

A Study of User Attitude Dynamics in a Computer Game

Yang Cao, Golha Sharifi, Yamini Upadrashta, Julita Vassileva

University of Saskatchewan, Computer Science Department,

57 Campus Drive, Saskatoon, Saskatchewan S7N 5A9, Canada

Keywords: Human Factors, Internet and Collaborative Computing, Agents for Internet Computing

Abstract: When designing a distributed system where a certain level of cooperation among real people is important, for example CSCW systems, systems supporting workflow processes and peer-to-peer (P2P) systems, it is important to study the evolution of relationships among the users. People develop attitudes to other people and reciprocate the attitudes of other people when they are able to observe them. We are interested to find out how the design of the environment, specifically the feedback mechanisms and the visualization may influence this process. For this purpose we designed a web-based multi-player computer game, which requires the players to represent explicitly their attitudes to other players and allows studying the evolution of interpersonal relationships in a group of players. Two versions of the game deploying different visualization techniques were compared with respect to the dynamics of attitude change and type of reactions. The results show that there are strong individual differences in the way people react to success and failure and how they attribute blame and change their attitude to other people involved in the situation. Also the level and way of visualizing the other players' attitude influences significantly the dynamics of attitude change.

1 INTRODUCTION

There are many examples of solid user communities that formed around pieces of technology (e.g. slashdot.com), but there are many more examples of failed ones. Exactly what went right in the thriving communities and what went wrong in the others is difficult to analyze. In our experience developing and deploying I-Help (Greer et al., 2001), a multi-agent environment supporting synchronous and asynchronous peer-help in a University environment, we discovered widely varying levels of user participation in different classes. It seems that not so much technical, but a complex interaction of social factors played a significant role, like rewards (in terms of marks, virtual money or reputation/visibility in the group), attitudes (pre-existing interpersonal relationships among users), and personal beliefs (e.g. altruism). This experience taught us that it is important to study the sociological aspects of cooperation, and that the application should model and support the existing relationships among people, organizational structures (Artikis et al., 2002 ; Sierra and Noriega, 2002) and incentives for cooperative action (Golle et

al., 2001). In the study described here we focus on the following general questions:

- how people develop interpersonal relationships when interacting in a computer-based multi-user environment,
- what is the role of individuality in attributing praise / blame in case of success/ failure, and
- does the design of the environment, especially the feedback given to the user about the other users' attitudes influence the reciprocation of attitudes quantitatively or qualitatively.

A multi-player game environment was designed as a tool to study these questions. It requires the players to represent explicitly their attitudes to the other players and to change their attitudes towards the other players depending on the outcome of the game and their realization of the others' attitudes towards themselves. Different ways of visualizing the others' attitude (text vs. animated face displaying emotion - smiley) were applied in two different versions of the game.

2 RELATED LITARATURE

There are many studies on the evolution of cooperation the area of groupware and CSCW. Methods have been proposed to support and manage collaboration by suggesting appropriate roles, detecting and helping resolve conflicts, and assigning tasks depending on the expertise of the users (Jermann et al., 2001). Enviroments exist that create awareness of the other participants' actions or focus of attention (Gutwin et al., 1995), or study the participation rate and role taking through analysis of the types of speech acts (Soller 2002, Soller et al., 2002) and user actions (Muehlenbrock and Hoppe, 1999), and create models of how these acts relate to effective collaboration and provide guidance about what acitivities the participants should engage in to improve collaboration (Barros and Verdejo, 2000). However, organizational rules alone do not necessarily yield the desired result, as a self-organizing dynamic may appear in the organization which guides the system away from the desired path (Hummel and Schoder, 1995). Such dynamic most often results from personal attitudes and relationships. Most of existing CSCW work is applied to settings where implicit social structures already exist, i.e. the users know each other in advance and have established relationships and status. With the advance of telework environments, there will be an increased need for CSCW environments supporting collaboration between users who have never met face to face and who don't know each other. Building up attitudes and social relationships in such environments happens exclusively during the process of collaboration, mediated through the collaborative environment and can therefore be strongly influenced by the design of the environment.

