correlation between their accesses to the status view and their total contributions, comments, and ratings was higher: 0.33 for total contributions; 0.36 for comments; and 0.48 for ratings. The most likely explanation is that users who were interested to compare themselves with others in terms of status thought that the best way to increase their participation was by giving more ratings and comments rather than by contributing more papers or re-sharing papers from others. The relatively higher and similar correlations between the numbers of ratings and comments given by these users with both the number of their accesses to the general view and the number of access to the view by status is that they could not reach the view by status without passing through the default general view, so they accessed both of the views each time.

As Figure 8 shows, as the number of contributions increased after introducing the visualization, their quality deteriorated. Several users found ways to game the system and exaggerate their nodes in order to gain higher status and visibility. For example, two users immediately downloaded a lot of articles from others and re-shared the copies, most

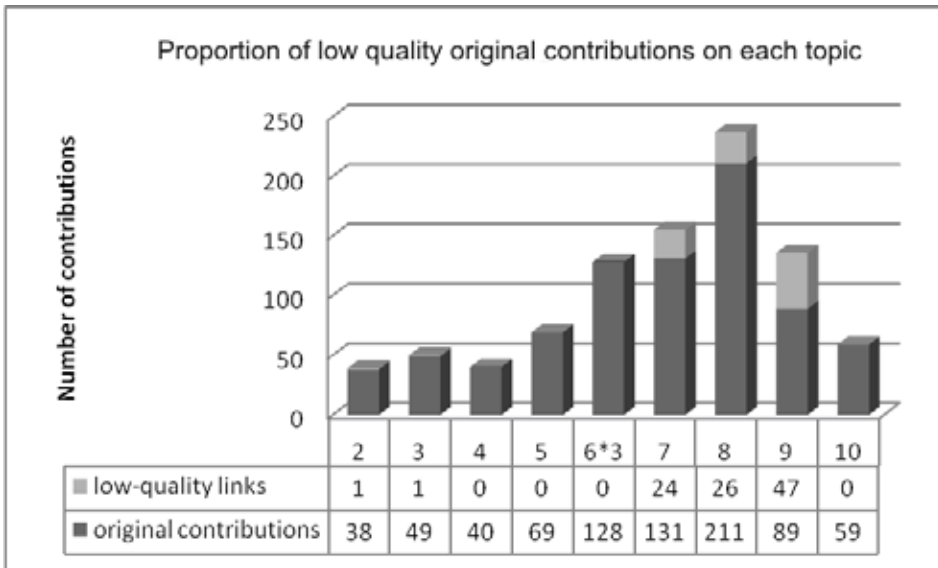


Figure 8. Gaming the system: proportion of low-quality (original) contributions over the topics (the quality and relevance of all contributions was evaluated by the instructor).

likely without ever reading them. These two users shared 570 and 559 articles in total, of which respectively 542 and 531 were shared after the visualization was introduced. Only 16 and 41 of these, respectively, were original contributions; the rest were files downloaded from others.

Qualitative Results. The questionnaire was designed ad hoc and was not based on any existing tool for evaluation of user experience. The reason for this ad-hoc design was that we were not aiming to carry out a formal evaluation of the design, which would have been very hard if not impossible. Our goal was more explorative and formative as we wanted to learn about the users' experience and motivation with the visualization. For example, to check if social comparison was taking place, one of the questions in the questionnaire asked if the user remembered the username of the person who shared the most articles. With some small variance in the ordering, all of the respondents correctly gave the user names of the top 5 contributors.

We tested the questions with a pilot group of graduate students, as well as with some undergraduate students (not members of the class) before we finalized it. There were several types of questions, explained to the users in the beginning of the questionnaire. For some questions, the respondents could select the answer directly from a list of answers provided below the question. For some questions, they had to rank the answers by putting numbers 1, 2, 3, etc., in the box provided for each answer, where 1 was the strongest rank. For other questions, they had to select a number between -2 and +2, where negative numbers were used for a poor evaluation or "disagree," positive numbers

– for a good evaluation or “agree,” and zero meant “neutral,” “don’t know” or “don’t care.” The questionnaire was offered online through the course Web site for a period of two weeks between the end of classes and the final exam.

The questions regarding the visualization together with the results (what percentage of the 32 students chose each answer) are listed below:

The first two questions aimed to find out if the chosen way of visualizing user contributions triggered social comparison and resulting participation actions to improve the user status in the community. Since the questions asked for the user’s perception and possible actions, rather than what the user did in reality (which could be checked from the database where all user actions were recorded), they were phrased conditionally:

1. *What would your reaction be if you saw yourself as one of the smallest nodes in the visualization?*
 - a. *Take immediate action: share more links to make your node larger.* (35%)
 - b. *Think of sharing more links, but later.* (19%)
 - c. *Feel unhappy, but do nothing.* (9%)
 - d. *Feel that the system is unfair, so it doesn’t make sense to contribute.* (9%)
 - e. *Do not care, so will do nothing.* (16%)
 - f. *Other, please specify:* (10%)

