Join the Club: Enabling Self-Organizing Groups on the Net

Stephen Fickas (Computer Science)

Holly Arrow (Psychology)

John Orbell (Political Science)

University of Oregon
Eugene, OR 97403
fickas@cs.uoregon.edu

Abstract

The goal of this project is to establish how social agents can help humans form groups and achieve beneficial outcomes in their task-oriented use of cyberspace. To achieve this goal, the social challenges faced by humans trying to form groups, and the design problems of creating agents that can operate effectively in a mixed human-agent context, must both be addressed.

The approach in the proposed project is to develop task-oriented social agents that combine three abilities: the ability to search out, evaluate, and share useful information that is tailored to the needs of people forming groups; the ability to operate in a mixed setting that includes people and other agents; and the ability to facilitate productive interaction among human members of an on-line group once it has formed. This work will draw on findings from the social sciences to understand the difficulties faced by work groups in general and by computer-mediated groups in particular, and to design social agents that will be accepted by people as specialized members of these groups that can assist them in choosing suitable partners and interacting effectively.

Overall, this project is based on the premise that a new social economics, distinct from the substantially bureaucratic structures characteristic of the industrial revolution, is in the process of developing, and that this will be, in substantial part, web based. One can see the first glimmers of this trend in the growing number of freelance sites springing up on the web. However, for an effective web-based work force to arise that moves beyond single freelancers bidding on smallish contracts or single buyers bidding on small quantities of merchandise, these individual actors must find an effective way to self-organize into collectives: temporary project teams, buying collectives, and virtual companies. When this happens, we will see the emergence of a full-fledged cyberspace economy, in contrast to the more conventional transfer of existing transactions between established industrial age companies and suppliers from phone, fax, and conventional mail onto the Internet.

For this self-organization of individual actors into small, medium, and large commercial collectives to occur, people must deal with the problems of group formation characteristic of face-to-face (FTF) settings, as well as tackling new problems arising from the use of cyberspace. When the field of potential partners is no longer restricted to those with whom one comes into FTF contact, the opportunities are greatly increased, but so is the complexity of finding, evaluating, and negotiating arrangements with potential partners. This project proposes to study how specialized agents can assist individuals to find one other, form a group, and work together effectively. Agents should be able to help people handle the problems of information overload in finding suitable partners, and may also promote positive social dynamics by reminding group members of simple procedures that help in building trust, ensuring cooperation, and facilitating smooth coordination on tasks. Project results will provide a clearer understanding of how to design agents that will facilitate the on-line formation of work groups and organizations in the new social economy of work.

Introduction

The web has opened up new opportunities for groups to form “bottom-up” for their economic gain. However, our laboratory results to date, using face-to-face groups, shows that the removal of a top-down structure may stymie group formation or lead to less than optimal group-outcomes. We believe that software agents may step in to mitigate such problems. In this paper, we describe our work on bottom-up group formation, and the role we see for software agents. We then layout a set of laboratory experiments, drawing on multiple disciplines, that can lead to a solid footing for deploying software agents in new economic settings.

Groups and Agents in Cyberspace

One of the most striking social developments in the on-line world is the spontaneous proliferation of self-organized groups. In the first wave of mass Internet expansion, huge numbers of discussion groups, common interest groups, and social support groups formed using bulletin boards and chat rooms (Kiesler, 1997). The benefits from most of these groups are primarily social; no tangible goods are created or exchanged. As the Internet is increasingly used for mass commerce in addition to socializing, task-oriented and economically oriented exchanges are proliferating as well. Examples include on-line auction sites such as eBay; sites where people seek advice-for-pay from experts, such as expertcentral.com and expertcity.com; sites where specific jobs are contracted out, such as freelance.com; and sites where virtual groups can coordinate their work, such as eproject.com.

In the social setting of chat rooms and bulletin boards, software programs known as bots were developed that interact with humans in recreational social settings (Powers, 1997). In the commercial sector of the Web, personal agents search out and compare product information, helping consumers find and select among products. A third type of agent that has emerged is a specialized program that both processes information for users and facilitates social interaction between people. The programs discussed by Carnevale and Probst (1997), for example, assist negotiators in finding solutions attractive to both parties, and also help structure the interaction between negotiators to improve the chances of a good outcome.

We propose a project that will develop goal oriented social agents that combine aspects of all three agent-types. Like the social "chatterbots," our agents will have the ability to operate in a mixed setting that includes people and other agents. Like the personal Web agents, our agents will be able to search out, evaluate, and share useful information. Rather than being purely personal assistants, however, they will assist groups of people. Like the specialized mediator bots, the agents will be designed to facilitate productive interaction among people -- in this case human members of an on-line group. We draw on findings from the social sciences to understand the challenges faced by work groups in general and by computer-mediated groups in particular, and to design social agents that will be accepted by people as specialized members of these groups that can assist them in completing their tasks. We draw on studies of agents in complex systems that include humans and agents to inform the computer science challenges in designing agents that can operate effectively both in a controlled laboratory environment and on the Internet.

