Agents in Groupware Systems

Jacques Wainer
Universidade Estadual de Campinas
Institute of Computing
wainer@dcc.unicamp.br

Clarence (Skip) Ellis
University of Colorado
Department of Computer Science
skip@colorado.edu

Abstract

This paper describes four classes of groups and discusses groupware tools for these classes. It concentrates on a particular set of groupware functionalities, agents, and their application to a particular class of groups, amorphous groups, of which the Internet is the clearer example.

Keywords: groupware model, agents, Internet.

1 Introduction

[EW94] proposes a model of groupware in which the central functionalities of a groupware system are classified into three different groups, or **aspects**. The **keeper** aspect groups functionalities related to the operations on shared data. The **coordinator** aspect groups the functionalities related to the temporal sequencing of activities (performed by the groupware users). And the **communicator** aspect groups the functionalities that allow unrestricted communication among the groupware users.

Some groupware systems tend to provide functionalities that, in their majority, belong to only one of these aspects. Thus, a synchronous editor such as GROVE [EGR90], provides mainly keeper functionalities, and a workflow system [JB96] provides mainly coordinator functionalities. We will call a groupware system that provides mainly keeper functionalities, a keeper; and a groupware system that provides mainly communicator functionalities, a communicator; and similarly to coordinators.

But more modern groupware system have other functionalities besides the ones defined by three aspects above. In [Ell97], these functionalities are grouped into the **agent** aspect. The term agent is used in a much broader sense than in domains such as Artificial Intelligence and Social Sciences. Agents, in our sense, are functionalities, that may be implemented by an program that can be considered autonomous and can be modeled as having mental states, but it can also be implemented as a segment of other programs (that implement the other functionalities).

In this paper we will further develop the notion of agents, in particular agent that are appropriate for a kind of group that we call amorphous groups, which includes the Internet.

2 Four families of groups

We find it useful to classify groups into four classes, based on the working mode, perspectives, and attitudes of people belonging to these classes. The classes are not disjoint; in fact the classes themselves can be seen as extremes of a continuum that links these four extremes as figure 1 illustrates. The four classes are **teams**, **organizations**, **social interaction groups**, and **amorphous groups**.

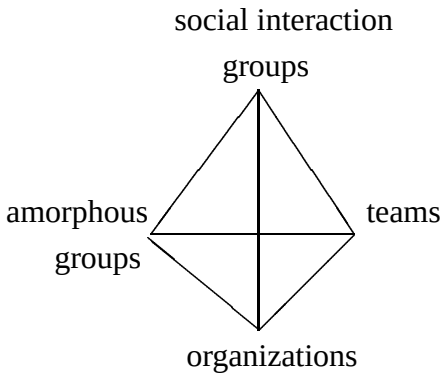

Figure 1: The spectrum of groups

It is important to notice that our classification of groups is centered on a prototypical "working mode", with the goal of analyzing what groupware tools are appropriate for each group. For example, to group all organizations into a single class disregards much of the research in organizations from the management perspective: Mintzberg's five classes of organizations [MQ88], or Harrison's four organization cultures [HS95], or Burns and Stalker's division between Mechanistic and Organismic organizations [BS94].

2.1 Teams

A team is a cohesive, mutually knowledgeable group of people working on a particular task. Prototypically teams have a small number of members, that know and trust each other and that have a common goal (the accomplishment of the task at hand). Also usually the team members are peers, that is, there is no hierarchical or power structure among them. Examples of teams are two authors writing a paper, a football team playing a game, engineers designing the transmission for a new car.

It is important to notice that our definition of team is related to a task, and do not extend after the termination of a task. Thus, under this view one should not say that a set of people make a good team because they can do a lot of things together, but one should say that when doing X (performing task X), that set of people worked as a team.

Typically, teams work in **meetings**: the members of the team get together (at the same time and place) to work on the task. The meeting itself does not follow a predefined sequence of steps, but is dynamically planned to fulfill the needs of the people and the task. The work division among the team members is also not predefined, but evolves as the task evolves.

The work of each team member is highly dependent to each other's work: a member must be aware of what the other team members are doing because it will usually impact on his own

work. On the other hand, each member trust that the other members will warn him when a major disruptive action is about to be taken.

