

Introducing Instant Messaging and Chat in the Workplace

James D. Herbsleb

David L. Atkins

Bell Laboratories, Lucent Technologies
2701 Lucent Lane
Lisle, IL 60532
+1 630 713 1868
jherbsleb, datkins@lucent.com

David G. Boyer

Avaya Communications
101 Crawfords Corner Rd
Holmdel, NJ 07733
dgboyer@avaya.com

Mark Handel

Thomas A. Finholt

School of Information
University of Michigan
3218 School of Information North
Ann Arbor, MI 48109
handel, finholt@umich.edu

ABSTRACT

We report on our experiences of introducing an instant messaging and group chat application into geographically distributed workgroups. We describe a number of issues we encountered, including privacy concerns, individual versus group training, and focusing on teams or individuals. The perception of the tool's utility was a complex issue, depending both on users' views of the importance of informal communication, and their perceptions of the nature of cross-site communication issues. Finally, we conclude with a discussion of critical mass, which is related to the features each user actually uses. More generally, we encountered a dilemma that imposes serious challenges for user-centered design of groupware systems.

Keywords

Instant Messaging, Presence Awareness, Groupware, Technology Diffusion, Chat, Distributed Teams

INTRODUCTION

Instant messaging (IM), after experiencing enormous popularity among recreational users, is beginning to move into the workplace [1, 13, 15], following on the heels of other forms of text-based computer-mediated communication (CMC) such as e-mail and MUD rooms [2]. To date, there have been few published studies of the use of interactive text communication such as IM and chat outside of research groups who were motivated to use it by their desire to explore the technology. As with many types of groupware, one would expect that adoption is not a given among work groups more generally. Interactive text potentially suffers from many of the classic issues of groupware adoption [3]. For example, there are likely to be some individuals whose attention is in more demand than others. IM may impose an undue burden on them, but mostly provide benefits for others, i.e., those who want to reach them. While other forms of CMC, such as e-mail, also have this potential, the synchronous nature of IM makes it harder to ignore.

Independent of its actual utility in the workplace, which is largely unknown at this point, interactive text may be

perceived in a way that makes adoption unlikely. Many people have exposure to interactive text primarily through teenage users who presumably exchange gossip and rumors, talk about their personal lives, and use IM to do all the things teenagers do [4]. Workplace expectations are likely to be either that interactive text is a waste of time, or possibly even socially undesirable.

In addition, in order to diffuse through the workplace, interactive text must achieve collective adoption within a community of interest, as has been shown for other interactive communication technologies [12,14]. Specifically, if collective adoption is described as an "accelerating production function" – where greater number of users leads to more value for potential users – then Markus [12] identifies two key factors in producing collective adoption. First, users must choose to absorb certain costs associated with use, such as acquiring new skills, purchasing hardware or software, and exercising "communication discipline" – meaning regularly reading and responding promptly to communications. Second, there must be initial variation in users' abilities to contribute to and benefit from use of interactive communication systems. Therefore, collective adoption starts when a highly interested group of people – i.e., the "critical mass" – perceives differential value in using a given technology, and through their use demonstrates utility to others, who in turn become likelier to adopt the technology.

With interactive text, then, if there are no other users, it is uninteresting to message with and be aware of oneself. It is not clear, however, what defines the community of interest for interactive text systems, or where critical mass must arise to assure successful adoption. Specifically, is the community defined by the larger organization, where simply attaining some minimum level of adoption will lead to success? Or, is the community of interest defined more precisely, such as within particular sub-units or working groups? Resolving these issues is critical, for example, in determining whether deployment proceeds according to a "one size fits all" approach or whether deployment is tailored to match local characteristics.

Our expectation is that the ability of interactive text to support informal, spontaneous, and opportunistic communication should make it particularly suitable for geographically distributed teams. Previous research has shown that such teams suffer badly from the absence of these

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2002, April 20-25, 2002, Minneapolis, Minnesota, USA.