Our focus is on groups that perform tasks that require the pooling of resources by multiple people or agents. We are interested in addressing the difficulties that arise in coordination, trust, and the establishment and enforcement of norms to regulate behavior in groups that emerge from the "bottom up." These self-organized groups (called clubs in economic theory) must develop and enforce behavioral norms on their own to regulate behavior. Problems in this process can be exacerbated in the on-line environment, in which people have reduced social presence and find it more difficult to influence others. We believe that carefully designed agents can help smooth the difficult challenges faced in the formation and coordination of task-oriented virtual groups. In this paper, we provide background for our beliefs, and an experimental strategy for testing our ideas out.

Copyright © 2000, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Using the Social Sciences as a Foundation

In this section, we draw on the social sciences to understand why economically oriented groups form, how they form, and the difficulties they have in coordinating the actions and interests of their members to produce beneficial outcomes.

Coordination of interests and actions in self-organized groups

Social science research has identified many factors that affect whether work groups that meet face to face will be more or less successful in (1) completing their tasks and in (2) establishing and maintaining patterns of interaction that allow the group to resolve conflicts of interest among members and between members and the group as a whole (see Forsyth, 1999; Hackman, 1990; McGrath, 1984). A common problem in groups occurs when members exploit one another rather than collaborating effectively. Research has identified some beneficial social processes that improve teamwork and help groups overcome the problem of selfish behavior. In classic social dilemmas, for example—situations in which cooperative behavior by all members yields the best outcome, but all members have an incentive to behave selfishly--group discussion of the problem followed by universal promises to act cooperatively "solves" the dilemma in up to 85% of individuals (Orbell, Dawes, & van de Kragt, 1990). Reciprocal strategies of rewarding good behavior and punishing bad behavior are also effective in "stamping out" selfish behavior by group members (Axelrod 1984, Trivers 1971).

Research comparing face-to-face (FTF) and computer-mediated (CM) groups suggests that collaborating effectively and solving these social problems can be even more difficult for CM groups. Coordination difficulties are more severe for CM groups, conflict is often higher, and leadership is harder to establish and maintain (McGrath & Hollingshead, 1994). The quality of task performance, as compared to FTF groups, varies, with mixed results depending on the nature of the task.

One major thread of our research is to search for ways that social agents can help people in distributed, computer-mediated groups solve coordination problems by promoting positive social processes such as universal promising. While there is new and exciting work in the area of agent-to-agent interaction, our particular interest is in the means by which a software agent can fit in to a task-oriented group with human members and help these groups complete their tasks. To solve this problem, we must design agents that will be accepted by humans as group members, and that have the capacity to take on specialized task and social roles within a group.

The socialization of group members involves a process of mutual accommodation by which new members and established members adjust to one another as the new members become socialized and fit in to the group (Arrow, 1998; Moreland & Levine, 1988). This process of fitting in is more difficult for members who are demographically different from existing members—for example, for the only male in an otherwise female group (Arrow, 1998). We propose to extend work on socialization, fit, and mutual accommodation to the problem of designing social agents that, despite being radically different from humans in design and abilities, will still be accepted by humans as useful group members.

As demonstrated by the example of bot wars on IRC chat and MOOs (Powers, 1997), rogue agents can seriously disrupt social systems in cyberspace. A more typical problem in cyberspace, however, is misbehaving humans who violate the norms that regulate most people's behavior. A second thread of our research is to design social agents that can assist humans in choosing wisely among potential group members. When people are part of a community that is large enough and fluid enough that they don't know most people very well, information about past behavior of potential partners can be useful in screening out "bad apples." We believe agents that can search out and evaluate information on potential partners can serve as useful advisors to people searching for fellow group members.

Club theory: An Economic point of view

The framework of club theory developed by economists speaks to our concern with how dispersed resources can be organized in an optimal manner for production of some good, and how that can happen short of centralized direction. A club is "...a voluntary group of individuals who derive mutual benefit from sharing one or more of the following: production costs, the members' characteristics, or a good characterized by excludable benefits." (Corrnes & Sandler 1996) Club goods are like private goods insofar as they are subject to exclusion, but they are like public goods insofar as, once produced, they are available to other individuals. Examples range from swimming pools, through groups that simply enjoy each other's company, to groups that are engaged in some kind of productive endeavor (our primary concern).

Club theory assumes that individuals have tastes (in this case for particular club goods, provided at some level of quality) and that they search the market for providers of such goods that are consistent with their own tastes, and with their own willingness and capacity to pay. Distinct from that model, however, the individual's resources are assumed insufficient to provide the good acting alone, meaning that it is necessary to combine his or her resources with those of others in sufficient numbers to ensure that the good is produced.

