

What Is Chat Doing in the Workplace?

Mark Handel

School of Information, University of Michigan
3218 Computing Center Building
Ann Arbor, MI 48109
handel@umich.edu

James D. Herbsleb

School of Computer Science, Carnegie Mellon University
Pittsburgh, PA, USA 15213-3891
+1 412 268-8933
jherbsleb@acm.org

ABSTRACT

We report an empirical study of a synchronous messaging application with group-oriented functionality designed to support teams in the workplace. In particular, the tool supports group chat windows that allow members of a group to communicate with text that persists for about a day. We describe the experience of 6 globally-distributed work groups who used the tool over a period of 17 months. An analysis of use shows that the group functionality was used primarily for bursts of synchronous conversations and occasional asynchronous exchanges. The content was primarily focused on work tasks, and negotiating availability, with a smattering of non-work topics and humor. Nearly all groups were remarkably similar in the content of their group chat, although the research group chatted far more frequently than the others. We conclude with suggestions for future research, and a discussion of the place of team-oriented synchronous messaging tools in the workplace.

Keywords

Presence, awareness, instant messaging, chat, MUD, teams, groupware

INTRODUCTION

It has become commonplace for teams to work across distances – so common, in fact, that new terminology such as “virtual teams” and globalwork [25] are becoming widely accepted. Many economic factors drive this style of working [7]. Despite naïve claims in some quarters that working across distance is not essentially different from co-located work [6], careful analysis and examination of the evidence indicates otherwise [26].

One major difference is that geographically distributed teams experience a very substantial reduction in frequency of communication, particularly informal “corridor talk” [2, 20]. These communication issues appear to slow distributed work down very substantially [14], and give rise to a number of difficult coordination issues.

An awareness of what one’s distant colleagues are doing, and their availability for interaction, are key parts of

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.

CSCW’02, November 16–20, 2002, New Orleans, Louisiana, USA.
Copyright 2002 ACM 1-58113-560-2/02/0011...\$5.00

improving multi-site work. There have been a number of approaches to informal communication and awareness over distance [21, 32], including various applications of video [1, 10, 11, 24], open audio channels [16], and text [8]. Each has addressed some parts of the awareness and informal communication problem, but the results have been mixed.

E-mail, the form of computer-mediated communication in widest use today, has only limited potential for signaling awareness and availability. E-mail messages are generally delayed a minimum of several seconds, and often several minutes or longer, before being delivered. Further, there is little expectation that users will read and respond to e-mail immediately when received. E-mail is generally used as an asynchronous communication medium.

Increasingly, forms of messaging that are closer to synchronous, “real-time” communication are being brought into the workplace. On the heels of tremendous popularity among recreational users, synchronous messaging applications are beginning to show up at work. Previous research [5, 17, 23] has shown that synchronous messaging in the workplace has a number of uses, including opportunistic interactions, broadcasting of information or questions, and a “signaling” function in which people negotiate availability for other interactions. Use also tends to vary from conversation-like turn-taking to a more asynchronous mode in which minutes or hours may pass between conversational turns.

The relatively unstructured nature of synchronous messaging allows it to potentially support many different types of communication. At this point, we do not know the extent to which synchronous messaging is used for carrying out work, for discussing non-work topics, for negotiating availability, and so on. We also do not know if various groups tend to use synchronous messaging in similar ways, or if there are substantial differences, depending perhaps on work culture, the nature of the work itself, and individual personalities. Several important studies of instant messaging have looked primarily at use by research groups that developed the tools, (e.g., [31]), and it is not clear how typical this use is in a general workplace population. Other studies (e.g., [23]) have provided insights from interviews but have not looked comprehensively at actual message content. How synchronous messaging is used, and whether different groups use it in markedly different ways, are important

issues for such things as determining corporate policy for employee use of synchronous messaging, and for deciding on features that might enhance some uses at the expense of others.

Another interesting dimension in how synchronous messaging is used has to do with variation over the time of day [3]. Recent evidence indicates that interruptibility, for example [16] varies substantially, with managers being least interruptible in early morning and late afternoon. Since synchronous messaging can provide a lightweight way of ascertaining availability, one might expect that its use for negotiating availability might vary in similar ways. The topics discussed, e.g., work-related and recreational, might also reveal regular daily patterns.

Other synchronous or nearly-synchronous forms of messaging, such as workplace-oriented MUD rooms [8], are also finding favor in some workplaces, and serve a variety of functions, including presence awareness. Like e-mail, MUDs and chat rooms can preserve exchanges of messages for later reading and responding, hence have potential as asynchronous media. They may also, however, be used synchronously, to support an ongoing conversation, and to signal that one is present and/or available for other interactions.