In the optional comments given for this question, the students tended towards options *a.* and *d.*, i.e., contribute more papers, try to game the system or expressed concerns about the quality of the shared articles and the ease with which the system could be gamed, as they noticed that not all the users with bigger nodes were making useful contributions. Examples of these comments are listed below (original spelling preserved): “*Creerate the 1000 pound gorilla tactic,*” “*I would feel that some users are “bloating” their ratings by sharing a lot of poor-quality papers, and would try and find more high-quality papers myself,*” “*I would still try to upload good quality links irregardless of quantity,*” “*again, it may affect my final mark,*” “*The quality of a large number of links tend to bad, in particular from those that have shared a large number of articles. It’s very discouraging.*”

2. *If you saw yourself as one of the largest nodes in general, what would you do?*
 - a. *Feel proud of your status and try to contribute even more.* (41%)
 - b. *Feel proud, but at the same time, in some sense ‘exploited’. The others are not bringing in so much, so you will stop or decrease your contributions.* (6%)
 - c. *Feel worried, you may be raising the bar too high and others may hate you or you may be perceived as an ‘overachiever’ by the others.* (6%)
 - d. *Feel nothing, since it is not important to you.* (19%)
 - e. *Other, please specify:* (25%)

The comments to option *e.* agreed with option *a.*, but some added a disclaimer that this would be only the case if they contributed good-quality papers. Here are the comments

(original spelling preserved): “*Feel proud, but i would only contribute the links that I felt needed to be contributed. Other people’s contribution would not matter.*,” “*Keep adding enough to stay at that level.*,” “*the node brings pressure, if i do not contribute any more, my node will be smaller.*,” “*It would probably depend on if I was trying to find quality articles and do a good job, which would make me feel good (and if I got good ratings I would).*,” “*It depends on the quality of the links I had contributed. If they were of high quality, then I would feel proud. If not, then I would feel underserv.*,” “*Feel proud of my status but only contribute more if I have time.*”

The answers to the first two questions indicate that the visualization stimulated or facilitated social comparison and competition among the students.

The next question aimed to check if the users preferred to compare themselves with others with respect to a particular area of interest, as was the case in the previous study.

3. *Would your response to the previous question be different depending on whether the current week is on a topic you are more interested in?*
- a. *Yes (55%)*
 - b. *No (32%)*
 - c. *I don’t know. (13%)*

The fact that half of the users responded positively means that the users did want to compare with others depending on the area of interest, which confirms the findings from the previous case study. However, due to our design, they had to make one extra click to select the view. From the quantitative data we found that in reality less than 50% of the respondents selected a particular topic as a viewing criterion and most of them did this infrequently. The next question aimed at finding what the user perception was of the frequency of their customizing the visualization and using the different views provided.

4. *How often did you use the following visualization options?*

	Often	Seldom	Never
<i>Who is currently logged on.</i>	29%	32%	38%
General view.	58%	29%	12%
View for a given topic.	45%	45%	9%
View by original contributions.	45%	29%	25%
View by total number of shared links.	50%	16%	32%
View by frequency of log in.	22%	29%	48%
View by status (gold, silver, bronze).	45%	16%	38%
View comments from other users.	22%	41%	35%

The answers to this question show that most users correctly estimated their dominant use of the default view for topic (General view). However, half of the users (16) erroneously believed that they often viewed the participation by total contribution. In reality, only 13 users ever accessed the view by total contribution and only 5 of these accessed it more than 10 times over the four weeks. An equally high number of users (14, or

45%) believed that they have used the views of original contributions frequently, by status and by different topics. In reality, the view by original contributions was used most frequently (at any access to the visualization, since it was the default view) by all users. These 14 users did access the views by status and for given topics, but not that frequently. Ten of them accessed the view by status less than 10 times, 1 accessed it 14 times, and only 3 users accessed this view more than 20 times.

The purpose of the next question was to check if the visualization provided information about the other members of the community that the users were interested in. The first three answer options of the question provide additional indication about the degree to which the users were comparing themselves with the others.

5. *What would you like to know about other users?*

	-2	-1	0	1	2
<i>who is online</i>	12%	16%	25%	22%	22%
how much others contributed	3%	6%	12%	35%	41%
<i>are you a freeloader or a contributor</i>	16%	16%	12%	35%	19%
who downloaded from you	3%	9%	12%	35%	38%
who you downloaded from	3%	19%	19%	38%	19%
who gave similar ratings as you	3%	9%	22%	41%	22%
<i>the ratings or summaries of others</i>	3%	6%	29%	45%	16%

To sum up, about 76% of the subjects wanted to know how much others contributed and 54% were interested in knowing if others see them as freeloaders or as active contributors. Many users (73%) were interested to know who downloads their papers, 64% were interested to know who gives similar ratings to their ratings, and 61% wanted to be able to see the ratings and summaries of other users. The conclusion is that the majority of users were interested to know details about the contributions of others and to compare their own contributions to those of their peers.

Another interesting question was what information users were ready to reveal about themselves to others.