2.2 Organizations

Organizations are a group of people that are ruled by power structures and models. Prototypically, organizations have a large number of people that may not know each other, and that may not share a common goal. The organization as a whole is engaged in many activities and have many goals, and the members of the organization are also engaged in many activities and have different goals.

Members of an organization are related by different formal structures, such as hierarchical and power relations. Each member of the organization can play different roles, and the roles define a set of rights to perform some actions, to access some information and so on. Information about the identity of the organization members, the roles they can play, and all the formal structures that are embedded into the organizations are usually available as models (or directories).

The typical mode of work in an organization is **office work**, that is people working individually in one or more activities that are part of a larger, predefined procedure. Furthermore, because the members of the organization do not necessarily have the same goals, do not necessarily trust each other, are engaged in many different activities at the same time, one cannot expect that the organization's members will spontaneously coordinate their work. Thus both the work division and the sequencing of activities are predefined. The work is usually divided into individually executable activities, that are ordered in a way to guarantee that all information needed to perform an activity is already available when the activity is performed.

In terms of the management literature, our definition of organization corresponds to Mintzber's Machine Bureaucracy, to Burns and Stalker's mechanistic organization, and less precisely to Harrisson's organizations with role culture.

2.3 Social Interaction Groups

A social interaction group is a group of people that is socializing. Socializing groups are small sized, whose members know each other, but do not have a common goal (besides socializing itself). The typical mode of interaction for social groups is **conversation**, that is a unstructured, situated, highly interactive, communication activity.

We will not elaborate much on social interaction groups in this paper, which does not mean that this class of group is unimportant; on the contrary, the proliferation of chat rooms and conversation news groups in the Internet, the explosive use of e-mail for keeping in touch, are all examples of the importance of social interaction groups.

2.4 Amorphous Groups

Amorphous groups are groups formed by very large number of people that do not know each other, do not share common goals, are not embedded into structures, do not trust each other, and so on. In some way amorphous group do not have any of the properties that the other groups have. The only difference between a amorphous group and a mathematical set of people is that there is an **information space** associated with the amorphous group. We will discuss this

information space shortly. In order to belong to an amorphous group a person must provide a way being accessed by the other members, and must also provide some public information that allow the other members of the group to match that person to different descriptions. The public information about each member is the information space associated with the amorphous group.

For example, a city (or better the inhabitants of a city) are an amorphous group: each member can be accessed (contacted) by means of an address (and/or telephone number), and telephone directories, address directories, yellow pages, advertising in local newspapers, advertising posted on business windows, and so on, are all part of the information space associated with the amorphous group.

There are information spaces which are not associated with amorphous groups: a library is an example. A library has not only information (say a book about controlling insects in an egg-plant farm), but has meta information (does this library has a book about insects and eggplant farms, was such a book ever written, which library has the book, what books could contain such topic). But the library (or books) does not provide a way of accessing the author of the book. The point of the library is to to provide the user with enough information to access someone with knowledge on insects and egg-plant farming, but to provide him with the information on insect and egg-plant farming. Thus there is a distinction between the information space for the purpose of providing access means (as in the newspaper advertising) and the information space for the purpose of providing information itself (the library).

The main mode of work, or better mode of interaction, in an amorphous group is **search**, that is a individual searching in the information space of the group for ways of accessing whoever the person needs to access. The person needs to find a place that sell egg-plant safe insecticides, and the whole interaction with the amorphous group is to search through its information space for a member that satisfy the description of "an egg-plant safe insecticide seller".

The Internet is probably the most interesting amorphous group because its information space is computer readable, and so can be explored by programs (or agents in our terminology). Furthermore its information space has both the characteristics of information for the purposes of access (home pages, advertisement) and information for its own purposes (course notes, encyclopedia).

2.5 Intersections

As we mentioned, the classes of groups are not disjoint, but are extremes in a four-way spectrum (figure 1). There are interesting cases of groups that have some characteristics of two of those families, and in fact these groups with mixed characteristics are very common.