At this point, the problem becomes the one analyzed by the economic theory of public goods—that individuals who can gain access to enjoyment of the good often have an incentive to free ride by withholding their own contribution. Club theory provides an escape from this problem by the requirement that members contribute (at an agreed on, appropriate level) as a condition of their access to the club good, a criterion that can be met either by an "initiation fee" or by a "use fee" or some combination of the two. Critical to the theory is the individual's freedom to enter (when an appropriate payment is made) and to leave (should, for example, an excess of members—crowding—mean that provision of the good has declined relative to the member's taste). This granted, the economic theory of clubs addresses both how a wider population might end up with a mosaic of clubs (each, presumably, characterized by different quality and price) and, of particular importance to the current project, how particular club goods can be provided at a level of quality and price that suits the tastes of their current membership, and that is not eroded by free riding or overcrowding. Critically, both these things have happened as the product of a "bottom up" sorting of individuals, quite independent of any centralized control or direction. Since its original formulation by Buchanan (1965), club theory has been the subject of considerable formal development (usefully summarized in Cornes and Sandler 1996) but very little systematic, quantitative empirical test.

Club theory assumes that (1) resources (tastes, willingness and capacity to pay) are distributed within a population; but that (2) given tastes, individuals' willingness and capacity to pay is insufficient to produce the good in question alone, thus that resources held by several individuals must be bought together in a productive effort; and that (3) once formed, groups confront a problem of ensuring that, in fact, the constituent individuals do contribute at the levels necessary to provide the good that all want—that is to say, that individuals do not free ride. Club theory also directs attention to particular problems likely to come up in any effort to form productive groups and that might be mitigated by automated 'bots. First: granted the prior specification of a productive task and the components into which it might be broken down, 'bots could facilitate the search process for individuals with the appropriate skills, interests and "willingness to pay" (in terms of their time, and the use of their particular skills). Second, 'bots could facilitate the organization of such "found" individuals into the various productive activities required by the task. And third, 'bots could assist in monitoring and coordination of the various individuals' performance of whatever tasks comprise their contribution to the collective effort—in club terms, their payment of membership fees.

Club theory's efficiency conclusions are, characteristically, based on the end product of the sorting process by individuals into the set of clubs characterizing ("partitioning") the population. But a source of inefficiency is represented by the transaction costs of arriving at that outcome, and any mechanism by which such costs can be reduced represents a contribution to efficiency at both the club and the societal level.

Trust and social dilemmas: A Political Science point of view

"Trust" and "trustworthiness" bear on the success of group formation efforts in two ways. Within a Prisoner's Dilemma framework, a game theoretic perspective predicts that nobody will ever cooperate—the dominant incentive leads each individual to defect, and that everyone should expect everyone else to defect. But social dilemma research shows that people often do cooperate, and further, that they expect (trust) others to do so. Whatever the origin of those expectations, they are an important contributor to persuading the trusting person to cooperate him- or herself. Second, trust influences choices among potential partners prior to such a game actually being joined. As classically modeled, a cooperative enterprise involves the risk of a partner's "cheating" (taking the benefit without contributing to the work), thus there is a premium on selecting "trustworthy" partners for such enterprises—that is to say, partners who will, in fact, contribute to the work.

The extensive tradition of theoretical and empirical investigation into individual behavior in social dilemma (more generally, prisoner's dilemma) circumstances speaks to the role of trust and trustworthiness in resolving these problems. First, trust is sometimes seen as an emergent property of reciprocal relationships that, in themselves, provide private incentive that militate against cheating behavior. Axelrod's (1984) famous analysis of "tit for tat" is based on the assumption that, through iterated relationships, individuals are unable to win by cheating against a player who plays tit for tat, but Axelrod also emphasizes how trust can emerge as a side product of such an iterated sequence played in this manner. Other analyses have focused on the attributes of strangers (in "one shot" relationships) that subjects in laboratory experiments use as a base for trusting (cooperative) behavior (Mulford et al. 1998; Orbell, Dawes, and Schwartz-Shea 1994), but a distinct line of research has asked about the circumstances under which subjects are willing to trust individuals as partners in such games—making a bet, as it were, on their cooperation should the game be consummated (e.g., Orbell and Dawes 1991; Orbell and Dawes 1993). In general, these studies see trust as relevant to the problem of selecting partners for productive enterprises (as modeled in social dilemma terms), but also as relevant to the problem of provoking cooperative, productive contributions to a group effort once the group has been formed.

What are the Research Questions?