The ability to support both asynchronous and synchronous communication is potentially a very important feature of chat rooms and MUDs in the workplace. Unlike e-mail, they can support conversations, with conversational turns on the order of seconds, not minutes or hours. But unlike spoken conversation, other users, not present at the time, can catch up with the conversation later. They may or may not choose to participate, but they have the benefit of additional context information about their co-workers.

Despite the potential for providing an additional, effective communication channel, movement of synchronous messaging (chat, IM, MUDs) into the workplace has raised serious concerns in some quarters. For example, an article in a prominent business publication said,

Messages that pop up on screen at an inopportune moment (sometimes from the next cubicle) are destroying workers' concentration. Thoughtless text scrawled and sent in haste can spark online arguments. And in some offices, the question of who is privileged enough to receive certain instant messages is creating the kind of tortured pecking order last seen in high school. [30]

In addition, synchronous messaging, like e-mail, has the potential for "flaming" behavior – rude, impulsive messages and expressions of extreme views [19]. Flaming behavior has been well-documented in a variety of electronic communication media, including interactive discussion systems [19], newsgroups [18] and group

support systems [29]. Although some of the research suggests that flaming activity has more to do with the social context than the underlying medium [18], the reduced social cues in synchronous messaging may contribute to anti-social behavior within the environment.

In this paper we present an empirical study of how a synchronous messaging application was used over a period of 17 months by a number of different teams in a geographically distributed corporate setting. We call the application Rear View Mirror (RVM) – the idea is that it helps users keep track of what is going on around them while their primary attention is elsewhere [4]. RVM combines familiar IM and presence awareness features with novel team-oriented functionality. In particular, users can create "groups," each with its own dedicated and persistent chat window.

Specifically, our empirical study addresses these research questions:

Using group chat. *To what extent did users create groups, join groups and use group chat? To what extent was group chat used for synchronous and for asynchronous communication?*

Content of Chat. *What did users talk about in group chat? To what extent did users "flame"?*

Daily Rhythms of Chat. *Is there evidence that chat was used in different ways, or for different purposes at different times of day?*

Differences Between Groups. *Did different groups use group chat in different ways? Did the research team that developed the tool use it very differently from other groups?*

Rear View Mirror

RVM is an implementation of an IM and presence awareness system with novel features designed to support teams. It is built around a client-server model, with the individual clients running on the user's machine. Transparently to users, the clients connect to two different servers. One server maintains presence information, group information, and implemented our privacy model. The other server was an IRC chat server that provided reliable message transport. In order to support the widest variety of platforms, the RVM client is written mainly in Java (a small amount of native code is necessary in order to provide time-out functionality based on keyboard and mouse activity).

The identity of other users is shown by means of an iconized picture (see Figure 1). Each person's presence state is indicated by the border color around the picture, and more detail can be obtained with a mouse-over. Users can change their status explicitly by selecting from a status menu, or by setting screen-saver-like timeouts that change status after a user-settable period of mouse and keyboard inactivity.

RVM also provides support for groups. Groups are similar to and can serve much the same purpose as “buddy lists” in other messaging applications, but there are some important differences. Buddy lists are defined independently by each user, whereas the membership for a group is the same for all users. If the group administrator adds a member to a group, for example, the group membership changes for all users. In addition to joining groups, one can, in true “buddy list” fashion, add other individuals to one’s surface, see when they are logged in, and exchange instant messages with them.

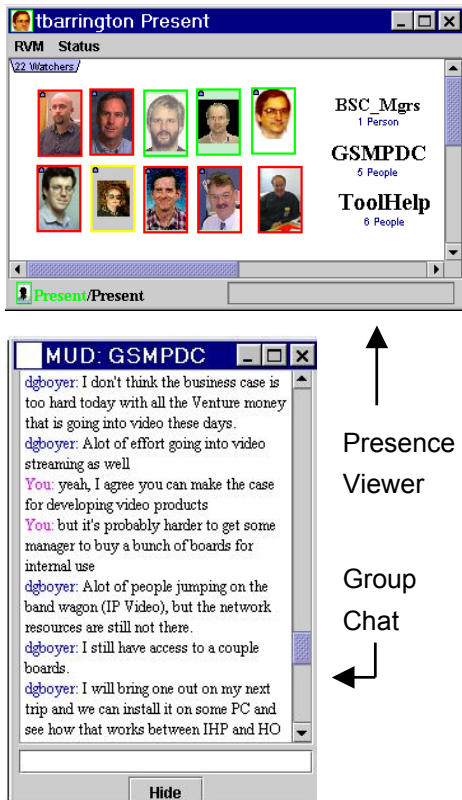


Figure 1. Rear View Mirror user interface.

Groups serve several purposes in RVM. For most of the period covered by this study, groups were central to our privacy model. A preliminary study indicated that many of our potential users were uneasy providing presence information that was available to everyone, but our experience was that forcing users to request individual permission from every person whose presence they wanted to observe was extremely clumsy [12]. Joining a group in RVM allows other members of that particular group to see one’s presence information, balancing privacy and convenience.