Copyright 2002 ACM 1-58113-453-3/02/0004...\$5.00.




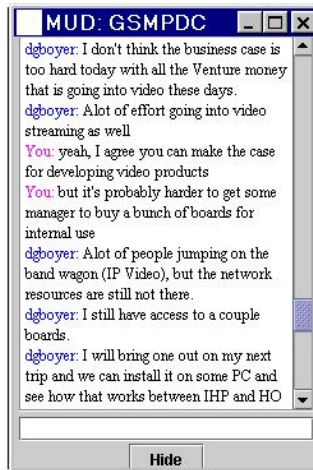


Figure 1. Screen shot of the RVM tool.

to one's presence viewer could be considered that user's buddy list, i.e., the users one potentially wants to interact with and be aware of. Each user has his/her own buddy list. Groups, on the other hand, have a consistent membership, defined at the system level, not at the level of each individual user. In RVM, any user can create any number of groups. By default, the creator is the group's administrator. Groups can be open (anyone can join) or closed (only those selected by the administrator can join). In order to be a member of a group, one must be permitted to join (if it is a closed group) and one must actually join by selecting an appropriate menu option and typing in the group's name.

Group chat is visible to all participants in a group, and only to members of that group. Group chat windows open automatically when one logs on to RVM, and the last *n* hours of conversation are displayed. (The length of time is adjustable by each group's administrator.) Group chat resembles, and in fact was originally inspired by, workplace MUD rooms [2]. It differs in that it is not threaded, has a single "room," has limited persistence, has only text chat functionality, and is intended to be continuously present.

We also wrote documentation for RVM, including an on-line user manual, and a 1-page quick reference card. Both prominently displayed e-mail addresses and phone numbers where help was available. We in fact received many e-mails, but very few phone calls, soliciting help with installation or with RVM features. Installation was accomplished with a standard installer – the user needed only to double-click on it to install the application. Registration of new users was handled by the research team for the first six months, after

which we deployed a web page where users could register themselves.

SITES AND DATA

Over the course of approximately 17 months, RVM was introduced into several work teams. All of these work groups were geographically distributed. The following table shows the sites and teams for which we report data in this study. In the next section, we describe how RVM was introduced to each of these groups.

Organization	Team	Sites
Wireless 1	Management 1	UK, Germany, France
	Systems Engineers	UK, Germany
	Quality	UK, Germany, France
	Test	UK, Germany, France
Network	Architects	Four US sites
Wireless 2	Management 2	Ireland, US

There was no overlap in people or even sites among the three organizations. Most individuals in the teams from Wireless 1 were acquainted with one another, and there was some overlap in membership among the teams. In particular, one first-level manager was a member of Management 1, Quality, and Test. Both of the management teams consisted of a second-level manager and the first-level managers that reported to him/her. Two of the other teams (Quality and Test) consisted of a first-level manager and two different teams that he managed. The other two groups (Systems Engineers and Architects) consisted of technical staff (although the staff eventually persuaded their first-level manager to join the Architects group). In the two management groups (but none of the others) RVM use was mandated by the ranking manager. The teams varied considerably in size over time, but all were generally in the range of 5-10 people.

We gathered usage data via automatic logging on the server, which included logins, logouts, joining and leaving groups, as well as group chat messages. In order to preserve users' privacy, we did not log instant messages. We also conducted about two dozen semi-structured interviews with users, and two small focus group sessions to get feedback.

INITIAL INTRODUCTION -- INDIVIDUAL SESSIONS

We planned carefully for an initial introduction of the tool at two sites in the Wireless 1 organization, one site in UK and one in Germany.

Targeting Key Users

In order to have the greatest impact, we used the results of a survey to target users who appeared to have the greatest need for a cross-site communication tool. Prior to the introduction effort, we conducted a survey, primarily to help us understand the prevalence of specific communication and coordination problems. We also used the survey to identify who communicates most frequently with whom across sites. We used the results of two additional questions about willingness to try new communication technologies to



functionality of the tool, but they were not shown *how to collaborate*¹ in any meaningful way with the tool.

