

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.

communication opportunities [5], particularly when they need to respond to unanticipated events. Such problems appear to greatly slow collaborative work that is split across sites, in comparison to comparable work where all activities are conducted at a single site [7]. An application that shows when colleagues are available, and makes it easy to non-intrusively signal the desire to communicate [16] should facilitate this type of communication.

As part of a research program designed to address and alleviate the communication and coordination problems of multi-site software development work, we deployed an interactive text tool that we developed, called Rear View Mirror (RVM). In this paper, we focus on adoption issues, describing the strategies that we pursued in order to facilitate adoption, the results of those strategies in terms of usage, and diffusion of the technology through organizations and teams.

REAR VIEW MIRROR

Background

The idea of introducing interactive text in the workplace came about as a result of our analysis of observational and interview data gathered at a dozen software development sites in North America, Europe, and Asia. Collectively, we spent months on site, and we conducted over two hundred semi-structured interviews over a period of approximately two years in order to fully understand the issues of multi-site projects, and the various approaches to communication and coordination used in the company throughout the world. The interviews included technical staff, supervisors, managers and executives.

The findings from these studies provided us with the set of problems we set out to address. Chief among these issues were substantial delay in any work split across sites [7], and difficulties in coping with unforeseen events [5], and the difficulty of finding and negotiating the availability of co-workers [6]. One factor that clearly contributed to these problems was the virtual absence of informal communication between sites. Such “corridor” or “water cooler” talk is essential for keeping people aware of events and concerns, who is doing what, and more generally, the context in which other people are working.

We developed RVM as one thread of activity designed to help alleviate these problems. The organizations into which we introduced RVM helped to sponsor this work, and the members were generally interested and cooperative in working with us. It is important to note, however, that we were not responding to a specific request for a tool of this type. Rather, after going to great lengths to try to identify and understand the problems these multi-site projects were experiencing, we proposed RVM as one part of a solution. In general, use of RVM was encouraged, but not mandated by managers (two exceptions will be noted below).

Low-fidelity prototype. We called our team-oriented interactive text tool Rear View Mirror, with the idea that it was small, unobtrusive, and allowed you to see what was going on around you with a glance. While the initial tool was in a very early stage of development, one member of the

research team visited two of the intended introduction sites. Using a low-fidelity paper mockup of the tool, he had 30-minute sessions with several potential users at each site, explaining the intended functionality, and asking if the potential user thought it would be useful, and if they had any concerns.

These initial feedback sessions were generally favorable, in that all the users indicated they thought such a tool would be useful for staying in touch with colleagues. Many users were not familiar with instant messaging or chat, and were initially confused about how they differed from other text messaging, such as e-mail.

One potentially serious concern surfaced in these initial sessions. Particularly at the German site, users had major reservations about the awareness aspects of the tool, and its potential use for the purpose of surveillance. We were told that the workers’ council would have to approve it, and that it might in fact be illegal under German law.

In response to this initial feedback, we concluded that we should plan to introduce people to the concept of instant messaging and, as opposed to merely introducing them to a new tool. We also resolved the concern over surveillance by building a security model into RVM such that a user could see another user’s presence states only if explicitly permitted by that user. This alleviated the privacy concern, and the work council representative was satisfied.

Application

RVM provides three primary types of functionality: presence awareness, instant messaging, and group chat. Users can add other users to their presence viewer (seen as iconized photos in Figure 1), and, if permitted, can see that user’s presence states. Presence states are indicated by the color of the border around that person’s image (e.g., green indicates present, yellow generally means temporarily not present, and red generally means away for more than an hour or two.) A mouseover on a person’s icon produces more information, e.g., “away for about an hour.”

Users can explicitly set their own presence to one of several pre-determined states (lunch, meeting, or busy, which turn the border yellow; gone for the day and vacation, which turn the border red). Users can also configure RVM to change border colors and display (on mouseover) user-defined messages after user-selected intervals of keyboard and mouse inactivity. Users often had fun with this feature. One user, for example, set RVM to automatically turn his border yellow and display the message, “Missing, presumed working” after an hour of inactivity.

The second primary type of functionality was instant messaging. For anyone whose icon appears in the presence viewer, one can open an IM window, type in a message, and a window pops up on the other user’s screen. Text messages can then be exchanged in the usual way.