6. *What would you like the other users to know about you?*

options	No	Neutral	Yes
<i>Your status level (gold, silver, bronze)</i>	6%	64%	29%
If you are currently online	9%	64%	25%
How many papers you contributed	6%	64%	29%
Who downloads the papers you contributed	6%	51%	41%
Who you download links from	9%	48%	41%
The ratings you give to papers downloaded from others	6%	48%	45%
Your comments on the links you share	3%	35%	61%
Information about my activities should be kept private	31%	54%	12%

Interestingly, users were less willing to divulge information about themselves than to learn about others. For example, while 73% of the users wanted to know who downloads papers from them, only 41% of users would like other users to know who downloaded papers from them.

The next question asked the users to evaluate different aspects of the system.

7. *Please, rank the following features of Comtella on a 5-point scale.*

Feature	-2	-1	0	1	2
<i>Overall</i>	0	6%	22%	64%	6%
As a support tool for the class	0	3%	9%	41%	45%
Usability	6%	19%	38%	32%	3%
Visualization attractiveness	3%	29%	32%	25%	9%
Visualization intuitiveness	6%	12%	35%	38%	6%
Visualization effect on your contribution level	6%	16%	22%	29%	25%
Fairness in ranking your contribution	6%	15%	25%	25%	25%

It seems that the system was viewed positively as a support tool for the class by the majority of the students. Two students (6%) disliked the visualization and responded strongly negatively to all related questions. Fifty four (54%) of the respondents thought that the visualization had an effect on their contribution level; 50% thought that the ranking of their contribution was fair; 44% thought that the visualization was intuitive; and 34% thought it was attractive.

Half (9) of the 18 comments entered by the users referred directly to the visualization. Of these, one-third (6) were positive (the visualization was “*a good idea*,” “*nice*,” “*intuitive and impressive*,” “*perfect*,” that “*it motivates people to submit links*,” and that it was “*exciting to have*”). Two comments explicitly praised the idea of using competition to encourage competition. One student suggested that “*relationships could be shown in the visualization rather than in the search results list*.” Another student suggested adding “*more sorting criteria*.” Three comments were critical about the aesthetics of the visualization (“*kinda bland*,” “*needs to be prettied up*,” “*not as attractive*” as the initially shown in class visualization used in the first version of Comtella). Half of the comments (9) addressed the problem of low quality contributions and that the mechanism for calculating the membership level has to be changed to encourage high quality contributions. One of these 9 students stated that “*Membership levels did not motivate me at all*.” Another one commented “*I didn’t download the latest release so I wouldn’t be tempted to fall into the role of a competitive article submitter. I would rather see quality articles than lots of articles*.”

DISCUSSION

The quantitative results showed a significant increase in group participation when the group was in the test condition with respect to all of the participation metrics introduced.




Figure 9. Relationships between Social Comparison, Participation and Visualization.

77% of the number of total contributions, 60% of the number of original contributions, 80% of the comments, and 87% of the ratings produced by the group of participants were made during the test condition, which lasted only 4 weeks (versus 6 weeks under the control condition). We observed peak values in participation by total contributions in the second week of the test condition which were 6 times higher than the maximum weekly contribution value reached under the control condition (see Figure 5). The maximum group participation value reached for original contributions was nearly 2 times higher than the maximum weekly value reached during the control condition (see Figure 6). Comparing the differences using Wilcoxon's Signed-Rank in individual number of logins during the control and test conditions, we also observed a significant effect of the introduction of the visualization. Logging into the system is also a measure of participation though we did not represent it the visualization.

On the individual level, we found a strong correlation (0.66) between the number of visualization accesses by users and the number of their original contributions. A strong correlation can indicate that either the visualization stimulates users to engage in social comparison which will cause them to participate more. Alternatively, it can indicate that the users that participate more engage in social comparison and use the visualization as a tool for this purpose (see Figure 9).

While there does not seem to be a causal relationship between visualization usage and participation, it seems likely that both of these variables are related through a hidden variable, the degree of social comparison in which users engage. If those users who accessed the visualization more often contributed more than those who did not access the visualization, this is an indication that the visualization stimulates or facilitates social comparison which has been shown in other studies (Shepherd et al., 1996) to be positively related to participation. The stronger correlation between the number of accesses to the view showing the community participation by original contributions with the number of original contributions is an indication that by making certain that a particular view is seen often (for example, by making it a default), one can increase the participation of a particular type.

The answers to the user questionnaire showed that the respondents indeed used the

visualization as a tool for social comparison. About 40% of the respondents of the first two questions agreed with the first answer option which indicated that the visualization was used for social comparison and its effect was increased participation. If we count the part of the respondents who chose option e. but in their comment indicated an agreement with the first option with a disclaimer regarding the quality of contributions, the majority of users acknowledged that the visualization would trigger their participation.

The gaming behavior observed under the test condition also provides evidence that the users engaged in competition, which is a form of social comparison. However, it is not clear if the main motivator in this case was the visualization, the status (bronze, silver or gold membership, that the users earned with their participation), or the combination of both. From the comments given by the students in the questionnaire, it seems that the status was the main motivator for the competition and the gaming. However, without visualization, the users would not have been aware of the status of their peers, and they may not have engaged in this competitive behavior.