Privacy versus setup time. Recall that our initial solution to the privacy concerns of potential users was to allow each user to select who could see his/her presence information, with a default setting that other users could see no information. While this did seem to put an end to privacy concerns, it made setting up new users inordinately difficult. When a new user installed the application, he/she was not able to see any information at all about other users, since these users had not yet permitted the new user to “see” them. In order to have presence information, the new user would have to contact each other user individually, and ask him/her to permit this. The procedure required considerable effort, and introduced substantial delay between the initial training and installation session and the first point at which something interesting happened with respect to presence, IM, and group chat. For many potential users, the amount of effort this required surpassed what they were willing to expend.

Adoption by Teams

We began to realize that where RVM was adopted, it was adopted by all or some substantial part of a team. Not all teams who began to use it continued, but nowhere did it seem to take off just by virtue of individuals finding each other and chatting. Figure 3 shows usage data (seven day rolling average) for four teams from Wireless 1. The Systems Engineering and Management 1 teams quit using the tool rather quickly, primarily because of reliability and usability problems. On the other hand, the Quality and Test teams, while starting more modestly, gradually added users, and endured for over a year. (As shown in Figure 2, use fell to near zero as various organizational changes affected all teams. We left off this “tail” to save space.)

RETOOLING, RETHINKING

As a result of our experiences with this initial round of RVM introduction, we made a number of changes in both the tool and the methods we used to introduce it.

Tool changes. In the weeks and months after the initial installation, we tested the RVM client and server extensively, fixing many bugs, and making the tool much more reliable and usable. We altered permissions to be group-based. After this change, all members of a given group, by default, can chat with each other and see each other’s presence information. The idea was to preserve privacy with respect to other RVM users in general, while making it easy to permit one’s team to see one’s presence information. Adding oneself to a group provides a much easier mechanism for making presence information available to others, thereby, we hoped, quickly giving new users a capability that would interest them.

It was also at this point that we introduced persistence in group chat. Prior to this, group chat had been available, but




Figure 3: Daily number of active users for selected RVM groups. Number of active users is measured by the number of group members who log in.

was not persistent, so a user could see only what was entered in the chat room while the user was actually logged on. There was no context and no history. With persistent group chat, when a user logged in he/she would see all chat that occurred for the last *n* hours (default of 24).

Installation and training. We decided to focus on teams, rather than pairs of people, and insofar as possible, we trained entire teams at once. We accomplished this in Wireless 2 by having 1-hour teleconferences (using a bridge to get the best audio quality possible). We were able to train most members of three teams in two 1-hour sessions. About half of the teleconference session was devoted to RVM, the other half focused on several other collaboration tools. For Network, we initially trained two members of the Architect team in person, who then went on to help their colleagues get set up and learn the tool.

We prepared for these sessions much as we had before, but the actual time and effort we spent during the installation and training sessions was much less, since we worked in groups. Prior to training, we set up groups we expected each user to be interested in based on information gained in interviews with contacts in the organization or with the new users themselves.

In addition to the installation and learning the tool functionality, we also had users chat briefly in group and individual chat windows, configure time-outs, and practice setting their presence status. The sessions were rather hurried, but we were constrained by the users’ very tight schedules. The users had telephone and e-mail contact information for help, as well as a quick reference card and web-based manual.

¹ Walt Scacchi’s succinct expression.



mail, for which there is a well-recognized need. Moreover, the most widely-publicized use of interactive text is teenagers gossiping via IM.

The utility of interactive text is subtle. “Water cooler” conversation is widely known by researchers to be a vital way for people to stay in touch, and pass along knowledge to colleagues [10]. But it is not necessarily clear to the workers themselves how important this sort of informal communication is. Rather, they often seem to perceive it as recreational, as a form of “goofing off.” These perceptions are not terribly important so long as a mechanism for supporting informal communication is intact, since such communication tends to happen spontaneously. But if members of distributed teams fail to see “water cooler” talk as “real work,” it may be difficult to persuade them that they need a tool such as RVM to replace something they may regard as frivolous.