Given the previous section as background, our overall goal is to establish how social agents can help humans form groups and achieve beneficial outcomes in their task-oriented use of cyberspace. We are concerned with tasks that require the pooling of resources by multiple people or agents, and with the difficulties in coordination, trust, and the establishment and enforcement of norms to regulate behavior. We assume that the group members themselves must develop and enforce behavioral norms—i.e., their conversation and behavior will not be "policed" by any outsiders, and in particular, will not have a top-down management structure (Fukuyama, 1999). To achieve our goal, we must pay attention to the social problems faced by humans trying to form groups, and to the design problems of creating agents that can operate effectively in a mixed human-agent context. The main issues for humans are:

1. How can a person determine which potential members will be compatible and effective members of a group? This is both a selection issue (who is trustworthy and will contribute toward a beneficial outcome?) and a composition issue (which potential sets—including oneself--have complementary resources and will work effectively together?).

2. Once a good set of potential members has been identified, how can a person help to pull these members together into a functioning group? This involves both communication (contacting potential members about the proposed group) and negotiation (potential members may have offers from others trying to form groups, or may have different preferences re the total group composition). We presume (based on past results) that people will take many different approaches to solving these problems, which will vary in effectiveness.

3. Once self-identified members have formed the group, how can members guide the group to a beneficial outcome? This involves coordination of action (to produce the desired product), coordination of interests (when members have conflicting preferences) and coordination of understandings (about norms, roles, procedures, and standards). It also involves strategies for enforcing norms if members violate them. Although there are many tactics that can help promote beneficial outcomes, group members do not necessarily use them.

We have little information on how the introduction of agents as potential group members will change the dynamics in mixed human-agent groups, but believe the answers will depend on the careful design of agents, based on results of empirical research in the controlled group formation environment we have developed. The primary questions for human-agent interaction are:

1. How will software agents be treated in groups including one or more humans? What attitudes do humans bring to working with a computer-controlled entity in their midst? What trust issues will agents need to overcome to be effective group members?

2. Can software agents not only be accepted, but operate as effective social agents? Can they help guide a group towards a successful outcome by suggesting tactics that help promote effective group processes? For example, universal promising is an effective solution to the kind of conflict of interests among members known as a social dilemma. However, in naturally occurring FTF human groups, this tactic is proposed in only about 50 percent of groups. If software agents were programmed to suggest this approach whenever a social dilemma was detected, would groups accept and adopt this tactic?

3. Assuming that software agents can be effective social agents, can we extend their domain of use to the Internet as a whole? Can they facilitate group formation among people already conducting business on task-oriented web sites, and serve as specialized members of these new groups? How will designs that work in a controlled laboratory site need to be adapted for effective interaction with people on the Internet?

These three questions have lead us to propose a set of laboratory and field experiments. These will be taken up shortly, but first we provide some background research.

Studying club formation in an FTF Setting

In the past two years, Arrow, Orbell, and colleagues have run a series of group formation experiments in a face-to-face (FTF) setting, using a laboratory paradigm for studying self-organized formation of small human groups. The paradigm, called social poker, uses playing cards to represent resources, and provides payoffs for groups who pool their resources to create specific card hands (Arrow, Bennett, Crosson, & Orbell 1999). Each member in a pool of players receives two or three playing cards, and the task of the players is to seek out others with complementary cards and form hands to earn money. In contrast to players of regular poker, no one in the social poker game has enough cards to form a hand by themselves. Players may not trade their resources (cards), but must form a group and pool them. Hands are turned in to earn money for the group, and group members either (in one version of the game) decide together how to divide the money or (in another version) make private claims on these earnings, usually after discussing with other group members what they will do. If, collectively, the private claims exceed what the group has earned, a penalty is subtracted from each claim. If the claims do not exceed the total earned, each member gets what they claimed. After learning what they earned (but not what others in their group claimed or earned), players get new cards and once again seek out others to form a group, which may have the same members or different members from any groups in the previous round.

As people play repeated rounds of the game, they accumulate information on the behavior of others in forming groups, earning money, and proposing and following (or not following) norms for dividing the money. Research using the social poker paradigm allows for the study of individual choices in choosing partners, collective choices in allocating group earnings, and group strategies for solving social dilemmas in the private claims version.

Using the Social Poker paradigm, Ph.D. student Scott Crosson, working in collaboration with Orbell and Arrow, has tested the prediction from club theory that a population with freedom to select partners for mutually profitably clubs will partition itself into a set of clubs that is efficient, not only from the perspective of particular clubs, but also from the perspective of the wider population. In these experiments, which used the private claims version of the game, subjects were free to make claims on the group earnings (the "club good") that might, or might not, exceed their agreed-on share of that good. This overclaiming provision meant that club members confronted prisoner's dilemma logic: Each individual had an incentive to overclaim, but each instance of overclaiming eroded the value of the club good, and decreased each individual's "take home pay" from membership in a club.