Finally, RVM provides a group chat functionality. It is important to clarify what we mean by groups. Groups are quite distinct from “buddy lists.” The set of users one adds

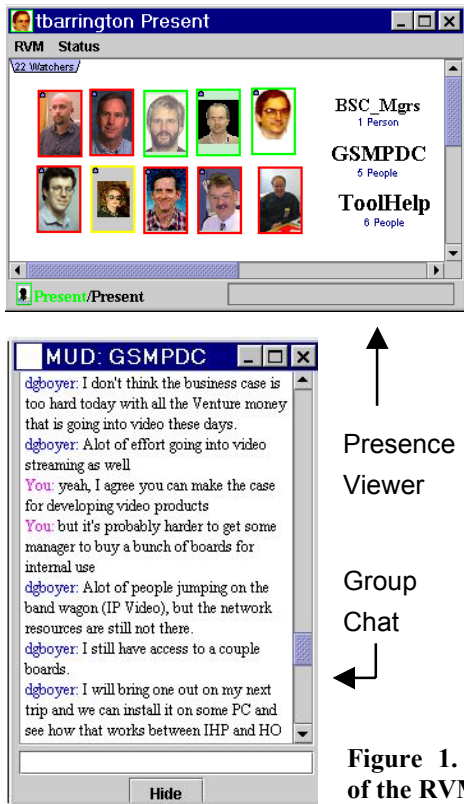


Figure 1. Screen shot of the RVM tool.

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

further refine this list to what we predicted would be our most likely adopters.

Installation and Training

Based on previous experience with training users on collaboration tools, we decided to make individual, one-hour appointments with each potential user. A full hour, we reasoned, would give us ample time to describe several collaboration tools to the user, install and test whichever tools the user wanted, and systematically train the user on the tools' functionality. The tools included RVM, an on-line calendar, and an expertise-finding tool. At the end of the hour, we left the user with a quick reference cards for all of the tools, and contact information for obtaining help by phone, e-mail, or a group chat in RVM.

We wanted to achieve critical mass quickly, so people would not lose interest in RVM before there were enough other users to make it interesting. To this end, the installation and training were carried out by two 2-person teams, during one intensive week, simultaneously at both sites. With our encouragement, the director of the department formed an RVM group with his direct reports. They were among the first who had the tool installed and began using it. After this initial week of installation and training we had one member of the research team at each site for an additional week. This person informally followed up, asking users if they were having problems, if they were finding the tools useful, and being generally available to troubleshoot problems.

RVM Adoption Issues

In contrast with our expectations, it proved relatively difficult to introduce RVM in the workplace. Looking at adoption at the organizational level, we computed a percentage figure to get a rough idea of the proportion of people who continued to use the application once they were trained and had it installed. The percentage is computed on a daily basis. The numerator is the number of people in an organization who log in on a given day, and the denominator is the *cumulative* number of people who had installed the application and logged on *as of that date* (demonstrating that they know of RVM, are able to log on, and have seen some of its basic features). So the numerator is the *actual* number of users on a given day, while the denominator is intended to represent the total number of *possible* user on that day.

As you can see in Figure 2, while we initially had about 50% use among those we

trained in Wireless 1, within 2 months this dropped off to about 20%, then gradually leveled off at about 10%. We had hoped to achieve an organization-wide critical mass, but instead we eventually lost about 90% of the possible users. There were a number of factors that caused users to decide not to use, or to discontinue use, of RVM. We describe these factors next.

Software engineering, usability engineering. The application we released initially was essentially an alpha version, and beset by many of the difficulties one would expect with software in this state. There were reliability problems with both client and server, as well as usability issues, with such things as confusing menu items, lack of feedback for some actions, and mysterious error messages. We mention these issues simply because they clearly played a role in discouraging use of the tool.

Choosing potential users. In most of our early introductions, we had focused our efforts on recruiting the most highly connected pairs of potential users who resided at different sites. Behind this strategy was the thought that we needed to achieve a critical mass relative to the larger organization as quickly as possible. We assumed that our installation effort needed to be sufficient to make a substantial proportion of the entire organization RVM users. We gave some attention to the teams that these pairs belonged to, primarily to have a starting point for demonstrating the team functionality to them. But we did not specifically target teams, i.e., try to get most or all of a team online.