While attitude formation has been studied in the area of social psychology, the CSCW literature to date has not paid much attention to the fact that people often think in terms of relationships with other people and their attitudes /feelings towards other people govern to a high extent their actions. Attitude formation is a complex process which has been modelled theoretically from different perspectives. For example, the balance theory (Prendinger and Ishizuka, 2002; Rist and Schmitt, 2002), symmetry theory, congruence theory and cognitive dissonance theories take a cognitive stance and explain how people's attitudes towards each other are influenced by their (shared or different) attitudes to important ideas, events or other people. A more pragmatic view is that human attitudes depend on past experience and reciprocation. For example, if somebody has behaved badly towards

another one in the past, it is very likely that the second one will develop a dislike to the first one (without even trying to judge the motives). While such behavior could be modeled theoretically (e.g. the reciprocating "tit-for-tat" strategy in the iterated Prisoner's Dilemma) (Axelrod, 1984) and can be implemented practically with machine learning techniques, attempts to explicitly represent relationships among users have been made only recently. Models of trust updated by reinforcement learning from experience (Yu and Singh, 2002a; Yu and Singh, 2002b) and /or reputation, using other agents as a source of indirect experience (gossip) (Conte and Paolucci, 2002) have been proposed recently in the area of multi-agent systems. These studies have been concerned with the emerging global properties of the system as a result of introducing trust relationships among agents (e.g. what types of equillibria can be reached, how robust is the agent society with respect to "cheaters").

Interpersonal relationships have been studied on a global scale by sociologists. Studies of social networks focus on the patterns of interactions within a group and analyze particular properties of the graph formed by the people (nodes) and their interactions (edges): density, cohesiveness, etc. There have been studies of CSCL environments using social analysis, for example (Nurmela et al., 1999), where the social network cohesiveness of the group is measured to identify the prominent participants in collaboration.

There has been a lot of interest recently in the area of social sciences in general, and particularly in the area of business management in the development of "social capital" in a community or workplace, resulting from positive weak ties (Granovetter, 1973) as a way to promote cooperation, information flow and innovation at the workplace. We believe that building social capital or positive relationships can be an important incentive in CSCW systems where there is no external source of motivation for the users to collaborate (Vassileva, 2002). While introducing currency and micro-payments can help motivate users to help each other (Golle et al., 2001), many users can actually feel repulsed from a money-oriented system (Shirky, 2000); something that we discovered also in our experience with I-Help (Vassileva, 2002). People can be motivated by the possibility to create relationships with other people, and by participating in an active network of relationships, creating thus a small world where recognition and being liked by peers are important factors for the individual (Vassileva, 2002).

We believe that representing and reasoning explicitly about attitudes and relationships among users could be applied in the areas of groupware and CSCW, and this approach can provide a way to

handle emerging self-organizing group dynamics. The design of the rules of interaction in the CSCW can encourage the development of positive attitudes and relationships and increase the motivation for the users to act and to cooperate.

Multi-player computer games provide a good context for exploring emerging social relationships. A Swedish research project on a game called "Kaktus" (Laakolahti and Persson, 2001), allows teenage users to experiment different social behaviours and respond to various social pressures. "Sims Online", a multi-player simulation game allows (according to the advertisement) to: "Build a network of friends to enhance your power, wealth, reputation and social standing." Multi-player action games such as "Dark Age of Camelot" provide a even better ground for studying dynamic social network issues. Recent surveys show that players are troubled by cheaters and saboteurs. The rules of the game (team-based player versus player conflict, no direct communication with the other teams, no ability to switch teams, etc.) set up a situation where given perfect game balance, on any given night, a player may lose 2/3 of his battles. This often leads to frustration and looking for someone to blame. Over time, teams that were intended to be unified against a common enemy end up fragmented into smaller, tighter communities that bicker among themselves, only to reunite eventually and repeat the cycle. Could some subtle alteration of the game rules break this cycle and create large, happy communities? Or do people naturally seek small circles of friends and find reasons to isolate themselves?