The gaming resulted in a high number of duplicates and low-quality resources, which made finding good resources very difficult. Quality control became a major problem after the visualization was introduced and this was reflected in the user comments in the questionnaire. In retrospect, it was an oversight in the visualization design that it did not emphasize the quality of shared resources. In this design the visualization did not encourage social comparison among users with respect to quality and this oversight was exploited by gamers. Since the only possible way to measure the quality of so many resources in a short time is through collaborative ratings, one important lesson for designers is that in addition to the four types of participation that we tried to encourage, another type of participation has to be encouraged - giving ratings.

One point that deserves special discussion is the declining participation in the 3rd and 4th week of the test condition. Such a decline may indicate that the effect of the visualization on user participation may be short-lived and not sustainable. There are five possible explanations:

- 1) The visualization had an effect by chance. The users' different inherent interests in the different topics and the nature and "hotness" of the topics in the real world may be responsible for the increased contributions. For example, it can be argued that the topics 7 and 8 ("Computers and Work" and "Broader Social Issues") were broader and it was easier to find relevant articles for them on the Web. It can be argued that the increase in the contributions, comments, and ratings may be caused by the subjects being more interested in topics discussed in the latter half of the term when the visualization was introduced. The decline in the last two topics was due to lower interest of the students in these topics.
- 2) The novelty effect is well known in the area of Human-Computer Interaction and

may account for the initial interest in the students to use the system with the new interface, which died off as they got familiar with the visualization.

- 3) It is hard to control the “noise” in a real world experiment. There are many factors that play a role in a real classroom. There is an increasing stress to finish the course-work in this and other classes towards the end of term. Most class projects in this and the other classes are due in the end of the term, including a large team-based class project for this class, which is worth 20% of the course grade. This limits the amount of time the students have to spend searching and contributing papers to Comtella.
- 4) Towards the end of the course the “score” had been settled already. Participants had been using the tool for a number of weeks and everyone knew which fellow students were the “high-achievers”. Social comparison was working, but it was not possible to join this “elite” group with a week’s effort. Also, due to the gaming, the “elite” group did not appear attractive anymore.
- 5) The gaming of the system led to too many contributions, a significant proportion of which were of poor quality. This led to overload, frustration, and withdrawal of most of the users from the system.

We will argue that 1) and 2) are not the reasons for the decline in contributions, but instead, the combination of 3), 4) and 5) led to the decline. If 1) and 2) were correct, they would invalidate our hypothesis that the visualization stimulates social comparison and participation, since according to 1) the effect observed was just coincidental – the topics discussed under the test condition just happened to be of greater interest for the students, and according to 2) the effect was due to the novelty of the interface, including the visualization, rather than to the visualization per se, so therefore can never be sustained.

Comparing the participation data from three offerings of the same class using different versions of Comtella can help to see if there is a pattern of high interest in topics 7 and 8 across different offerings and if there is a general pattern of decline in participation towards the end of the class. According to the class instructor, who has taught the class for four years, the topics in the first half of the class, i.e., topics 2 to 6, especially topic 6 “Computer Crime and Security” are typically the most interesting ones for students.

Figure 10 shows the distribution of contributions by topic from three offerings of the class using Comtella, in 2004 (using the system described in this paper), 2005, and 2006. Data from 2003 is available too but it cannot be used for comparison, since the students used personal Web sites to post links and this was not done on a weekly basis, but in accord: all but 2 students contributed all of their links for all topics in the last 10 days before the final exam. One can see that there is no pattern showing many contributions for topics 7 and 8 in any of the other offerings. Explanation 1) is not consistent with the data, so it can be refuted.

One can also see that in all three offerings there is a declining pattern in the num-

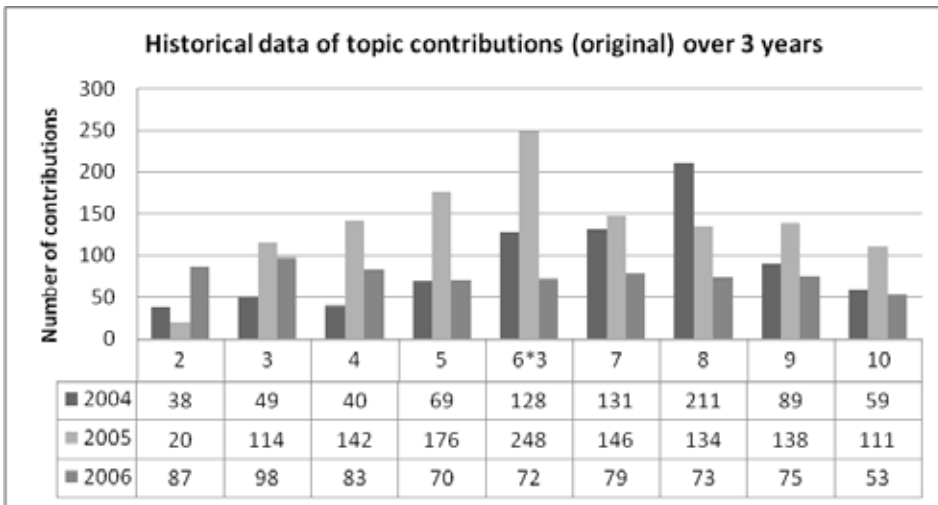


Figure 10. Original contributions for each topic given in three different offerings of the class, in 2004, 2005 and 2006.