Teams gain organizational characteristics when they become too large: the members will not all know each other and trust each other, or to be aware of what everyone is doing becomes to difficult a task. These large teams would create structures such as domains of interest, where subgroups of the team share the same interest/goal for a subpart of the task, or domains of action, where subgroups will have different rights to perform different activities within the task. In each case, the group will adopt, with more or less intensity, structures that limit and organize them. Software development, large engineering design, and large scientific experiments have these characteristics. The members of the teams are organized into subgroups, with different responsibilities and sometimes different rights.

Also, if the team members do not have all the same goals, or do not trust each other to do the right thing, then usually the team will adopt some methodology ([CB98, LT75] for example)

that specify how the work should be divided and how activities should be sequenced, which is also a organizational characteristic.

Communities seems to contain the characteristics of social interaction groups, amorphous groups, teams and organizations. [Wel98] cites some classes of communities: communities of practice, communities of inquiry, and distance learning.

Communities of practice is a group of people that share a set of activities and which interact to accomplish those activities. Communities of practice, although bound by the organization limits, are usually informal and spread across organizational substructures such as departments. Its members help each other achieve their possibly independent activities, and they learn from each other though advice and war stories [LW91].

Communities of inquire is a group of people with the same goal of learning more about a topic. Distance learning (for example [HHT95]) is a more formalized form of community of inquire (and thus possesses more organizational characteristics).

Communities have an amorphous group component because the members of the community may not know each other. It has a team aspect because they have to some degree a common goal, and it has a social interaction aspect because communities do not come together to accomplish collectively their common goal, but to interact and help each member individually achieve her goal.

Large organizations (or disorganized organizations) tend to have amorphous groups. If a member of an organization needs to find someone (or some information) for which there is no model (or directory) then the member has to use amorphous group methods. The fact that an organization cannot have formal structures and models for all contingencies seems to be the driving force behind the development organizational-wide Intranets; these intranet allow members of the organization to search (amorphous group like) for the people and information they need to solve an new problem.

2.6 Discussion

This model of groups is somewhat hybrid: it puts together structures that have since long been recognized as groups, such as teams and organizations, and social interactions "groups" which are better seen as functions. Social interaction is one of the things people do in teams and organizations. The point of this classification is that the computer support need of a group that is interacting socially and a team that are attempting to achieve a common goal are different, as will be discussed in the section below.

3 Groupware for different groups

As it turns out, there is some regularities between different groupware systems classes and different group classes, or better, groupware systems are designed with a particular group class in mind. Next section will discuss the agents appropriate for each class of group.

3.1 Teams

At the team level, groupware seems to be seen as tools: the team can opt to use or not the tool, the groupware is designed to serve very specific purposes, and usually a set of these tools is needed to accomplish a task.

Groupware tools for teams tend to fall into two main categories: keepers and communicators. Keepers are useful because groups have always a particular task, which usually involves the construction of some form of artifact (be it a design or a document). Thus providing the team with the appropriate keeper may help to achieve its goals. For example if the team's task is to design a bridge, a standard single-user CAD system would make the work more difficult. If the work has been divided among the team members, the CAD system itself becomes a bottleneck because no two users can operate on the artifact at the same time.

Communicators, on the other hand, allow the team members to expand the limits of meetings: with teleconference tools (text, audio, or video) the team members do not need to be at the same place to hold a meeting. And asynchronous tools like e-mail and bulletin board enable meetings not to happen at the same time.

3.2 Organizations

On the keeper side, configuration management systems, document management systems, and organizational memories (based on Lotus Notes, for example) are common examples of organization level groupware.

On the coordination side, workflow systems are gaining a lot of attention recently as a potentially very important groupware system for the improvement of efficiency in organizations. Software process management systems [FW96] (also called Process Centered Software Engineering tools) are also example of coordinators that seems appropriate for (software development) organizations.

At the organizational level, it is more usual to refer to groupware as systems. These are large and possible expensive software, whose use is mandatory for the organization members, and on which depend the correct execution of the organization's tasks. For example, if an organization has a document management system, not only all members of the organization that have to access the documents must do it through the system, but the correct behavior of the system becomes critical for the successful accomplishment of the tasks.

3.3 Amorphous and social interaction groups

At the amorphous group level and social interaction groups, the only appropriate groupware aspect is communicators. People in social interaction groups want to communicate with each other either synchronously (chat, MUDs, video conference, media spaces [NYNI96]) or asynchronously (e-mail, bulletin boards).