We propose a new way of exploring emerging interpersonal relationships in a computer mediated environment by using specially designed multi-player games. In this way we can capture the time evolution of social networks of real people, not artificial agents, as with social simulation. The players form relationships (even though only for a short period of time, in the context of the game) and are more willing to reveal their attitudes to each other in a context of a game than in a real environment. While it can be argued that the context of the game is different than the context of a real world collaboration environment, we believe that most multi-player games reveal individual characteristics of the players that can be seen also in their real-world encounters. The game allows to study the individual differences in the way people change their attitudes, which can help in designing individualized feedback in CSCW environments. The next section describes the design of a web-based multi-player game called "Who likes me".

3 GAME DESIGN

We want to study the evolution of personal relationships among a group of people using a multi-player web-based game. The rules of the game require the users to express and modify explicitly their attitude to the other players as a level of liking or disliking.

In each round of the game a player picks a destination player and has to send him/her a signed packet containing 100 units. However, the packet can't be sent directly to the destination, but by passing to one of the other players (the most liked one). If the selected player likes the originator of the packet, it passes it directly to another player (his/her most liked player), but if s/he doesn't like the originator he/she will take part of the packet proportional to the level of dislike and then pass it further. This process continues until the packet reaches the destination or is destroyed by the other players. After each rounds of the game, the player gets system feedback about what proportion of his packet reached the destination, feedback about the other players' attitudes towards him/her and is able to change his/her attitudes to the other players. After each player completes a given number of rounds (e.g. 10), the one who achieved the highest number of transported successfully units wins the game.

The success of a player in the game is determined by the attitude of the other players to him/her. It is advantageous if the player has a reciprocated positive relationship with at least one other player. However, this is not enough, since if the "friend" of the player passes his/her packet to another one who dislikes him/her, the packet can be destroyed nevertheless. Only through mutual liking and cooperation can all players achieve high scores (though in this case other factors will define who wins the game, e.g. who sends packets faster). However, the uncertainty in the other players' attitude towards oneself and the desire for reciprocation after unsuccessful rounds make the players increase or decrease their level of liking, which makes the game dynamic, unpredictable and interesting. Strategizing successfully in such a complex situation is practically impossible.

3.1 Game Rules

The requirement for the game is that there should be at least three people to play. The game starts by player A signing in the system. Player A will be provided with the list of pseudonyms of the current players and will be required to enter his/her attitude (how much he/she likes each of other player) as a natural number from 5 (strong like) to 1 (strong

dislike). Player A can start to play a round of the game by choosing one of the players as a destination. Player A sends a packet with containing 100 units to destination. The packet continually passes among the group of remaining players, until it reaches the destination (fully or partially) or is destroyed. Each intermediate player, receiving A's package takes away a number of parts proportional to the level of dislike it holds towards A. The round finishes when the packet reaches the destination player or is destroyed. At the end of the round, the player gets feedback about the success of his/her package and feedback generated by the system about the attitudes of the other players towards him/her. We chose to provide only a rough summary of the relations of the other players towards the player, deduced from the observation of how the packet travels and how much it loses. Only summary information "likes" or "dislikes" is presented to the player, but it is not clear if, for example, "dislikes" means 3 or 1. We designed two different ways of presenting the feedback to the user in two different versions of the game – textual and graphical (see Figure 1). After seeing the feedback, the player can change his/her attitudes to any of the other players (if s/he wishes) and play another round.

