An Agent Framework to Support Opportunistic Collaboration

Melfry Moreno[1], Adriana Vivacqua[1], Jano M. de Souza[1,2]

[1] COPPE/UFRJ – Computer Science Department, Graduate School of Engineering
[2] Institute of Mathematics - Federal University of Rio de Janeiro
PO Box 68513, Zip Code 21945-970, Rio de Janeiro, RJ, Brazil
{melfry, avivacqua, jano}@cos.ufrj.br

Abstract. In increasingly networked times, CSCW systems have become more common. In most of these, individuals work collaboratively from their personal computer terminals, unaware of their peers. In this scenario, opportunities for collaboration often go unnoticed. In this paper, we investigate aspects of unplanned cooperation and how it might be encouraged. We propose an agent framework to encourage and support unplanned cooperation between people not necessarily in the same team. Agents build user profiles and match users according to their interests, activities and opportunities for cooperation. By matching users' work contexts, needs and resources, we expect to uncover opportunities for collaboration that might otherwise go unnoticed. We believe that the notification of users of these opportunities will lead to more frequent collaboration between them.

1 Introduction

As more people and organizations become connected and cooperative work tools become commonplace, it becomes more common to find people working together in virtual environments. Most environments allow for message and file exchanges, discussions and co-editing. Some of these environments are media rich, including audio and video interaction in addition to standard tools. However, certain opportunities for interaction are lost in these environments, such as the informal hallway conversations and suggestions that may influence one's line of thought or work. When at the computer, a person's awareness of the environment is severely limited and the absence of environmental information often presents a setback.

Instant Messaging systems and their awareness tools have started changing that somewhat: people can now be aware of others they know who are online at the moment and contact them should they need assistance. However, excessive messaging can be disruptive in a work environment, as has been pointed out in [13]: to a large extent, information about the activities of others is irrelevant in the current working context and only hinders work. We believe more can be done to jump start collaboration. Through user profiling, interest and expertise management and context awareness, it might be possible to better leverage users' skills, competencies and available time and induce cooperative work. This is especially true in unstructured or

loosely structured work environments, where work groups and teams are highly reconfigurable and not necessarily pre-set from the start. The academic environment is one such example: research teams may be engaged in different lines of work and specialists may join the group and contribute at different points. They may work as temporary additions to the group (with the objective of solving a particular problem, for instance) or they may become permanently involved with the project as a whole.

Many opportunities for cooperation are wasted for lack of awareness that the opportunity even exists. Individuals don't know of others skills, interests, availability or willingness to collaborate in a project. Furthermore, it's not as easy to get to know and trust someone in the virtual world, or to bump into someone you know and might be able to collaborate with. Some systems allow for informal interaction in the hopes to stimulate conversation and creativity in their members, with various results.

The extra information provided by awareness systems leads to a need for careful control of information flow, so as not to disrupt users. The issues of what information to present, when, how and to whom have grown in importance. To address these issues, we turn to profiling and matchmaking techniques, in order to filter the amount of information to be provided and the moment and recipients of the information. Another issue invariably associated with the provision of awareness information is privacy violation: a user's privacy may be violated by making details of their activities available that should have been kept private. Every piece of information about a user that is made available to others is a potential privacy violation.

In this paper we propose an agent-based framework to support awareness and discovery of potential collaboration opportunities. We begin by introducing some background work (section 2), then move on to the framework in section 3, and wrap up in section 4.

2 Background Work

Several previous works have dealt with video interfaces and using video to support personal awareness and informal interactions. One such example is VideoWindow [9]. Piazza [6] enables people to be aware of others who are working on similar tasks when using their computers, thereby enabling unintended interactions. It also supports intentional contacts and planned meetings. PIÑAS [11] is a platform that provides potential and actual collaboration spaces, as well as specific services customized to support collaborative writing on the Web (Doc2U [11]).