Individual training sessions. The individual training sessions, based on our experience with training users on groupware applications, had several major shortcomings. When people first logged on, there were no other users who

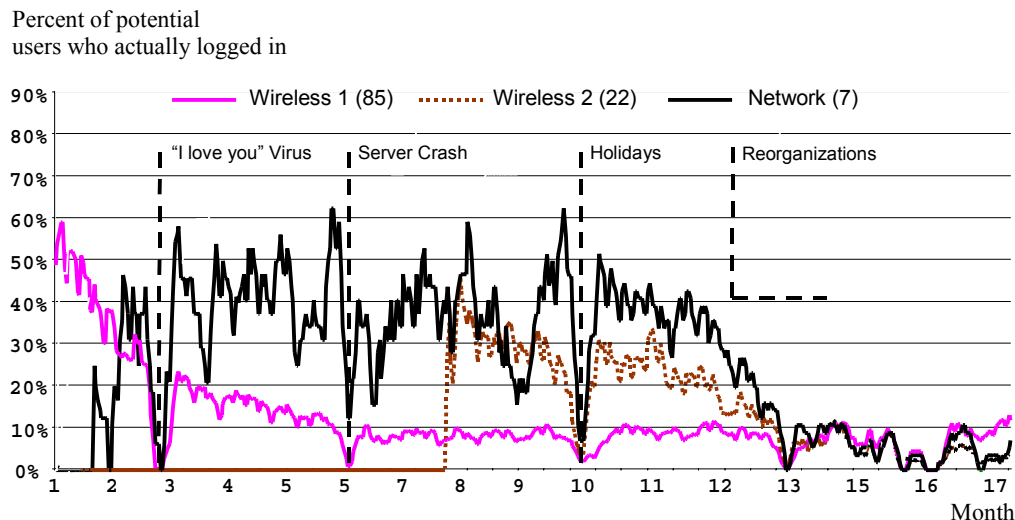


Figure 2: Daily use of RVM over time in each organization, as a percentage of the cumulative total of potential users in that organization. The figure also notes important external events that impacted RVM usage. The numbers in parentheses are the final cumulative total number of potential users in each organization.

had visible, meaningful presence information. There was usually no one else to chat with. The users learned the

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

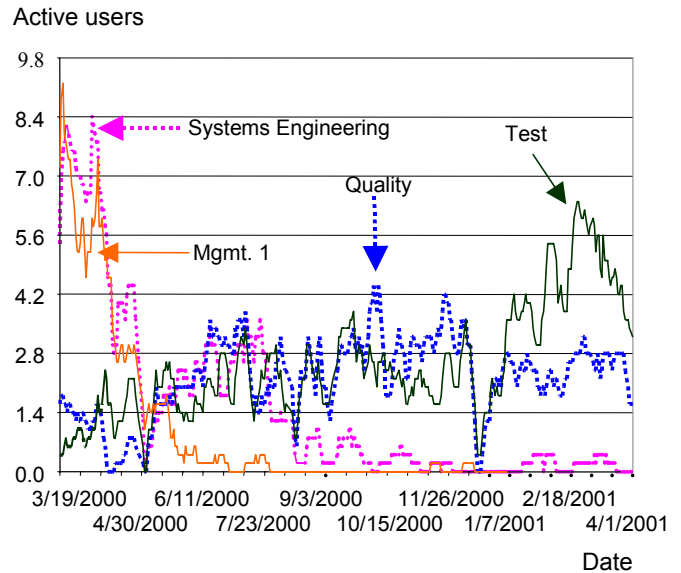


Figure 3: Daily number of active users for selected RVM groups. Number of active users is measured by the number of group members who log in.

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.

Adoption of RVM

These changes seemed to improve RVM’s reception, and Figure 2 bears this out. Wireless 2 and Network both sustained substantially higher rates of use than the Wireless 1 organization. Wireless 1 settled in at around a fairly consistent 8-10% of possible users logged in each day. Wireless 2, after briefly flirting with 40%, converged on about 20-35%. The Network organization appears to be the most volatile, jumping between about 20% and 60%. The differences in daily login rates between Wireless 1 and Wireless 2, and between Wireless 1 and Network are highly statistically significant, using student t-tests, 1-tailed. Wireless 2 and Network are not significantly different from each other.