Uncharitable attributions. In the many interviews we conducted before introducing any tools, it became clear that there was considerable friction between workers at different sites. The attributions made by remote team members about the causes of their remote colleagues’ “irritating” behavior are often such that additional opportunities to communicate do not seem desirable. Especially during the early stages of sites beginning to work together, we saw many instances of what we came to call “uncharitable” attributions about behavior at the other site. If e-mail was not answered promptly, it was because the remote person was not responsible, or did not respect the sender. If questions were misunderstood, or not answered fully, it was not a mere language issue, it was likely the competence, commitment, or diligence of the other person that was in question. After people got to know each other, such attributions began to disappear, but they were quite common early in cross-site relationships.

It is a well-known finding of social psychology that observers of action tend to attribute actions to personal characteristics of the actor, rather than to the actor’s situation, while actors tend to see their own behavior more as a product of the situation they are in [9]. One of the reasons for this difference is simply the greater amount of information the actor has about his/her situation [8]. Enhancing awareness of remote context may help address this asymmetry. Yet if the perceived problem is that remote team members are uncooperative, or uninterested, or hostile, or inept, then it is not clear why one would want *more* communication with them. Such perceptions, e.g., “they are uncooperative” do not lead in any straightforward way to “we need IM and chat for additional context.”

Critical mass

Interactive communication technology must have a critical mass of users in order to achieve the larger goal of collective adoption [12, 14]. In general, within a given community of interest, the more users of a new interactive technology, the higher its perceived value for non-users within the same community of interest. Our experience has shown that what

constitutes an effective critical mass is subtle and can vary dramatically depending on different definitions of the community of interest.

Who is part of my critical mass? In one sense, critical mass simply means a group of users for whom a given new interactive technology has differential benefit, which leads to higher levels of use of the new technology by that group. One assumption is that the criteria that determine benefit for potential members of a critical mass are universal within an organization. This simple view is not necessarily sufficient to explain the patterns of adoption we describe in this study, however. In particular, people seem to value very different parts of the functionality, and use the tool in very different ways – in part as a reflection of their membership in different organizational sub-units. For example, users in settings that primarily value the group chat capabilities must have a different kind of critical mass (i.e., some number of other users to actively participate in group chat rooms – an ongoing kind of communication discipline) compared to users in settings that primarily value the tool’s indication of presence (i.e., some number of users who merely agree to be visible to others – a one time decision to change an awareness setting). There is a potential asymmetry here, since the “group chat” users might also be part of the “presence only” users’ critical mass, while the reverse is not true. To further complicate the picture, the critical mass number for some features, e.g., chat, may be small. Perhaps the ability to chat with only 1 or 2 other people is sufficient [1]. Use purely for other features, e.g., presence information, on the other hand, might require many more users in order to be sufficiently valuable to keep users logging in regularly. Such wrinkles present a potentially complicated picture of what constitutes critical mass, especially when it is difficult to predict what combinations of features various users in different organizational settings will find valuable.

Another facet of critical mass, to which we fell victim, poses a particular dilemma when introducing any of a large number of groupware tools. There is widespread agreement among CHI professionals that it is essential to get feedback from real users in order to evaluate a design. For many groupware systems, there is really no way to do this other than producing a real, working system. Low-fidelity prototypes can be useful for evaluating interfaces, but understanding how the technology impacts social, political, privacy, and other concerns will require actual use over some period of time.

The problem is that only a few users will be innovators or early adopters, to use Rogers’s [14] terminology, willing to tolerate a technology that is not completely “cooked.” This is not such a problem for single-user applications, since one can learn from the experience of these few to produce a better application. But if the application requires collective adoption, and only a small percentage of users will be willing to tolerate its unpleasant features, creating sufficient critical mass to move toward collective adoption may be difficult or impossible. Under these circumstances, then,