The experiment involved three conditions: (1) A single shot game in which subjects played (and knew they were playing) only one round of the social poker game; (2) An iterated sequence, in which subjects in fact played four rounds in sequence, but did not know when the sequence would end; (3) An iterated game (also with four rounds) in which any individuals not included in a successful club would be paid a flat $1.00 fee (a form of "social insurance"). The prediction from club theory, as described above, was broadly supported in both the single shot condition and the social insurance condition; most clubs were "minimal" and thus efficient from both the club's perspective and the perspective of the wider society. In the iterated condition without social insurance, groups commonly included more members than they needed, diluting the share of the good each member received–both because there were more people to divide among and because the incidence of overclaiming in the larger-than-minimum size groups was higher.

The current proposal will extend the laboratory analysis in a manner that informs efforts to use the web as a basis for organizing the efforts of dispersed individuals to engage in the production of club goods. Rapid developments in e-commerce have already advanced market efficiency by vastly reducing transaction costs surrounding the exchange of private goods, yet we have little knowledge about how such reduced transaction costs will affect efforts to organize productive clubs on the web.

Analysis of data from the consensus version of the game, in which players must come to a binding agreement about how to divide up their earnings, illustrates the emergence of "turn-taking" arrangements over multiple round games. Because players must divide whole dollars, a strictly equal division is often not possible–when a 3-person group, for example, earns $10. Some groups resolve this by having members take turns getting the "extra dollar," which presupposes that the same group will reassemble in the next round. It demonstrates one path by which trust is built in self-organized groups. Other groups, while also subscribing to an equality norm, use strategies (such as drawing straws) that equalize opportunity for getting the extra dollar without presuming future interaction. Commitments to reassemble in the same group in the future, however, can decrease the efficiency of groups (with respect to creating the most valuable group goods) if member resources change over time so that they are no longer complementary. Experimental data suggests that while group members will initially reorganize into new configurations when the distribution of resources (cards) changes, they often return to their "old" groups in subsequent rounds, forgoing the extra earnings that new group configurations would offer.

Together, these experiments illustrate the combined operation of both economic and social forces in the formation and membership dynamics of groups. Given the reduced social presence of people in virtual (as compared to the face-to-face) interactions (Wellman, 1997), we are interested in how the balance of forces may be altered on the Internet.

A Step Towards Online Groups

During the past two years, Fickas has supervised the development of a first version collaboration framework to support the formation and interaction of task-oriented groups. This framework, called Confab, allows for distributed interaction among group members, and allows us to replicate the face-to-face social poker experiments online. The interface has been tested extensively with naïve human participants, who report that the current version is intuitive and easy to use, with only a short period of orientation and instruction. In the process of developing and testing the system for human-only social poker experiments, we have gained some preliminary informal information about probable behavior changes based on the reduced social presence of people in the virtual environment. It appears, for example, that people are less likely to include extraneous group members, and more likely to exploit one another, in an online setting as compared to face-to-face group formation.

The three authors have also been working informally for more than a year to explore the use of social agents in online, task-oriented domains. Confab is general in the sense that it supports any mixture of human and agents meeting online to conduct business. It has been used in a graduate seminar to build an agent-only system: all participants were online agents that interacted to form groups and strike deals. Finally, we have begun a proof of concept project to mix humans and agents: we have built an agent called a chaser that will follow humans around attempting to insert itself into a group (somewhat modeled after annoying people we have known).

How to Proceed

Our goal is to establish one or more field studies to evaluate the effectiveness of agents in self-organizing groups/clubs. Our starting point is the results we have obtained from FTF experiments, and the infrastructure we have built to support online groups. Given this, we see the need for a set of three largish studies to fill in the intermediate steps between our current position and our goal.

Replication of FTF results for group formation: First study

Although we have some very preliminary indications of trends in virtual group formation and interaction that differ in the Confab system as compared to FTF interaction, a systematic study is needed to establish similarities and differences in group formation for the two media. Completed FTF experiments, discussed above, provide data on the incidence of suboptimal group formation (groups that earn less than the maximum possible) on member attempts to exploit one another, and on the relative stability of group membership. A partial replication and extension of these experiments in the Confab system will allow us to pinpoint what problems are about the same in the two media, and what problems are worse when interaction is computer-mediated. In all, we predict the need for 450 participants for this experiment, given its multiple cells and iterated play.

Scaling up in complexity: Second study

One of the ways in which the standard social poker paradigm differs from the group formation challenges for people seeking out possible partners on the web is the big difference in the number of potential partners, the number of possible groups, and also the time scale of interaction. We are also interested in the development of trust (and mistrust) over time as people both interact directly and gather indirect information about others. In the second study, we will greatly increase the pool of possible partners and also extend the time scale to allow for information about past behavior of group members to accumulate and be transmitted between participants. These steps will increase the contextual realism of the paradigm in modeling group formation on the Internet, and also will allow for a more detailed study of the evolution of reputation and the impact of reputation on people's willingness to partner with others.