2.1 Awareness

Awareness has received much attention from the CSCW community in the past few years, as researchers realize the importance of being aware different aspects of the environment in a collaborative work setting. For instance, Pinheiro et al. propose a framework for past event awareness, in which users become informed of past occurrences, results and work history (including evolution of shared data, members' actions and decisions, etc.), so as to better collaborate in the present [12]. Other authors have proposed prospect awareness systems, to enable individuals to envision

potential benefits of collaboration in an attempt to motivate collaboration [5]. Many recent papers have addressed awareness in mobile computing environments, where location awareness is a central issue for collaboration [8]. Other research has focused on document or task based awareness, and providing information to users about who is working on the same document or performing similar tasks at a given moment [6, 10]. The most basic form of awareness is personal awareness, such as provided by current messenger systems, where a user specifies a list of contacts and the system displays their availability and status.

We believe the first step towards a successful collaboration is becoming aware of the opportunity to collaborate. We therefore focus on potential collaboration awareness, and aim to provide users' with information on opportunities for collaboration, a similar problem to that addressed in [11].

2.2 Unplanned Interactions

Kraut [9] presents a useful classification of the different types of interaction found in work environments: (1) *Scheduled*: conversations previously scheduled or arranged; (2) *Intended*: one where the initiator set out specifically to visit another party; (3) *Opportunistic*: one in which the initiator had planned to talk to other participants sometime and took advantage of a chance encounter to have the conversation; (4) *Spontaneous*: a spontaneous interaction in which the initiator had not planned to talk with other participants.

He also points out that the majority of conversations in organizations are informal in nature and that these are usually short, involve two people and build upon previous discussions. They tend to happen because one person happens to be near another at a time when one wants to ask for or provide information. Studies show that these informal interactions play a central role in helping workers learn, understand, adapt and apply formal procedures and processes [6]. Few systems have focused on support for opportunistic and spontaneous interactions.

In [10], virtual proximity is defined as situations in which users access the same data or users invoke the same application in the virtual environment. We take a similar approach, using an individual's current context (what one is currently working on) to search for others who might be interesting to talk to, within that context. Our system also takes the user's context into account, providing only information that is relevant to the user at the moment, so as to not worsen the problem of information overload or disrupt the user's flow of work or line of thought.

3 CUMBIA Framework

We have envisioned an agent-supported peer-to-peer architecture, CUMBIA, where each user has a cluster of agents to help with knowledge management and collaboration tasks. These agents are in charge of identifying potential cooperation situations and trying to make these come to fruition. Agent-oriented techniques are finding increased usage in a range of telecommunication, commercial, and industrial applications [7], as developers and designers realize it's potential. Agents are

especially suited to the construction of complex peer to peer systems, because they are lightweight, permit parallelization and easy reconfiguration of the system. Agents have been used in groupware for a long time due to their social abilities [2]. A recent survey of the application of agents in groupware and CSCW can be found in [3]. Systems such as Personal Assistant [4] and COLLABORATOR [2] are some successful examples of agent approaches used in developing collaborative tools. AwServer, and E_Places are good examples of agent-based awareness work [1].

3.1 Agent Architecture

In CUMBIA, each user has its own agency to assist with knowledge management and cooperation tasks. In this architecture, there are four agent teams to perform specific tasks as shown in Figure 1. Agent Teams and main functionalities are:

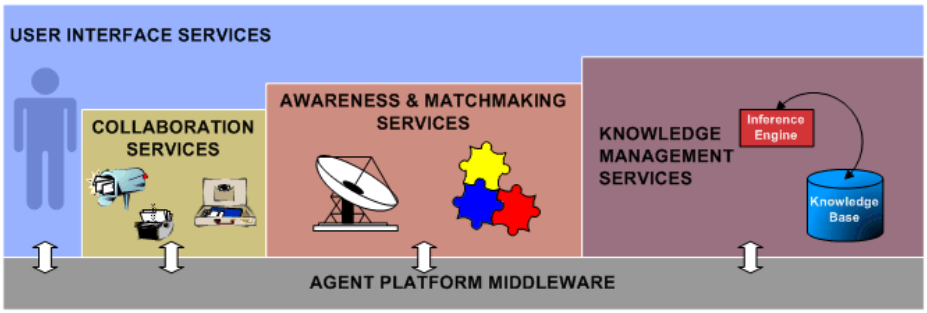

Fig. 1. CUMBIA Agent Architecture

- **User Interface**: displays information to the user and allowing the user to specify parameters and information to the other Agent Teams.
- **Collaboration**: allow for the easy and quick establishment of contact when the possibility for collaboration arises, provide tools for cooperation (forums, etc.)
- **Awareness and Matchmaking**: search for other users with whom it might be interesting to establish contact, contact other agents for their users' profiles and contexts, compare user profiles to current context and work environment.
- **Knowledge Management**: manage user's personal data and build initial profiles based upon it, keep track of documents, searches, ongoing collaborations and current research.

User interface agents perform all interface related functions, providing information to the users and requesting information from them. UI agents mediate requests between agents and users. We have not concentrated on the best ways to display information at this point, but agents will be able to decide on the more appropriate ways of displaying information when they receive it and decide on the proper time to display it, thus addressing the problems of what, how and when (given that who is fixed). We will be building upon work done in [1]. Knowledge Management, Awareness and Matchmaking and Collaboration Services are discussed in more detail next.

3.2 Knowledge Management Services

KM agents provide profiling functionality to the system. Basic units in the profile are projects and interests, which are interrelated. Part of the profile information has to be explicitly provided, and another part is automatically inferred. Users always have the last say on their profiles, being able to correct the information and determine which information can be made public and which can't.

In CUMBIA, a person is always working on a project. Thus, the basic context units are project definitions. Projects are related to documents, people, collaborations and research, but are inherent to each user. So, an ongoing cooperative project might (and almost certainly will) have two or more different project definitions associated with it, one for each user. This is consistent with the fact that each individual has their personal view of reality, and organizes their work accordingly. It is useful, of course, to keep track of the correspondence between individual projects that represent the same group endeavor.

Profiling Agents keep track of the following information:

- **Contact Information**: information necessary for one person, potential or actual collaborator, to contact the user: Name, Title, Email, Phone, etc.
- **Areas of Interest**: areas in which the user has some interest. May be automatically or manually setup, and ranked by interest and activity level
- Projects: projects the user is involved with. These are classified according to their activity status: Past, Present, and Future. Projects may have associated deadlines that can help prioritize agents' work.
- **People**: a user's contact list, this may be classified in different categories, such as personal or work contacts, previous, current or potential collaborators, researchers, etc. Contacts are related to projects in the context of collaborations entered and to areas of interest when the users have similar interests. Users can rate people regarding their relevance to an area of interest or project.
- **Web History**: agents track pages the user accesses when navigating the Internet. The user can rate sites according to their relevance to the work or project in progress (this information can also be inferred from time spent on documents, number of accesses, etc.)

Profile information is saved in a Knowledge Base. The system logs message exchanges (text, email, discussion) and these are linked to projects, forming a personal project history. This history serves to inform future discussions, to assist users in establishing common understanding, and to allow users to "pick up from where they left off" when entering new interactions. In addition, it helps users establish patterns of cooperation, determining which users have been cooperative in the past (and possibly favor these in the future) and which ones have consistently avoided interaction with the user (to possibly avoid these in the future).

Several sources of information can be used to build user profiles: email, bookmarks and publications read and written are just some of them. In many approaches, documents are analyzed for their text content and keywords are extracted and then rate each document according to its importance to the user. Document rating is based on number of accesses, length of access, type of access (read, write, print), and distribution (whether it was sent to someone else or not). Thus, documents in

users' profiles are ranked by popularity, and keyword importance is calculated accordingly. It is important to note that the system allows the user to establish what information is made publicly available for awareness and matchmaking purposes. The system is as transparent as possible, so users understand the process and how their information is being used.