Figure 4 shows daily logins for the two teams that adopted RVM in the second wave. The Architects team (in the Network organization) gradually increased use over time, while the Management 2 team (in the Wireless 2 organization) rose abruptly to a fairly high level of use, and stayed fairly constant throughout the period. The Architect group, much like the Test and Quality groups in the Wireless 1 organization, grew “bottom-up.” Two users interested in RVM began using it. As they gained experience and became convinced of its value, they became champions, recruiting the remaining team members. Interestingly, these two initial users had adjacent offices, though the rest of their group was widely distributed.




Figure 4: Daily number of active users for Architects / Mgmt 2 RVM Groups. Number of active users is measured by the number of group members who log in.

The Management 2 group, on the other hand, began using RVM because a senior manager mandated its use. This proved effective at quickly generating sustained use. However, the use was unusual in that this is the only group with substantial logins that did not use the group chat feature – only two messages were entered during the group’s

existence. And perhaps also reflecting the manner in which the group was initiated, some members revealed in interviews that they believed RVM was used primarily as a surveillance tool, to keep track of when people arrived and left. We have no evidence of such use, but given no other salient use, e.g., group chat, it is not surprising that some users reached this conclusion.

ISSUES IN WORKPLACE IM ADOPTION

The work groups who began to consistently use RVM clearly believed they were reaping benefits from the tool. For many users, group chat was by far the favorite feature. It provides a lightweight communication channel, a way to find out who is available, a trigger for opportunistic communication, some degree of team context, and a way to query one’s entire team at once.

In addition to group chat, many members of these groups made extensive use of IM for conversations that were too sensitive or private for group chat. Several individuals reported that it was in fact the presence awareness feature, rather than chat or IM that brought them to the tool. They could see who was around, and had a good idea of when they could reach people. One manager told us that he found RVM useful only because of the presence indication it provided. This allowed him to know when his colleagues were present, and “saved me a few steps” by telling him when he could walk down the corridor and find them in their offices. He did not have frequent communication with people at other sites.

Perceived Utility

There are several different reasons that people gave for decisions not to use RVM.

IM and chat as superfluous. It can be difficult to describe the advantages of interactive text communication to new users. If one needs to communicate, one can call or send an e-mail, which is what our users were accustomed to doing. We heard a number of comments to this effect, such as, “It’s just as easy to pick up a phone and call someone.” Several people said that IM and chat just didn’t fit into their established ways of communicating: “... [I] already was using e-mail, phone and voice mail.” Others feared being overwhelmed by yet other kinds of message. Feeling inundated with e-mail and voice mail, they saw IM and chat as further encroaching on their time to do “real” work.

Others focused their comments on the awareness functionality, i.e., being able to tell whether someone was currently logged on and active. Some users questioned the value of this type of awareness information, e.g., “It is not that valuable to know that someone is sitting at their desk.” Others found such functionality useful, especially in conjunction with teleconferences: “Sometimes used it to check on conference calls -- were they there?”

Is “water cooler” talk real work? IM is a tool for which most workplace users do not have a clear, well-defined need. One does not generally hear people in the workplace asking for IM tools to help them get their work done. It is not a direct replacement for any other tool, such as the phone or e-

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,

usable user feedback is likely to be quite elusive. This presents a genuine “groupware critical mass” dilemma. You need a robust, usable application in order to achieve critical mass – which will lead to collective adoption and the opportunity to observe the application in use by a significant population of users. On the other hand, you need to observe the application in use in order to refine the design into a usable, robust application.