In real life, a huge number of people are theoretically plausible as partners, but in practice, people who form collaborative relationships tend to seek out some subset of plausible partners from their broader social network. At any given time, however, preferred partners may or may not be available. They may be tied up with other projects, out of the country, or on vacation and not checking their e-mail. Thus one may need to seek out new partners for a new project. Instead of a pool of eight participants, from which a maximum of two groups can form, for this study we will recruit a pool of 160 people who will attempt to form groups with one another over multiple sessions, with each session involving a subset of 17 people from this larger pool.

Agents to the Rescue: Third study

In the third study, we will add agents to the mix. The purpose of this study is to determine which components or capabilities of agents will be most helpful in addressing the various problems of (1) failed attempts at group formation (2) suboptimal group formation and (3) overclaiming. The two main ways that agents can help group formation are by helping people identify suitable partners and by providing process information that facilitates cooperative behavior once a group has formed. The two basic dimensions we will vary for the agents are (1) the type of information they provide about other people and (2) whether they use an active or passive strategy in interacting with people.

Agents will all have information about which people are present for a group formation session and what resources (cards) each person has—just as the people do. In addition, all agents will have the ability to assess which groups are feasible for each participant and what the payoff is for each group. Another type of information that may be useful is cumulative information about the past behavior of each person—which groups they were in previously, how much those groups earned, and whether they claimed more than their share of the group resource or not. Human players, of course, have at least potential access to this information through either direct experience or by exchanging information with others. In short, people will gossip, and this will help shape the reputation of each person. However, people do not always remember past experiences or conversations perfectly, they do not necessarily tell all they know, and they sometimes deliberately spread misinformation about others. So we plan to outfit some of our agents with the ability to store and update information about the "record" of each participant in forming and interacting in groups.

Finally, agents may also have information about useful group strategies for solving group process problems. Universal promising is a good strategy for preventing exploitation, for example. Reciprocal strategies are good ways to build relationships and maintain an expectation of future interaction. Speaking up when a group member violates a norm is a good way to enforce good behavior. And so on.

The other dimension we will vary the agents on is passive versus active. Active strategy agents will be "busy-bodies" who volunteer information or suggestions without being cued. The passive strategy is the Star Trek model in which the computer has vast amounts of information but only speaks when spoken to. It seems plausible to us that people may appreciate an active strategy for some kinds of information, but prefer a passive strategy for other kinds of information.

Outcome data will include how much people make use of agents with different designs, people's subjective ratings of what features they found useful, attractive, or annoying, and the incidence of group formation and process problems with agents of different designs present. The data from study two will serve as a baseline for determining whether the agents are making things better or worse.

As with study two, we will recruit a large pool of participants (120-160) and have them interact in sets of 17. We expect to test four to six different bot designs (with different mixes of active and passive strategies and different levels and types of information. Outcome measures such as the rate of group formation, the level of group earnings, and the incidence of overclaiming, will be compared to study two to measure possible improvements attributable to the agents. We will also gather data on people's perceptions of one another and of the different agents. Interaction logs will also reveal qualitative data showing to what extent people made use of the information the agents offered.

Field Studies

Our goal is to reach a point where we can evaluate our ideas in a realistic setting. In this section, we provide a flavor of the type of field studies we would use today. We have chosen two different types of sites to draw out the range of possibilities. The first type of site is one predisposed to have a place for our agents, that of *virtual meeting rooms*. If we were to choose an exemplar site today, www.eproject.com would be a good. It provides a site where meeting rooms can be generated on the fly. Each meeting room supplies the type of virtual tools you would expect for holding online meetings, e.g., document handling and version control, discussion list, messaging, tasks and calendars. Entry to the team/room is by invitation only. We hypothesize that we can use our experimental results to build a social agent for such a site that (1) helps draw and filter members to fill out a team, and (2) helps a team be successful once it has formed. Again using eproject as an example, once the founder(s) of a room/team are established, our agent could search the growing number of freelancer sites on the web for candidate team members compatible with the existing team. Once a team has formed, our agent could monitor the completion of assigned tasks, and work with the team given slipped deadlines.

The second type of site of interest to us is a harder nut to crack*: open software development groups*. We have an interest in the open source movement, i.e., sites that use the Internet to coordinate a group of worldwide programmers working on a common task. We have participated in these type of sites, as well as running one ourselves. We believe the open style has enormous promise as a means of forming bottom-up teams to tackle difficult software problems. We have done some research into the current set-up of open source projects on the net. What we found is that several sites get much publicity, e.g. Linux, Mozilla, Apache. These sites are well established and have a strong group management structure. The remaining sites are often struggling. They have difficulty attracting group members and then difficulty keeping group members on-task and participating. We see these sites as prime candidates for our agents. If we were to choose today, sites such as www.tigirs.com or www.sourcexchange.com would be candidates. Each site has a commercial bent, thus separating itself from the somewhat volunteerism focus of the early founders of the open source movement. Tigris, in particular, uses a mixture of paid and volunteer programmers on its projects; Sourcexchange attempts to attract paid freelance programmers, solely. We would expect that our agents, integrated into sites like these, would play a similar role as in a virtual meeting room, i.e., help recruit team members, help with smooth running of the group. However, open source development sites are not in the team-formation business in the way virtual meeting room sites are. Instead, they are focused on the bottomline of cranking out and testing code. That a *group* of programmers might be needed to do this is not the focus but a necessary evil. We see the insertion of an agent on to such a site and having it succeed as a real challenge. It provides the other end of the spectrum of self-organizing group sites on the web.