At the amorphous group level, communicators are also the appropriate groupware once the person knows who to contact.But for these two groups, agents, which are discussed below, are also appropriate groupware.

4 Agents for different groups

Since the term agent is used in this paper to encapsulate all functionalities that do not fall within the three aspects, it is difficult to create a taxonomy of agents. But there seems to be different uses for agents depending on the type of group. We will discuss these main agents uses in the different groups, with the exception of social interaction groups.

4.1 Agents in teams

Not many agents are used in groupware tools to support teams, and if one is used, it tends to be very specific for the task at hand. One example of agents for teams are **critics** [FNO+93], AI-based programs that criticizes, verifies, and makes suggestions to the artifact being constructed (using a keeper). These critics behave as limited team members, they "know" about the domain and make specific contributions to accomplish it.

4.2 Agents in organizations

In organizations, the most common agent is an **access mediator**. Access mediators are a generalization of access control, instead of just controlling the access of a person to a program, service, or data, an access mediator classifies "something" that is being pushed onto it, and based on this classification allows total, partial, or no access to the program, service or data. In the case of a standard access control mechanism, what is being pushed is an "access request"; the access mediator would use some model to classify this access request as coming from a particular class of sources and allow access according to such class. One of the classes that an access mediator should determine is that of *foe*. No access is allowed for requests coming from foe.

For example, a database may classify its users according to three different categories and allow access to different tables for each category. The access mediator would classify an incoming request for access according to the table that relates each user to his category. Users that are not in any category are classified as foe, and denied access.

But in general an access mediator deals with more then access request to a service; the access mediator classifies request, information, notifications, and so on that are being *pushed* into any person or service. An access mediator may serve as a filter for information being generated in a very large design project (by a team with many organizational characteristics). The access mediator may know that notification of changes performed in module X23 are not interesting to this particular member of the team, but that notifications of changes to module X50 are very important to him. Thus the access mediator classify notifications from module X23 as foe, and notification from module X50 as important, and deal with the classified notifications appropriately.

4.3 Agents for amorphous groups

As we mentioned above, agents seems to be the main class of groupware applications for amorphous groups. We will in this section try to define some classes of agents.

In the Internet, which we will use as our main example of amorphous group, a person provides three forms of information. There is the **public information**, that is, information that a user wants other members of the group to access, and in fact hopes that the others will access. It corresponds to pages in the WWW, or files in a anonymous FTP area. The user is interested that his public information reaches the largest number of users (that may be interested in it). The same is true for a provider of services; the company wants everybody that may be interested (or may become interested in that service) to know about that it. In a city, public information would correspond to advertisement, telephone directories, address directories, and so on.

The other form of information is **voluntary information**, that is information that the user provides to some service (or person) for a particular, agreed upon reason and do not expect that

that information will be used for other purposes. A user providing to a search engine a list of topic he is interested in is providing voluntary information: he wants that service to know what he is interested in, but does not expect that this information be used for anything else.

The final form of information are **"fingerprints"**, that is marks that the user leaves when he access other information or services, and that the user is not aware of it. There are many examples of fingerprints in the Internet: when an user accesses a Web server, his client may tell the server who the user is, from where he is accessing it, and so on. or when a user sends someone an e-mail, a machine in the route of that message may collect information about that user. In this paper voluntary information that is used for purposes beyond the ones agree upon will also be classified as fingerprints.

We will discuss below some classes of agents that are useful for amorphous groups, and in particular, useful in the Internet. Some of these agents already exist, some do not exist yet, but they are all conceivable under current technology.

4.3.1 Information gatherer

A large class of agents in the Internet are information gatherers of some sort. In fact most of them can be described by a combination of mode (active or passive), type of information it gathers (public, voluntary or fingerprints), and propose of the information (copy, extract part of the information, classify the information, classify the source of the information, among others).

The mode describe if the agent actively search for the information or just collect the information that on purpose are send to it. The type of information collected is the three type already discussed. And the purpose of the information is a description of why the information was gathered.

Let us discuss some examples. Search engines are active, public information gathered for classification of the information: they actively gather WWW pages (public information) and classify the pages according to words present in the page. A cache system (or a mirroring system) is an active, public information gatherer for copy.