3.2 Agents Represent Players

Personal agents represent each player in the game, thus saving the player from having to consider individually each passed packet and ensuring consistency in the forwarding of packages according to the attitudes of the user towards the other players. The personal agent maintains a list of attitudes $\{a_1, a_2, \dots, a_k\}$ of the player towards the other k players. A number $a_i \in \{1, 2, 3, 4, 5\}$ where 1 (negative, dislike) to 5 (positive attitude, like) represents each attitude. The player assigns each the value of his/her attitude to each of the other players, thus "instructing" his/her agent how to play the game on her behalf. During the course of the game, the agents decide to whom to pass each packet sent to them and how much to take away from it, depending on the value of the attitude of the user towards the originator of the package. The packet is sent to the agent of the most liked player $M \mid a_M = \max_i \{a_1, a_2, \dots, a_k\}$. If the player dislikes completely the originator R of the package, i.e. $a_R = 1$, the agent will destroy the packet, i.e. it will not pass it further. Otherwise, the agent takes away n parts of the package where $n = 5 - a_R$ and a_R is the value of the attitude of the player to the originator R of the package. The agents do not reveal the attitudes of

their players to either other agents or to the system. In summary, the rules for the agents to play are:

1. To preserve privacy the system is not allowed to access the players' attitudes.
2. The agent that starts the round cannot send its packet directly to the destination.
3. An agent of player A will not send a package to the agent of a player B that A dislikes (i.e. there is $a_B = 1$ in A's attitude model).
4. Each agent of player A selects to pass the package to the agent of the player M to whom the user with the highest attitude value i.e. $M \mid a_M = \max\{a_1, a_2, \dots, a_k\}$.
5. To prevent infinite loops in the game:
 - The agent will not send the packet back to its sender or to the owner of the packet.
 - The agent selects a new agent to send the packet when it receives the packet from the same two previous senders twice.
6. If the player's packet is destroyed, the player's agent will not pass any packet to the first agent that received its previous packet.
7. If the player dislikes everybody (i.e. his/her attitude to every other player is 1), it cannot play.
8. The initial packet for each round of play has a value of 100.
9. The agent who receives the packet will destroy the packet if its player dislikes (at level 1) the packet's owner.
10. The agent who receives the packet will decrement the value of the packet by 5 minus the value of its attitude to the packet's originator.

4 EXPERIMENTAL RESULTS

We carried out experiments with the game to test the following hypotheses:

- Individuals react differently, but consistently to success and failure when changing their attitudes to the other people involved in the situation;
- People reciprocate the attitudes of other people, when they become aware of them;
- The way feedback about other people's attitudes is given plays a role in the way people reciprocate and in the dynamics of the attitudes.