3.3 Awareness and Matchmaking Services

We identify opportunities for collaboration by matching a user's current context with other users who might be interesting to collaborate with, thus "recommending collaboration".

Given a user's context, agents search the profile information for project-related information, building a "project profile". This "project profile" will have information on a user's resources for this projects (documents, references, links) and a user's needs for the project (which can be inferred by searches performed and references downloaded). Agents will search for other users who might be able to provide some of the user's needs. This match is made according to both users' contexts (if one doesn't have to leave their current context to help another, that is preferable than having to change contexts). The match also takes into account the potential for reciprocity (does the other user need something the first user can provide?), in an attempt to provide some motivation for cooperation.

When an opportunity for collaboration comes up, a user is notified. Opportunities are time sensitive, and the user should be informed of the potential for reciprocity (if any) and should be given information on the other user that includes past partnerships and cooperative behavior. Furthermore, it is important to make the initiation of collaboration as effortless as possible. So, when agents detect that some information or document a user has might be useful to another user, they can automatically suggest what to send it be sent. In this fashion, a user has less to worry about regarding finding appropriate information. Users may always choose to engage in longer interaction, entering a chat or message exchange.

To speed up the searches, matches can also be made "offline": these pre-matches are run in the background to find users who might be interested in any of the projects or areas of interest the user is involved in, building and storing simplified models of other users.

3.4 Collaboration Services

Once a collaboration opportunity has been identified, an individual may become an incidental collaborator or an active collaborator. Incidental collaborators are those who provide occasional suggestions and sometimes attend meetings. Active collaborators become inserted in the project, and will have to deal with its schedules and deadlines. It is important to know each participant's status and to know whether any tasks are dependent on them. It may be useful to integrate some project management capabilities to assist with active collaborations. In our system, we provide all the standard collaboration support tools, such as discussion lists,

messaging systems, shared whiteboards, file sharing mechanisms and email. Most of these already exist as modular solutions which can be plugged in, so we will progressively add tools and services to the system as they become necessary.

4 Conclusion and Further Work

One of the biggest problems when moving to a virtual work environment is that individuals lose opportunities for spontaneous, unplanned interactions. These opportunistic interactions usually happen when individuals are co-located and bump into each other in hallways or peek into each other's office for a quick exchange. This has caused the recent emphasis in awareness systems and recognition of the importance of informal interactions. There seems to be a general feeling in the CSCW community that "awareness information is good", so it is provided.

However, many systems proposed provide awareness without any focus to it. Systems provide users with information on others' current online status, past or present activities or objects being used, but provide no extra impetus towards cooperation, hoping that it will happen just because users can now observe each other (to an extent, it does). We believe that more can be done to encourage cooperation: systems can provide additional information on what information is relevant to the other users at the moment and how one can cooperate. It seems that most systems don't exactly have a purpose for providing awareness information. We focus on the opportunities that arise from users' work at their stations. Through careful analysis of users' work activities we can establish what topics are relevant and where they might need assistance. We can also establish what users are proficient in what areas and in this manner find possible cooperation opportunities between them. Most systems lack the ability to uncover opportunities for cooperation. Furthermore, by focusing on the reasons for providing awareness information and working towards that end, we expect to reduce the information flow and direct users towards opportunities for cooperation. This is better than providing a lot of extra awareness information that users will have to sift through, and consider whether each one is useful. Further work is also needed in the areas of user motivation for cooperation and user trust.

We adopt an agent-based approach to solving this problem, giving each user an Agent Team to assist with knowledge and information management, as well as search for peers who might be interesting to make contact with. We believe the first step towards establishing cooperation is identifying the potential collaborative situation, and the next step is making the act of cooperating as effortless as possible. Profiles contain a wealth of information so we can test different matching techniques and variables to determine which work best. Additionally, all information and matching rules may be parameterized by the user, to achieve better results. There is still much work to be done, namely in the areas of context inference and rule building for matchmaking. We will be testing and improving on matching and profiling methods as the project evolves, to see what rules work best.