Summary

Our work is based on the premise that a new social economics, distinct from the substantially bureaucratic structures characteristic of the industrial revolution, is in the process of developing, and that this will be, in substantial part, web based. One can see the first glimmers of this trend in the growing number of freelance sites springing up on the web. The advantages of the new information economy have had ample press. However, for an effective web-based work force to arise, one that moves beyond single freelancers bidding on smallish contracts or single buyers bidding on small quantities of merchandise, these individual actors must find an effective way to self-organize into collectives—temporary project teams, buying collectives, and virtual companies. When this happens, we will see the emergence of a full-fledged cyberspace economy, in contrast to the less revolutionary transfer of established transactions between established industrial age companies and suppliers from phone, fax, and conventional mail onto the Internet.

For this self-organization of individual actors into small, medium, and large commercial collectives to occur, people must deal with the problems of group formation characteristic of FTF settings, as well as tackling new problems arising from the use of cyberspace. When the field of potential partners is no longer restricted to those with whom one comes into FTF contact, the opportunities are greatly increased, but so is the complexity of finding, evaluating, and negotiating arrangements with potential partners. We are interested in the preconditions necessary for allowing individuals to find each other, to form a group, and finally to work together effectively as a group. We focus in particular on the possible uses of software agents to assist in handling the problems of information overload and promoting positive social dynamics. Our targets are sites where people with specialized skills and complementary work goals can find partners and form collective enterprises.

Related Work

There is a large and growing body of work on software agents. Our interest is in human-to-agent issues, and in particular, the role agents can play in a group that includes humans. There are existing social agents on the web, Yenta and Cobot being two that have made a splash in the popular press (NYT, February 10, 2000). Their approach is to insert an agent into a group of humans and watch what happens. There are clearly insights that can be gained by this seat-of-the-pants style, and we plan to continue to track the projects that use this approach for any results that can inform our lab experiments.

We share many of the scientific goals of the agent simulation community (Sichman et al, 1998) (Luna and Stefansson, 2000). In particular, we are interested in running experiments and then analyzing the event logs that are generated. In our case, experiments may involve all humans, all agents, and various mixtures.

We expect to make use of the agent wrapper technology suggested by (Tambe et al, 2000) when we move from laboratory experiments to field studies.

A review of the club literature has shown that, while there are many case analyses conducted in terms of club theory, there has been only one paper attempting a systematic (Battalio, Kagel, and Phillips 1986) laboratory analysis of various predictions made by club theory; that literature is almost exclusively theoretical. Current work by Crosson, Arrow and Orbell (2000) within the social poker paradigm is addressing this gap, suggesting important qualifications to club theory—notably, with respect to club members' predicted willingness to exclude individuals without the resources to contribute productively to the club good. Our interest is in extending this laboratory analysis, but to do so in a manner that will use the web as a basis for organizing the efforts of dispersed individuals to engage in the production of club goods. Rapid developments in e-commerce have already advanced market efficiency by vastly reducing transaction costs surrounding the exchange of private goods, yet we have little knowledge about how such reduced transaction costs will affect efforts to organize productive clubs on the web.

Although social interaction on the Internet has become a popular topic for research in the past five years (e.g., Kiesler, 1997; Wallace, 1999), studies of group formation have focused on the proliferation of casual discussion groups on USENET (e.g., Baym, 1997), not on project groups, which face much greater challenges for coordination. Some work has appeared on the formation of new relationships (see McKenna & Bargh, 2000, for a recent review of this literature), but the focus is on dyadic matching of romantic partners, not on the formation of work groups. Related to this recreational use of agents is work that uses the Turing Test as a gauge; (Foner 2000) provides an overview. Experimental work on group formation has focused almost entirely on coalition formation. The standard laboratory paradigm (e.g., Kravitz, 1987) allows for a single coalition to form from a small pool; in contrast, our interest is in the formation of multiple clubs from a larger pool of potential group members, and the social poker paradigm is designed to model this situation.

References

Arrow, H. (1998). Standing out and fitting in: Composition effects on newcomer socialization. In M. Neale, E. Mannix & D. H Gruenfeld, (Eds.), *Research on Managing Groups and Teams,* Vol. 1, *Composition* (pp. 59-80). Greenwich, CT: JAI Press.