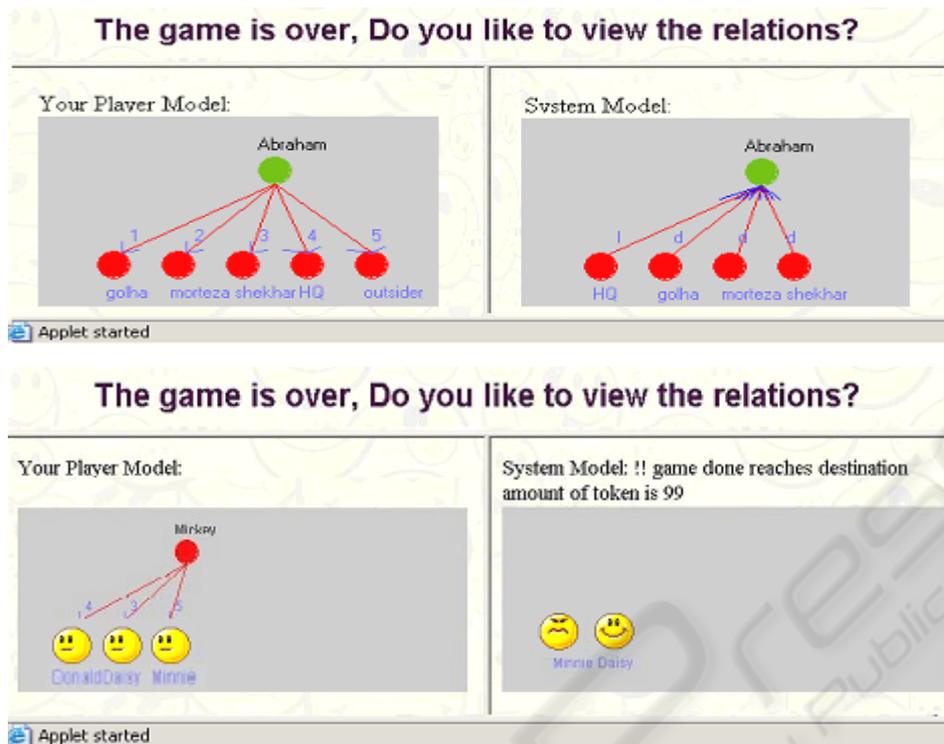


Figure 1: System feedback about the other players' attitudes towards the player (textual and smiley versions)

To test the third hypothesis, we experimented with two versions of the game, one with textual feedback and one with feedback visualized with smileys (shown in Figure 1). The preliminary results generated by two approximately 45-minute experiments with the two versions of the game are summarised below.

Six participants played fifty rounds of the text version of the game in total (i.e. packages sent by different players) and answered survey forms in the end. Seven different participants played fourty rounds each with the smiley version (i.e. seventy rounds in total). The participants had different gender, age, and ethnic background. The group using the smiley version was formed by computer science graduate students, while the group with the text version was mixed. In each set of experiments the participants did not know each other (aliases were used). The players were given a general introduction about the game and the basic rules.

While the rounds of the game were not synchronized across the players, there were five to six players playing at the same time. The routes for a packet to reach its destination were different for the different rounds. The shortest route was a package passed and destroyed by one player and the longest one involved all six players several times and reaching the destination without being destroyed. Three kinds of results were possible in each round:

the packets reached the destination completely; the packets did not reach the destination because they were destroyed, and the packets reached the destination partially.

The following cases reoccurred during the game:

- When everyone in the group strongly disliked the originator, the packet couldn't be send to the destination (direct consequence from rule 8).
- The shortest route of a package happens in two cases: when a sender passes the packet to the most liked other player and that player dislikes strongly the sender, the packet is destroyed immediately. The second case is when the sender selects a player who likes strongly the destination-player, because it passes the packet directly to the destination (see rule 4).
- The longest route happens in a group where no one prefers the destination player to the other players (if the sender likes at least one player and no one dislikes strongly the originator). In this situation, the message is passed continuously in the entire group according to rule 5 until it finally reaches the destination.
- If a sender has good relations with others and s/he selects to pass the package to a player who also has good relations with others, the packet is delivered to destination successfully. However, if the sender selects a player who doesn't like the others, the packet will not be sent to

destination even the sender has good relations with others.

These cases were possible to deduce from the rules of the game and the questionnaires showed that players were aware of them when interpreting the results of each round and were trying to strategize. The only way for players to strategize was by changing their attitudes towards the other players, using information from the system about the success of their package sending at each round and the system-generated record of the like/dislike attitudes of the other players towards them. In the next sections, we present the main observations from the two experiments.

4.1 Setting the Initial Attitudes to the Other Players

Across the two versions, players were fairly consistent in choosing their initial attitude (positive, negative or neutral). In the text version, 45% of all initial attitude choices (i.e. each player's choice of attitude towards each other player) was positive (levels 4 or 5), while in the smiley version, 49% of the initial choices of attitude were positive. In both the text version and smiley version the initial negative attitude selection percentage was 17%. In the text version 38%, and in the smiley version 34%, of the initial choices were neutral. From these numbers it seems that the players had neutral to positive attitude disposition at start. Next we shall see that they were fairly conservative in changing their attitudes.

4.2 Dynamics of Attitude Change

There were 174 opportunities for attitude change in total in the text version and 234 opportunities in the smiley version. The total number of opportunities is calculated as the sum of all feedback stages for each player multiplied by the number of other players at each stage. The breakdown of different scales of attitude changes is consistent across the two versions. Most players keep attitudes to other players constant most of the time – 66.7% and 70% of all opportunities for change of attitude for each player (to all other players at each round) were not used in the text version and the smiley version of the game, respectively. Gradual change with one level of liking/disliking makes 11% of all changes of attitude in the textual version; it is slightly less common (with 4%) than radical change of attitude (with 2 or 3 levels). Gradual change (12% of all changes of attitude) is slightly more common (with 1%) than the radical change in the experiment with

the smileys. Drastic change (from level 5 to level 1 or reverse) makes around 6% in both versions (6.3% in the text feedback version and 6.4% in the smiley version). In the cases when drastic change of attitude took place, it was mostly negative (64% of all drastic changes in the text version and 87% in the smiley version were negative).