We are implementing the first prototype, for an academic work environment, and should have some initial results soon. Academic work environments are usually loosely structured, in that there is some structure but it is flexible and adaptable.

Usually, several opportunities for spontaneous collaboration exist. In this environment, group formation happens as common interest appears and individuals come together to work for a period of time (the duration of a project) and disband later (but ties remain, as does the possibility of further collaboration). Additionally, cooperation may be externally triggered, for instance, with the appearance of a new funding opportunity. External funding agencies may provide guidelines for projects (which usually involve): in this case it becomes important to find and organize a group of interested, qualified people to form a group and write project proposals to take advantage of the opportunities. We expect this to be a good application domain, and one which will provide plenty of useful information for our project.

References

1. Alarcón, R. and Fuller, D. Intelligent Awareness in Support of Collaborative Virtual Work Groups. In: Haake, J. M. and Pino, J. A. (Eds.) CRIWG 2002, LNCS 2440, pp. 147 – 167, Springer-Verlag, 2002
2. Bergenti, F., Garijo, M., Poggi, A., Somacher, M. and Velasco J.R. Enhancing Collaborative Work through Agents. VIII Convegno dell'Associazione Italiana per l'Intelligenza Artificiale, 2002.
3. Ellis, C.A. e Wainer, J. Groupware and Computer Supported Cooperative Work. In Weiss, G. (Ed.) Multiagent, Systems, MIT Press, 1999.
4. Enembreck, F. and Barthès, J. P. Personal Assistant to Improve CSCW. Proceedings of the 7th International CSCWD, Rio de Janeiro, 2002, pp 329 – 335.
5. Hoffman, M and Herrmann, T. Prospect Awareness – Envisioning the Benefits of Collaborative Work. Available online at: http://iundg.informatik.uni-dortmund.de/iug-home/people/MH/ProspectAwareness/PAhome.html
6. Isaacs, E.A., Tang, J.C. and Morris, T. Piazza: A desktop Environment Supporting Impromtu and Planned Interactions. Proc. CSCW'96, Cambridge, MA, 1996
7. Jennings, N.R. An Agent-Based Approach for Building Complex Software Systems. Communications of the ACM, April 2001/Vol. 44, No. 4
8. Kortuen, G., Gellersen, H.W., Billinghurst, M. Mobile Ad Hoc Collaboration. Proceedings of CHI 2002, April 2002, Minneapolis, USA
9. Kraut, R., Fish, R., Root, B., and Chalfonte, B. Informal communication in organizations: Form, function and technology. In S. Oskamp and S. Spacapan (Eds.), People's reactions to technology in factories, offices and aerospace, The Claremont Symposium on Applied Social Psychology, Sage Publications, 1990. pp. 145-199.
10. Matsuura, N., Fujino, G., Okada, K. and Matsushita, Y. An Approach to Encounters and Interaction in a Virtual Environment. Proceedings of the 1993 ACM Conference on Computer Science, Indianapolis, Indiana, United States, 1993.
11. Morán, A. L., Favela, J., Martínez-Enríquez, A. M. and Decouchant, D. Before Getting There: Potential and Actual Collaboration. In: Haake, J. M. and Pino, J. A. (Eds.) CRIWG 2002, LNCS 2440, pp. 147 – 167, Spring-Verlag, 2002
12. Pinheiro, M.K., Lima, J.V. and Borges, M.R.S. A Framework for Awareness Support in Groupware Systems. Proc. 7th International Conference on Computer Supported Cooperative Work in Design – CSCWD'2002, Rio de Janeiro, Brazil, Sept 2002
13. Sohlenkamp, M. Supporting Group Awareness in Multi-User Environments through Perceptualization. GMD Research Series Report, No 6, 1999. Fachbereich Mathematik-Informatik der Universität – Gesamthochschule, Paderborn, 1998.