Another agent are the active, fingerprint gatherers for extraction. These agents gather fingerprint (and public) information (in news articles and in WWW pages) in order to collect e-mail addresses. Such e-mail addresses can be used by a marketing company to send unsolicited, advertisement e-mail, or to provide a service of finding people's e-mail given information about domain, and name (Netfind [Net]).

There are agents that passively collect voluntary information in order to classify the source of information. These include GroupLens [Gro] and MovieLens [Mov] GroupLens users send to the agent their rating of Usenet articles. The agent uses this information to classify the user, and then it is able to tell one of its users that "people like her" liked this article that she did not read.

The taxonomy allow for agents of which we have not heard yet, but are possibly in the future. Things like passive fingerprint for source classification: a business in the net collects cookies from potential customers that access their pages in order to try to classify that customer's interest.

4.3.2 Access mediators

Access mediators are also important agents in amorphous group. They control and restrict the access of people to information, services, and other agents. On the other hand they also control

the access of agents to these resources.

An important access mediator in the Internet are filters: more and more unsolicited information are being pushed onto users (see pushers below), in the form of e-mail, inappropriately posted messages in bulletin boards and so on. Filters allow the user to control what information actually reaches him. Intelligent, adaptive filters would certainly protect the user from unsolicited information.

4.3.3 Pushers

Pushers are agents that try to push information onto the other members of the group. Programs that send the same mail message to thousands of people (SPAMs or UBE - unsolicited bulk e-mail), program that post the same message in hundreds of news groups are all examples of pushers. Pushers are used by members of the amorphous group to make it sure that their public information is received by the largest number of people. It is conceivable that there will be other examples of pushers: pushing messages in chat rooms, sending video advertisement to public video conference servers, and so on.

Any advances on pushers would certainly demand new advances in access mediators that would cancel their effect.

4.3.4 Data synchronizer

One important aspect of the Internet, specifically the WWW, is that pieces of information make reference to other non-local pieces of information. But the information being referred to can change. The data synchronizer agent would guarantee that if a particular WWW page changed location, all pages referring to it should also refer to the new location, instead of the old one. Thus the data synchronizer would guarantee that information collected and stored by information gatherers would be correct, and that pages that are linked to other pages have the correct addresses.

4.3.5 Anonymity servers

Another class of agents in the Internet are the anonymity servers. These agents will protect the user from leaving fingerprint. Examples are anonymous re-mailers [Rem], anonymous news posting, and indirect WWW access [ano].

4.4 Agents for social interaction groups

Agents should have a limited role in pure social interaction groups: people want to interact with other people and use some form of communicator to achieve that. But agents can play a role in keeping the group interacting and in evaluating how well the group interact. For example, Herry Lieberman (personal communication) report of an agent that proposes topics in a chat room.

One can conceive also, agents that look over social interaction groups and evaluate how well the group is interacting, and proposing ways to improve the interaction.

For communities, agents seems to be more common. Agents are mainly used to overcome the amorphous group aspect of communities: people in a community do not necessarily know each other. These agents for communities are many case matchers: the user declares to the

agents her interests (or profile) and the matcher suggests either a set of people with a similar profile to the user, or information derived from those persons. GroupLens [Gro] and MovieLens [Mov] are matchers that return to the user not the identification of people with similar profiles but respectively UserNews articles and movies that those people liked and that the user has not read/seen. Another agent suggests entries in the matched people bookmarks [MC98], or suggests WWW pages that have been positively classified by the matched people [bcif98].

Matchers usually have a information gatherer component. Some may collect voluntary information on the user's profile (GroupLens and MovieLens). Others may collect the users public information (WWW home pages) to infer the users profiles.

5 Conclusions

The ideas presented in this paper are still tentative. They reflect the current state of the authors discussions on extending the three aspect model of groupware [EW94]. The authors realized that the three-aspect theory seems to be appropriate in team and organizational groups but fail to explain the groupware tools available in the Internet. Furthermore it fails to explain the increase in (intelligent) agent research and prototypes that support collaborative activities.

We believe that despite the tentativeness of the model, it does carry some contributions. It is not always clear to the groupware practitioner (but see [Gru94]) that groupware can be seen (and developed) differently from the team and from the organizational points of view. We have never seen an attempt to define amorphous groups such as the Internet as groups, nor to classify *from a groupware perspective* the types of groups and the tools that are appropriate for each of them. Of course, this first classification should now be compared and contrasted with group classifications proposed by disciplines within the Social Sciences.