4.3 Typical Reactions

One typical reaction is drastically reducing the level of liking to the most liked person after a partial or complete failure to deliver the packet.

This reaction was observed particularly frequently for specific players (e.g. all six drastic changes made by Goofy and all four drastic changes made by Daisy in the smiley version were negative and came in response to partial failure to deliver a packet, see figure 2). Three of the four drastic changes made by Abraham in the text version were of this type (figure 3). HQ had five drastic changes of attitude, two of which were negative and three – positive (figure 3).

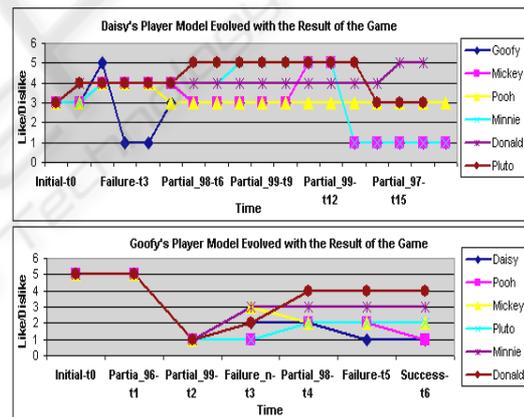


Figure 2: The evolution of attitudes of Goofy and Daisy in the smiley version of the game

Another characteristic reaction was to blame everyone for failing to deliver a packet, as did Abraham in the textual feedback version (figure 3). He reacted to the fact that his package was destroyed by changing his attitude to all other players to “strong dislike” towards the end of the game. After realizing that he will not be able to play anymore, he changed his attitudes to the other players assigning random values. He commented in the questionnaire afterwards that he was annoyed with the other players and didn't know what he should think about them in the end of the game. Two players demonstrated a similar drastic reaction also in the experiment with the smiley feedback – see the evolution of Goofy's and Daisy's attitudes shown in

figure 2. Goofy drastically reducing his attitudes to all players after a partial success and had to increase them again (to randomly chosen levels) to be able to play. Daisy reduced drastically her level of liking to three other players (Goofy, Mickey and Minnie), who were among the four most liked players after a series of consecutive partial deliveries.

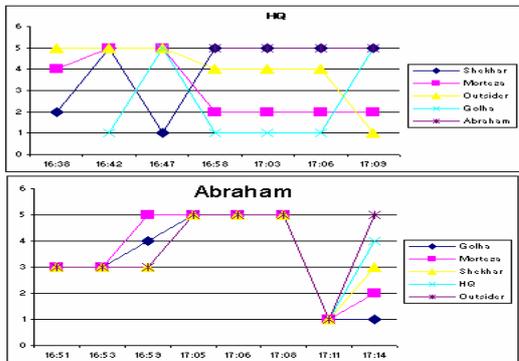


Figure 3: The evolution of HQ's and Abraham's levels of attitude towards the other players (textual version)

4.4 Reciprocation

Comparing the evolution of attitudes of two players towards each other (Figure 4) we see that some of them follow a pattern of reciprocity, delayed with several minutes because of the delay in feedback (only after a round of game the participant can see the system's evaluation of the others' attitudes towards him/her) and the asynchronous rounds across the players. We observed a pronounced difference between the two versions. To measure the reciprocation in attitudes between each couple of players, we mapped the evolution of the mutual attitudes of every pair of players as shown on Figure 4 and counted the changes in the same direction (i.e. converging) over the total number of attitude changes. Applying this measure for each pair of players, we obtained an average of 43.7% (median 50%) reciprocating changes across the players the text feedback version and average of 77% (median 73%) of reciprocating changes in the smiley version. This shows that the smiley feedback visualising the attitude of the the other players stimulates significantly more reciprocation expressed in changing the attitude in the same direction. The reason is probably that the smiley visualization of the attitudes of the other players is more intuitive and requires less cognitive processing, thus allowing a faster, more spontaneous reaction. In contrast, the textual version required more cognitive processing and probably some of the attitudes remained not-noticed by the players, who focussed their attention on one or two other players only.