There are several possible solutions, although none is completely adequate. We used them all in this study. One common solution is to have the team of developers use the tool. While they are likely to be sufficiently tolerant, they may be atypical of other users, hence the application refined around their use may still run afoul of social or political issues, or may not fully address needs of naïve users. Another possibility is to try to find an extremely tolerant user group. We found this in the Architects group. They hung in there through many application crashes, a bug that degraded performance of the user’s machine, very awkward administrative procedures for adding new users, and more. By the nature of their jobs, and their tool-related responsibilities, they all seemed to be innovators or early adopters. Finally, one can try to arrange for management pressure for adoption, in order to force the creation of critical mass (but see [11]). This can be effective, but can also be dangerous. Managers may be reluctant to put their own credibility on the line for a tool that you and they know will have significant problems. And use under management pressure may be markedly different from spontaneous use, as we saw in the Management 2 group.

Our experience with RVM in the workplace has convinced us that its combination of features has some potential to help distributed teams overcome the lack of context and absence of informal communication, two of the problems that make distributed work difficult. There are formidable adoption issues, however, that are likely to keep such tools out of many workplaces for some time. Notably, the characteristics that produce a critical mass of users can be local and highly variable – which increases the cost of deploying RVM-like tools by removing general approaches to achieving collective adoption.

REFERENCES

1. Bradner, E., Kellogg, W., & Erickson, T. (1999). The adoption and use of BABBLE: A field study of chat in the workplace. In *ECSCW* (pp. 139-157). Lyngby, Denmark: Center for Tele-Information.
2. Churchill, E. F. and S. Bly. It’s all in the words: Supporting work activities with lightweight tools. In *GROUP ’99*. 1999. Phoenix, AZ 40-49.
3. Grudin, J. “Groupware and Cooperative Work: Problems and Prospects.” In R.M. Baecker (ed.), *Readings in Groupware and CSCW*, San Francisco: Morgan Kaufman, pp. 97-105, 1993.
4. Grinter, R. E. and M. Eldridge (2001) “y do tngrs luv 2 txt msg?” In *ECSCW* Bonn, Germany: September 18-20
5. Herbsleb, J. D. and R. E. Grinter. Splitting the Organization and Integrating the Code: Conway’s Law Revisited. In *International Conference on Software Engineering*. Los Angeles, CA. 85-95, 1999.
6. Herbsleb, J.D., Mockus, A., Finholt, T.A., & Grinter, R.E. An Empirical Study of Global Software Development: Distance and Speed. In *International Conference on Software Engineering*. Toronto, Canada, May 15-18, 2001, pp. 81-90.
7. Herbsleb, J.D., Mockus, A., Finholt, T.A., & Grinter, R.E. Distance, Dependencies, and Delay in a Global Collaboration. *CSCW 2000* pp. 319-328, Philadelphia, PA, Dec. 2-7, 2000.
8. Jones, E.E. (1979). The rocky road from acts to dispositions. *American Psychologist*, 34, 107- 117.
9. Jones, E. E., & Nisbett, R. E. (1972). “The actor and the observer: Divergent perceptions of the causes of behavior.” In E. E. Jones et al (Eds.), *Attribution: Perceiving the causes of behavior*. Morristown: General Learning Press.
10. Kraut, R.E., Fish, R.S., Root, R.W., & Chalfonte, B.L. (1990). Informal communication in organizations: Form, function and technology. In S. Oskamp & S. Spacapan, (Eds.), *Human reactions to technology: Claremont Symposium on Applied Social Psychology* (pp. 145-199). Beverly Hills, CA: Sage Publications.
11. Mark, Gloria, Poltrock, Steven. Diffusion of a Collaborative Technology Across Distance. To appear, *Group 2000*.
12. Markus, M. (1990). Toward a “critical mass” theory of interactive media. In J. Fulk & C. Steinfield (Eds.), *Organizations and communication technology* (pp. 194-218). Newbury Park, CA: Sage.
13. Nardi, B., Whittaker, S., Bradner, E. Interaction and Outeraction: Instant Messaging in Action. In *CSCW 2000*, 79-88, 2000.
14. Rogers, Everett M., *Diffusion of Innovations*. New York: Free Press, 1995.
15. Tang, J., N. Yankelovich, J. Begole, et. al. ConNexus to Awareness: Extending Awareness to mobile Users. In *CHI 2001*. New York: ACM Press.
16. Whittaker, S., Frohlich, D., & Daly-Jones, O. (1994). Informal workplace communication: What is it like and how might we support it? *CHI ’94* (pp. 131-137). New York: ACM Press.