On the other hand, the idea that other forms of groups can be seen as a combination of the characteristics of two of the classes above is appealing, and should have interesting consequences. If, for example, a software development team is indeed a combination of team and organization, then the development of a groupware tool for such a group has to incorporate both the organizational aspect (roles, access mediators, workflows) and the team aspect (communication tools, maybe synchronous access to objects, the ability to change the rules in force).

Acknowledgments

The authors would like to thank the participants of the 1st International Workshop on Community Ware, help in Paris, July 1998, where a first version of this paper was presented, for interesting discussions specially regarding how communities would fit into the group classification. The authors would also like to thanks the comments of the anonymous reviews, which greatly contributed for this paper.

References

[ano] Anonymizer, inc. http://www.anonymizer.com/.

[bcif98] Content based collaborative information filtering. J. delgado and n. ishi and t. ura. In *Pre-Proceedings of the 1st Workshop on Agentsin Community Ware*, Paris, 1998.

[BS94] T. Burns and G. Stalker. *The Management of Innovation*. Oxford Univ Press, 1994.

[CB98] J. Conklin and M. Begeman. gIBIS: An hypertext tool for exploring policy discussion. 1998.

[EGR90] C. Ellis, S. Gibbs, and G. Rein. Design and use of a group editor. In G. Cockton, editor, *Engineering for Human Computer Interaction*. North-Holland, Amsterdam, 1990.

[Ell97] C. Ellis. Team automata for groupware systems. In *ACM Conference on Supporting Group Work (GROUP'97)*, pages 415–424. ACM Press, 1997.

[EW94] C. Ellis and J. Wainer. A conceptual model of groupware. In *Proceedings of ACM CSCW'94 Conference on Computer-Supported Cooperative Work*, Models of Cooperative Work, pages 79–88, 1994.

[FNO+93] G. Fischer, K. Nakakoji, J. Oswald, G. Stahl, and T. Sumner. Embedding computer-based critics in the context of design. In *Human Factors in Computing Systems, INTERCHI'93 Conference Proceedings*, pages 157–164. ACM, 1993.

[FW96] A. Fuggetta and A. Wolf, editors. *Trends in Software Process*. Wiley, 1996.

[Gro] http://www.cs.umn.edu/Research/GroupLens/.

[Gru94] J. Grudin. Computer-supported cooperative work: History and focus. *IEEE Computer*, 27(5):19–26, 1994.

[HHT95] L. Harasim, S. Hiltz, and M. Turoff. *Learning Networks*. MIT Press, 1995.

[HS95] R. Harrison and E. Schein. *The Collected Papers of Roger Harrison*. Jossey-Bass Publishers, 1995.

[JB96] S. Jablonski and C. Bussler. *Workflow Management Systems: Modeling, Architecture, and Implementation*. International Thomson Computer Press, 1996.

[LT75] H. Linstone and M. Turoff. *The Delphi Method: Techniques and Applications*. Addison-Wsley Publishing Co., 1975.

[LW91] J. Lave and E. Wenger. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, 1991.

[MC98] N. Mathe and J. Chen. Organizing and sharing information on the www using a multiagent system. In *Pre-Proceedings of the 1st Workshop on Agentsin Community Ware*, Paris, 1998.

[Mov] http://movielens.umn.edu/.

[MQ88] H. Mintzberg and J. Quinn. *The Strategy Process: Concepts, Contexts and Cases*. Prentice hall, 1988.

[Net] Internet resource discovery project. http://www.cs.colorado.edu/homes/schwartz/public_html/resource

[NYNI96] H. Nakamishi, C. Yoshida, T. Nishimura, and T. Ishida. Free walk: Supporting casual meeting in a network. In ACM Press, editor, *Proc. of the CSCW'96*, pages 308–314, 1996.

[Rem] Raph remailer list. http://www.cs.berkeley.edu/ raph/remailer-list.html.

[Wel98] B. Wellman. Virtual community: Introducing a new siggroup focus area. *ACM SIGGROUP Bulletin*, 19(1), April 1998.