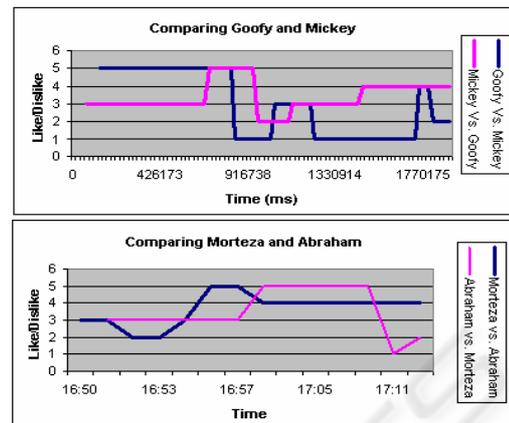


Figure 4: Evolution of the mutual attitudes between two players in the textual and the smiley version.

5 DISCUSSION AND FUTURE

Even though this experiment is too small to bring conclusive results, it gives some evidence in support of our three hypotheses and demonstrates the wealth of data that can be retrieved from a simple multi-player game. Our results indicate that individuality plays an important role in how people change attitudes in response to events resulting from the attitudes of other people. Probably people differ also in the way they assign blame for a situation, which they can not understand because of the complex interaction of the factors involved. One reaction is to blame everyone involved; another – to blame the closest person involved. Such individual differences need to be considered when designing feedback about the actions of other players, for example, providing less feedback for users who tend to react drastically or selecting appropriate visualization to encourage cooperation among users. In the future we will repeat the experiment analyzing the data available to each player at each point when they decide to change attitude and use a think-aloud protocol.

6 CONCLUSIONS

This paper argues for the importance of considering interpersonal relationships emerging among the users of multi-user applications, and for the use of computer games to investigate emerging user attitudes towards each other. Interpersonal relationships among users emerge in any social system, including those mediated by technology, and they play an important role in the patterns of

interaction among people. There are not enough studies of how people actually develop attitudes to each other in the context of a computer supported interaction environments and how these attitudes evolve in time in response to system-mediate events and realizing others' attitude towards oneself. The way the system mediates the user's perception of success and failure, as well as the attitudes of other users influences the way people act. We propose using specifically designed computer games as tools to investigate the dynamics of interpersonal attitudes and we show an example of such a game, together with the initial experimental results. Clearly, more work is needed to generate constructive results to guide system design, and we will be working in cooperation with social psychologists towards this goal.