Arrow, H., Bennett, R. E., Crosson, S., & Orbell, J. (1999). Social poker: A paradigm for studying the formation of self-organized groups. Technical Report 99-01. Eugene, OR: Institute for Cognitive and Decision Sciences, University of Oregon.

Axelrod, Robert. (1984). *The Evolution of Cooperation.* New York: Basic Books.

Baym, N. K. (1997). Interpreting soap operas and creating community: Inside an electronic fan culture. In S. Kiesler (Ed.), *Culture of the Internet* (pp. 103-120) Mahwah, NJ: Lawrence Erlbaum.

Buchanan,James. (1965). An economic theory of clubs. *Econometrica* 32 (February):1 14.

Cornes, Richard, and Todd Sandler. (1996*). The Theory of Externalities, Public Goods and Club Goods*. New York: Cambridge University Press.

Crosson, Scott, Orbell, John & Arrow, Holly. (2000). Voluntary crowding in response to social welfare concerns of club members. In preparation.

Carnevale, P. J., & Probst, T. M. (1997). Conflict on the Internet. In S. Kiesler , (Ed.), *Culture of the Internet* (pp. 233-255). Mahwah, NJ: Lawrence Erlbaum.

Foner, L. (2000). Are we having fun yet? Using social agents in social domains, *Human Cognition and Social Agent Technology*, Dautenhahn (ed), John Benjamins Publishing, 1/2000.

Forsyth, D. R. (1999). *Group dynamics*, 3d ed. Belmont, CA: Wadsworth.

Fukuyama, F., (1999) The Great disruption: human nature and the reconstruction of social order, *Atlantic Monthly*, May, 1999, pp. 55-80.

Hackman, J. R. (Ed.), (1990). Groups that work (and those that don't). San Francisco: Jossey-Bass.

Kiesler, Sara (Ed.) (1997). *Culture of the Internet*. Mahwah, NJ: Lawrence Erlbaum.

Kravitz, D. A. (1987). Size of smallest coalition as a source of power in coalition bargaining. *European Journal of Social Psychology*, 17, 1-21.

Luna, F., Stefansson, B. (2000) Economic Simulations in Swarm: *Agent-Based Modelling*, Kluwer Academic Publishers, 2000

McGrath, J. E. (1984). *Groups: Interaction and performance*. Englewood Cliffs, NJ: Prentice_Hall.

McGrath, J.E., & Hollingshead, A. B. (1994). *Groups interacting with technology: Ideas, issues, evidence, and an agenda*. Newbury Park, CA: Sage.

McKenna, K.Y.A., & Bargh, J. A. (2000). Plan 9 from cyberspace: The implications of the Internet for personality and social psychology. *Personality and Social Psychology Review*, 4, 57-75.

Moreland, R. L., & Levine, J. M. (1988). Group dynamics over time: Development and socialization in small groups. In J. E. McGrath (Ed.), *The social psychology of time: New perspectives* (pp. 151-181). Newbury Park, CA: Sage.

Mulford, Matthew, Orbell, John, Shatto, Catherine, & Stockard, Jean. (1998). Physical attractiveness, opportunity, and success in everyday exchange. *American Journal of Sociology*: 103: pp. 1565-1592.

Orbell, John, Dawes, Robyn, & van de Kragt, Alphons. (1990). The limits of multilateral promising. *Ethics*: 100, pp. 616-627.

Orbell, John, & Dawes, Robyn. (1991). A 'Cognitive Miser' theory of cooperators' advantage. *American Political Science Review* 85: 515-528

Orbell, John, & Dawes, Robyn. (1993). Social welfare, cooperators' advantage and the option of not playing the game. *American Sociological Review* 58: 787-800

Orbell, John, Dawes, Robyn, & Schwartz-Shea, Peregrine. (1994). Trust, social categories, and individuals: The case of gender. *Motivation and Emotion*, 18, pp. 109-128.

Powers, James. (1997). *Irc & Online Chat*, Abacus Software, September, 1997.

Sichman, J., Conte, R., Gilbert, N., (1998), Multi-Agent Systems and Agent-Based Simulation, *Lecture Notes in Artificial Intelligence* 1534, Springer, 1998

Tambe, M., Pynadath, D. and Chauvat, N. (2000) Building dynamic agent organizations in cyberspace, *IEEE Internet Computing* (to appear)

Trivers, Robert. (1971). The Evolution of reciprocal altruism. *Quarterly Review of Biology*: 46, pp. 35-57.

Wallace, P. (1999). *The psychology of the Internet*. Cambridge: Cambridge University Press.

Wellman, B. (1997). An electronic group is virtually a social network. In S. Kiesler (Ed.), *Culture of the Internet* (pp. 179-208). Mahwah, NJ: Lawrence Erlbaum.