ber of contributions for the topics discussed in the second half of the class, with slightly higher values for week 9, which is consistent with 3) and 4). If we exclude the topic 2 which is in the beginning of the class, when the students are still unfamiliar with the system and typically² also share fewer resources than for any other topic, the contributions for the last topic in 2005 and 2006 are the lowest among the contributions for topics 3 to 9 in a given year's offering. However, we observe an exception in the case of the 2004 offering: the number of contributions for the last topic is not only *not* the minimum of all contributions, but it is higher than the contributions of topics 3 and 4, which are, according to the class instructor among the most interesting and resource-rich topics. This is an indication that the visualization had a positive effect on the participation and mitigated a little bit the usual decline in contributions for the last topic.

Unfortunately, the design of our experiment did not allow studying the long-term effect of the visualization on participation, so we cannot exclude completely the novelty effect explanation. To correct this, we designed the case studies differently with the next versions of social visualization. In 2005, we divided the group of users in two groups: one with the test condition and one with the control condition and swapped them after the 6th week (Vassileva and Sun, to appear). In 2006, we divided the group in two – a test and control group with the test group using the visualization consistently for all of the 10 topics (Webster and Vassileva, 2006). In neither case did we observe a novelty effect in the group with the test condition. This leads us to believe that there was no novelty effect in the case study of 2004 as well.

2. The exception in 2006 was due to the fact that the system was introduced early in the week for topic 1 and the students were already familiar with the system when topic 2 was discussed.

The most likely reason for the decline in participation is 5) - the gaming of the system which caused too many contributions and duplicates to flood the system and the resulting decline in quality. This made it hard for the users to find good articles, and their resulting disappointment was reflected in the negative comments in the post-experiment questionnaire. A similar phenomenon called “information overload” (Shenk 1997) has been observed in other online communities. It is typically a result of an influx of many new contributions with medium or low quality, which makes the users feel swamped by a mass of unwanted information. Jones and Rafaeli (1999) found that the users’ most common response to it is to reduce or end their participation in the community, both as contributors and as consumers. Social comparison in such an environment evolves into a senseless competition of out-smarting (gaming) the system, which is repulsive for most of the users. The majority of comments given by the respondents of our questionnaire expressed exactly this frustration and repulsion. Therefore, we believe that the decrease in participation for the last two topics compared to the previous two topics is because of the information overload emerging in the system due to the gaming, aggravated by the general trend seen in Figure 10 for lower contributions towards the end of the class, as a result of external factors: course-work and possibly the lower interestingness of topic or availability of Web-resources, i.e., a combination of 3) and 1). While the external factors cannot be eliminated, to ensure stable user participation in the online community, it is necessary to avoid the information overload by controlling the quantity of user contributions in the system, motivating users to contribute high-quality resources and inhibiting the inferior contributions.

The results of the study were somewhat ambiguous regarding the value of customization. The quantitative experimental results showed that the different views in user-controlled customization tended not to be used, which in retrospect is not surprising, considering that it violated the “minimum user effort” principle in human-computer interaction design. On the other side, the comments of the users, as well as their answers to question 3 suggested that they enjoyed the idea that they can customize their views, and believed that they had used customization more than they did in reality. Also the relatively higher correlation between users’ access to the status view and the number of ratings and comments they gave indicates that the different views can stimulate different kind of participation. While Erickson’s (2003) first guideline for social visualization design states that user-customizable views are not needed, we will be careful generalizing from our experience. Our results do not say anything about how automatic customization of the view would have worked. For example, an adaptive customization that showed a view of the users sorted by different criteria, depending on what category of participation the individual user is currently lagging behind in, may have been motivating and acceptable for users (though special effort would have to be made to avoid confusion in the user about what s/he is currently viewing).

The user feedback suggested that the visualization was neither intuitive nor very at-

tractive. The star metaphor was barely recognizable in the second design. In fact, we are not sure if the graphical representation had any advantage over presenting just a sorted list of contributions for each user or the same information in a tabular form. Future work is needed to compare the effect of visualization versus presenting the same information in a tabular form. Yet analyzing a table may have taken more time and effort and as a result may not have been used so often, which would have reduced the motivational effect. Nevertheless, it seems that to justify the effort of designing a social visualization, the graphical language used has to be richer in meaning, to convey more information, and meet the expectations of the users.