REFERENCES

- Artikis, A., Pitt, J., Sergot, M., 2002. Animated Specification of Computational Societies. *Proc. Autonomous Agents and Multi-Agent Systems Conference, AAMAS'2002*, ACM Press, 1053-1061.
- Axelrod, R., 1984. *The evolution of cooperation*. New York: Basic Books.
- Barros, B., Verdejo, M. F., 2000. Analyzing student interaction process in order to improve collaboration: the DEGREE approach. *International Journal of AI in Education*, 11, 221-241.
- Conte, R., Paolucci, M., 2002. Reputation in Artificial Societies. *Social Beliefs for Social Order*, Kluwer.
- Golle, Ph., Leyton-Brown, K., Mironov, I., 2001. Incentives for Sharing in Peer-to-Peer Networks. *Proceedings Electronic Commerce EC'01*, October 12-17, 2001, Tampa, Florida, ACM press, 264-267.
- Granovetter, M., 1973. The Strength of Weak Ties. *American Journal of Sociology*, 78, 1360-80.
- Greer, J., McCalla, G., Vassileva, J., Deters, R., Bull, S., Kettel, L., 2001. Lessons Learned in Deploying a Multi-Agent Learning Support System: The I-Help Experience. *Proceedings of AI in Education AIED'01*, San Antonio, IOS Press: Amsterdam, 410-421.
- Gutwin, C., Stark, G., Greenberg, S., 1995. Support for Workspace Awareness in Educational Groupware. In *Proc. CSCL'95, 1st Conference on Computer Supported Collaborative Learning*, 147-156.
- Hummel, T., Schoder, D., 1995. Supporting Lateral Cooperation through CSCW Applications: An Empirically Motivated Explanatory Approach. *Supplement to the Proceedings of the Fourth European Conference on Computer Supported Cooperative Work (ECSCW)*, 11-15 September 1995, Stockholm, Sweden
- Jermann, P., Soller, A., Muehlenbrock, M., 2001. From Mirroring to Guiding: A Review of State of the Art Technology for Supporting Collaborative Learning. *Proceedings of the First European Conference on Computer-Supported Collaborative Learning*, Maastricht, The Netherlands, 324-331.
- Laaksohlahti, J., Persson, P., 2001. "Kactus" project, <http://www.sics.se/humle/projects/Kactus/>
- Muehlebrock, M., Hoppe, H.U., 1999. Computer-Supported Interaction Analysis of Group Problem Solving. *Proceedings Computer Supported Collaborative Learning Conference, CSCL'99*, 398-405.
- Nurmela, K., Lehtinen, E., Palonen, T., 1999. Evaluating CSCL Log Files by Social Network Analysis. In *Proceedings of the Computer Support for Collaborative Learning (CSCL) 1999 Conference*, C. Hoadley & J. Roschelle (Eds.) Dec. 12-15, Stanford University, Palo Alto, California. Mahwah, NJ: Lawrence Erlbaum Associates.
- Prendinger, H., Ishizuka, M., 2002. Evolving Social Relationships with Animate Characters. *Proceedings of Animating Expressive Characters for Social Interactions, AISB'02 Convention*, London, April.
- Rist, T., Schmitt, M., 2002. Avatar Arena: An Attempt to Apply Socio-Physiological Concepts of Cognitive Consistency in Avatar-Avatar Negotiation Scenarios. In *Proceedings of Animating Expressive Characters for Social Interactions, AISB'02 Convention*, London.
- Shirky, C., 2000. *In Praise of Freeloaders*. The O'Reilly Network. Available on line at: http://www.oreillynet.com/pub/a/p2p/2000/12/01/shirky_freeloading.html
- Sierra, C., Noriega, P., 2002. Electronic Institutions: Future Trends and Challenges. *Proceeding Workshop on Cooperative Information Agents CIA'02*, Madrid.
- Soller, A., 2002. Computational Analysis of Knowledge Sharing. In *Collaborative Distance Learning Doctoral Dissertation*. University of Pittsburgh.
- Soller, A., Wiebe, J., Lesgold, A., 2002. A Machine Learning Approach to Assessing Knowledge Sharing during Collaborative Learning Activities. *Proceedings of Computer-Support for Collaborative Learning, CSCL2002*, Boulder, CO, 128-137.
- Vassileva, J., 2002. Motivating Participation in Peer to Peer Communities. In P. Petta, S. Ossowski and F. Zambonelli (eds.) *Proceedings of the Workshop on Emergent Societies in the Agent World, ESAW'02*, Madrid, <http://www.ai.univie.ac.at/%7Epaolo/conf/esaw02/esaw02accpapers.html>
- Yu, B., Singh, M., 2002a. Distributed Reputation Management for Electronic Commerce. *Computational Intelligence*.
- Yu, B., Singh, M., 2002b. Emergence of Agent-Based Referral Networks. *Proceedings of Autonomous Agents and Multi-Agent Systems Conference AAMAS'02*, ACM Press, 1268-1269.