CONCLUSIONS AND FUTURE WORK

There are several general lessons learned from the two design studies presented in the paper. The experimental results and the user feedback received through the questionnaire in the second study clearly showed that the motivational visualization effectively encouraged social comparison and competition, which resulted in increased participation. The implications for system designers are that they can encourage user participation in desired activities by showing a representation of the contributions of the community members along these activities. Care needs to be exercised in the selection of what types of participation need to be encouraged, and how they will be measured and represented, considering that if the visualization encourages social comparison and competition, some users will try to game the system. We cannot recommend or advise against customization based on our experience; however, care needs to be exercised in the selection of the default view in a customizable visualization, since it will likely be the one used by most of the users for social comparison. The selection visualization metaphor and graphical language depends on the kind of community, the knowledge of the users, their general motivation for using the system and the available resources. Size is generally an intuitive representation of participation. Showing all users instead of just those on line was an important improvement. Finding a good way to show who is currently online is important in ensuring awareness of presence at the moment. However, in a purely asynchronous system like the one used in the second study, co-presence is not very relevant for the users. The system used in the first case study, which depended entirely on co-presence, along with other systems that allow chat or synchronous collaboration, will benefit more from visualizing co-presence. All of Erickson's (2003) guidelines have proven quite useful for us.

Motivating social comparison in the quality of the contributions, comments, and ratings became an important future direction of research. The only way to ensure quality control on such a massive scale is through collaborative ratings. However, motivating users to submit ratings is a difficult task. Stimulating competition in the number of given ratings may be hard in a system where their primary form of contribution is resources, not ratings. Therefore, including special and immediate rewards for rating can help to stimulate users to give ratings. This direction of research has been taken in Cheng and Vassileva (2006).

To stimulate users to make high-quality contributions, one has to encourage social comparison and competition in terms of quality. Developing an intuitive visualization of both quality and quantity of contributions is an interesting area addressed by our research (Vassileva and Sun, 2007).

One should not forget that the ultimate purpose of the Comtella system (apart from being an experimental tool for testing motivational approaches for user participation) was to facilitate students in finding and reading fresh materials related to the topics of their studies or research. One student commented that she did not like the competition for bringing resources, since people cared more about finding resources than about reading them. It seems that reading articles found by others is also an important way of participation, even if it is “invisible” from the viewpoint of the community, since it does not contribute to enlarging the pool of shared resources. It is important to remember that the resources are only valuable for the community, if they are *read* by the community members. Otherwise, participation becomes a game with no higher, in this case educational, purpose. Therefore, we should try to find ways of encouraging people to participate even as a “silent audience,” if not as active contributors. One way to encourage such participation, as indicated by the students answers to the question if they want to know who reads the paper contributed by them is through fostering reciprocity in user relationships. Current research in our lab (Webster and Vassileva, 2006) investigates this problem.

ACKNOWLEDGEMENTS

This research was supported by the Natural Sciences and Engineering Research Council of Canada under the Individual Discovery Grant program. The authors wish to thank the two anonymous reviewers for the very thoughtful, detailed, and constructive reviews which helped to significantly improve the paper from its original version. Thanks also go to Prof. Winfried Grassmann and Ran Cheng for their help in the statistical analysis of the results; to Ran Cheng, Weidong Han and Jeremy Long who along with Lingling Sun implemented Comtella; to 16 graduate students and faculty and the 32 undergraduate students who participated in the case studies.

REFERENCES

- Asch, S. E. “Effects of Group Pressure upon the Modification and Distortion of Judgments,” *Groups, Leadership, and Men*, 1951, pp. 177-190.
- Alexander C., Ishikawa S., Silverstein M., Jacobson M., King I. F., and Angel S. *A Pattern Language: Towns, Buildings, Constructions*. Oxford University Press, New York, 1977.
- Bretzke H. and Vassileva J. “Motivating Cooperation on Peer to Peer Networks,” *Proceedings User Modeling UM '03*, LNAI 2702, Springer Verlag: Berlin-Heidelberg 2003. pp.218-227.
- Cheng R., Vassileva, J. “Adaptive Reward Mechanism for Sustainable Online Learning Community,” *Proceedings AI in Education AIED'2005*, Amsterdam, July 18-22, 2005, IOS Press: Amsterdam, pp. 152-159.

- Cheng R. and Vassileva J. "Design and Evaluation of an Adaptive Incentive Mechanism for Sustained Educational Online Communities," *User Modeling and User-Adapted Interaction*, 16 (2/3), 2006, pp 321-348.
- Eick S. G. and Wills G. J. "Navigating Large Networks with Hierarchies," *Proceedings Visualization '93*, IEEE Press, 1993, pp. 204-210.
- Erickson T. "Designing Visualizations of Social Activity: Six Claims," in *Proceedings CHI' 2003* Ft. Lauderdale Florida, ACM 2003, USA.
- Erickson T., Huang W., Danis C., and Kellogg W. A. "A Social Proxy for Distributed Tasks: Design and Evaluation of a Working Prototype," ACM, in *Proceedings CHI 2004*, Vienna, Austria, ACM, 2004.
- Erickson T. and Kellogg W. A. "Social Translucence: Using Minimalist Visualizations of Social Activity to Support Collective Interaction," in *Designing Information Spaces: The Social Navigation Approach* (eds. K. Hook, D. Benyon, A. Munroe), Springer-Verlag: London, 2003, pp. 17-41.
- Festinger L. "Theory of Social Comparison Processes," *Human Relations* (7), 1954, pp. 117-140.
- Garcia S., Tor A. "Rankings and Competition: Social Comparison in the Shadow of Standards," *Social Science Research Network Library*, 2005. (available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=880505#PaperDownload)
- Humphrey M. "Creating Reusable Visualizations with the Relational Visualization Notation," in *Proceedings of 11th IEEE Visualization 2000 Conference (VIS 2000)*, October 08 – 13, 2000, Salt Lake City, UT.
- Jones, Q., and Rafaeli, S. "User Population and User Contributions to Virtual Publics: A Systems Model," *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*, Phoenix, Arizona, 1999, 239-248.
- Lamme V. A. F. "Why Visual Attention and Awareness are Different," *TRENDS in Cognitive Sciences* (7:1), January 2003.
- LimeWire. *The Gnutella Protocol Specification v0.4 – Document Revision 1.2.*: retrieved from http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf
- Pattison T., Vernik R., and Phillips M. "Information Visualization Using Composable Layouts and Visual Sets," *Conferences in Research and Practice in Information Technology*, Vol. 9. Commonwealth of Australia, 2001.
- Shepherd, M.M., Briggs R.O., Reinig B.A., Yen J., and Nunamaker J.F., Jr. "Invoking Social Comparison to Improve Electronic Brainstorming: Beyond Anonymity," *Journal of Management Information Systems*. 12(3), 1996, pp. 155-170.
- Shenk, D. *Data Smog: Surviving the Information Glut*, Harper Collins, New York, 1997.
- Sprenger T. C., Brunella R., and Gross M. H. "H-BLOB: A Hierarchical Visual Clustering Method Using Implicit Surfaces," in *Proceedings of the 11th IEEE Visualization 2000 Conference (VIS 2000)* October 08 – 13, 2000, Salt Lake City, UT.
- Uther J. and Kay J. "VIUM, a Web-Based Visualization of Large User Models," in *Proceedings User Modeling, UM'03*, Springer LNAI 2702, Berlin-Heidelberg, 2003, pp. 198-202.
- Vassileva J. "Motivating Participation in Peer to Peer Communities," in *Proceedings of the Workshop on Emerging Societies of Agents. ESAW'02*, Madrid, 2002. <http://www.ai.univie.ac.at/%7Epaolo/conf/esaw02/preproc/E0029.pdf>.
- Viegas F. B. and Donath J. S. "Chat Circles," in *Proceedings of the SIG CHI conference on Human factors in computing systems*, ACM: Boston, 1999.
- Vassileva J. and Sun L. "An Improved Design and a Case Study of a Social Visualization Encouraging Participation in Online Communities," in *Proceedings 13th International Workshop on Groupware: Design, Implementation, and Use, CRIWG 2007*, Bariloche, Argentina.
- Webster A. and Vassileva J. "Visualizing Personal Relations in Online Communities," in *Proceedings Adaptive Hypermedia and Adaptive Web-Based Systems*, Dublin, Ireland, LNCS 4018Springer Verlag: Berlin-Heidelberg, 2006, pp. 223-233.

APPENDIX 1: Questionnaire used in the study

COMTELLA – 490 USER SURVEY

The goal of this survey is to help us evaluate your experience of the Comtella system. This survey contains 35 questions and will take you about 60 to 90 minutes to complete. It contains two parts: the first one is brief, and after you have completed it, you will be able to proceed with the next one.

There are several types of questions. For some you can select the answer directly from those listed below. For some questions you need to rank the answers by putting numbers 1, 2, 3, etc. in the box provided for each answer. 1 is the strongest rank. For other questions you need to select a number between -2 and +2. Negative numbers are used for a poor evaluation, and Positive numbers are used for a good evaluation. Zero means neutral, “don’t know” or “don’t care”.

Thank you for your participation!

PART I (FOR EVERYONE)

Please, try to quickly answer the questions below to the best of your knowledge (it is not a test).

1. What happens if you click on a link from the search results?
2. Is Comtella centralized or decentralized system?
3. Please, list the bronze search features.
4. Please, list the extra functions a silver member gets.
5. Please, list the extra functions a gold member gets.
6. Who was the user who shared most links in Comtella?
7. What happens if you pass with your mouse over a node in the visualization?
8. What happens if you click a node in the visualization?
9. How many ways are there to view the community?
10. What do you think one needs to do to gain higher membership class in Comtella?
11. What would you rather do now: would you rather go to a party or stay home and watch a movie / TV?

If you have managed to answer at least one of the previous questions, GOTO part II.

If not, click << here – GOTO part III>>.

PART II (FOR PEOPLE WHO HAVE USED COMTELLA)

1. Please, indicate your general reaction to the Comtella system:

(select one number from -2 – very poor to +2 – very good):

-2 -1 0 +1 +2

2. Are you going to use the system again, after the class is over (if it is available)?
(pick one answer only)

- a) Yes
- b) No
- c) Maybe

3. Would you recommend the system to be used in another class?

- a) Yes
- b) No
- c) If you answered “Yes”, could you specify in which class:

4. Why did you use Comtella?

(please, rank at least 3 of the answers below, most important is 1, less important - 2, etc.)

- a) It appears interesting.
- b) To check what my fellow students contribute and compare with them.
- c) To find useful links from other students
- d) To find a good article to write my weekly summary.
- e) Because the professor asked us to use it for the weekly coursework
- f) I am not using Comtella

5. Please rank the following aspects of Comtella in 5-point scale:

(-2 – very poor, +2 – very good)

	-2	-1	0	+1	+2
As a support tool for the 490 class					
Usability (Interface)					
Reliability (crashes etc.)					
Visualization attractiveness					
Visualization usefulness					
Visualization intuitiveness					
Visualization effect on your contribution level					
Fairness in ranking your level of contribution					

6. What would you like to know about the other users? *(Please, answer each question by choosing one number from -2 – strongly disagree to +2 – strongly agree)*

	-2	-1	0	+1	+2
who is currently on line.					
how much the other users contribute.					
who downloads from me the links that I have contributed.					
who I download most links from.					
who gives similar ratings like me.					
if other users perceive me as a freeloader or as an active contributor					

7. Did you prefer to use Comtella in the beginning of the term when there were fewer links available, or later, when there were more (*pick one answer*)?

- a) When there were fewer links, since it later it became very hard to find good links, the number of results was overwhelming
- b) When there were more links, since there was a better choice
- c) The number of links didn't really matter.

8. What would you like the other users to know about you? (*Please, answer each question by choosing one number from -2 – strongly disagree to +2 – strongly agree*)

	-2	-1	0	+1	+2
if you are currently on line.					
how many links you contribute.					
the ratings you give to links downloaded from other users					
your comments on the links you share					

9. How often did you use the following visualization options: (*pick one answer for each feature/row*)

	Often	Seldom	Never
Who is currently logged on			
General view			
View for a given week			
View by original contributions			
View by total number of shared links			
View by frequency of usage			
View by status (gold, silver, bronze)			
View comments of other users			

10. What would be your reaction if you saw from the number of query results that there are already a large number of contributions in a given week (*rank those of the answers that you think apply to you*):

- a) you decide to find some links too to compete with the others
- b) there is enough already, no need to add more
- c) the instructor will not notice your participation, why bother?
- d) try to find some better quality papers to share.
- e) Other: Please, specify:

11. What would be your reaction if you saw yourself as one of the **smallest** nodes in the visualization? (*pick one answer only*):

- a) Take immediate action: share more links to make your node larger.
- b) Think that you should probably share more links, but later.
- c) Feel unhappy but do nothing.

- d) Feel that the system is unfair, so it doesn't make sense to contribute.
- e) Do not care, so will do nothing.
- f) Feel offended - will quit using the system, at least for a while.
- g) Other – please, specify:

12. If you saw yourself as one of the **largest** nodes, would you (*pick one answer only*):

- a) Feel proud of your status and try to contribute even more.
- b) Feel proud, but also in some sense “exploited”, stop bringing more links
- c) Feel worried, you may be raising the bar too high and the others may hate you or you may be perceived as “overachiever” by the others.
- d) Feel nothing, since it is not important for me
- g) Other – please, specify:

13. Would your response to the previous question be different depending on whether the current week is on a topic you are more interested in?

- a) Yes
- b) No
- c) I don't know.

14. Would you prefer to have your servent running on your local machine (as other P2P systems, like KaZaA) rather than having to log in a server?

- a) Yes
- b) No

If you answered “Yes”, please answer also the following two questions:

Do you keep at home a machine that is running constantly?

- a) Yes
- b) No
- c) Does not apply

Would you keep your servent running constantly so that others can download the links that you share?

- a) Yes
- b) No

15. On average, how often did you use Comtella? (*choose one answer only*)

- a) Never, or maybe just once or twice.
- b) Once per week
- c) Twice per week
- d) More then twice and less then 7
- e) Every day of the week.

16. When you are logged on, you are actively using Comtella
 - a) less than 20% of the time
 - b) 20-40% of the time
 - c) 40-60% of the time
 - d) 60-80% of the time
 - e) more than 80% of the time
17. What features would you like to see in the next version of Comtella?
18. Would you like to give any comment or suggestion about the Visualization or about the Membership Level?
19. Would you like to give any comment, criticism or suggestion about Comtella in general?
20. Do you remember the screen name (alias, or user id) of the top users in Comtella?

PART III (FOR USERS WHO COULDN'T REPLY TO ANY OF THE QUESTIONS ON THE PRE-TEST)

1. Can you explain briefly the reasons why you didn't use actively Comtella?
2. Would you like to give any comment or suggestion about what kind of system you would have preferred to use, if any system at all?
3. What features would you like to see in the next version of Comtella?

This is the end of the questionnaire. Thank you very much for taking the time to answer all the questions!

Julita Vassileva, Lingling Sun, Ran Cheng, Weidong Han