Groups in RVM are also associated with group chat windows that have persistent content. This provides a mechanism for teams to communicate with all other team members, and a way for a person to clearly separate communications associated with different teams. Each time a group member logs in, the group chat window opens, displaying (by default) the last days’ chat

messages. When a user logs in, he/she does not join a chat “context-free.” The user sees recent exchanges, and is able to respond with relevant comments. In addition, for users who are separated by time zones, persistent group chat helps to provide information about what happened at other sites before or after that user’s work hours. With respect to both the persistence and the group nature of chat, RVM has strong similarities to Babble [5].

Finally, we intended that users would, over time, absorb substantial context information from group chat. We suggested to users that they make preferential use of this public chat space even for (non-sensitive) messages targeted at some subset of the group, or an individual. The idea was to provide additional context, i.e., some general awareness of the questions, concerns, interests, and even personalities of other group members.

Any RVM user can create any number of groups. By default, the creator becomes the group’s administrator. Groups can either be permissive (anyone can join) or the group administrator can restrict membership to a specific list of users. A current list of all existing groups is available by selecting the appropriate menu item. The administrator can also determine how long chat in that group’s window will persist.

One of the trade-offs in an application like RVM is between the synchronous messaging functionality and the interruption to the user when new messages come in. We tried to strike a balance between these competing demands in RVM. New messages increment a counter displayed as the icon (in the upper left corner) of the chat window. Even when a chat window is minimized, a user can see the number of new messages since the last time the chat window was on top. In addition, certain events (people logging on and off, new private chats, and new messages) also can play sounds which can be turned on and off by the user.

In the next section we describe the teams that used the RVM group functionality, the data we collected, and how we analyzed the data.

SITES AND METHODS

RVM was originally developed for use within a software development organization primarily located in England and Germany, part of a large, multi-national telecommunications equipment manufacturer, working on software for wireless communications systems. Since the start of this effort, the wireless organization has grown to encompass sites in France and India as well, and has more than doubled in size. The organization has a wide range of activities, some which require very close collaboration across sites and others that are more loosely coupled. We report primarily on four teams in wireless that used RVM: a quality team, a knowledge management team, and two test teams. These teams ranged in size from four to 28 members.

In addition to the wireless organization, RVM had trials in other organizations: an architecture team in optical networking, and a research team. The research team was made up of members of a research project looking at the problem of geographically distributed software development. It included the authors, and was the source of the RVM software. The research team had members in England and Ireland, as well as four US sites. The architecture team had six members in four locations, all in the US.

Data collection

We gathered data about RVM usage through a variety of methods. The most detailed data came from log files generated through RVM usage. These log files recorded virtually every action that every user performed in RVM, including logging in and out; creating, joining, and leaving groups; changes to buddy lists; and changing status. Group chat was logged (with users’ knowledge), but IM was not, since we believed users would have a reasonable expectation that IM would be private, compared to the relatively public group chats.

When calculating the number of logins for a particular user or group, we count at most 1 login per user per day. On any given day, in other words, a given user either logged in or did not log in. We excluded multiple logins as a potential source of noise, since various network problems, rebooting, etc., cause (uninteresting) logins to show up in the data. In computing the number of chat messages, we considered the text preceding a return as one message. Like many synchronous messaging systems, RVM sends text to the server (and then to other clients involved in the chat) each time the user types a return.

In addition to the log files, we also performed about two dozen semi-structured interviews and had one small focus group session with six users. We interviewed both regular users of RVM (as evidenced by the usage logs) and people who stopped using RVM. These semi-structured interviews happened both face to face, and over the telephone, guided by an outline of questions and topics. The focus group was conducted face-to-face, with a group of active RVM users, all from the same team. These interviews and focus groups were primarily designed to get feedback about RVM usability, ideas for new features, and to get some insight about how RVM was being used in the field.

In this paper, we concentrate primarily on presenting the quantitative results from the log files. We made use of the interview and focus group data primarily to help us make sense of these quantitative results, and to confirm our interpretations of the data.

Categorizing group chat content

In order to understand how group chat was used by various groups, we extracted all group chat messages, for all groups, from our logs. In order to analyze these data,

we developed a set of mutually exclusive and exhaustive categories into which the messages were classified. In order to develop the category scheme, we constructed an initial set, using the literature about how IM is used in the workplace [23], and research about conversations among software developers [13, 27]. We also read all of the RVM chat messages, and used the results of interviews conducted with RVM users. The authors, of course, were also users of RVM for the entire period covered by the study.

The initial set of categories was used by the authors to independently code samples of 200 contiguous messages, all taken from groups other than the research group. Contiguous messages were used in order to be able to see the context in which the message was sent. This allowed us to code ambiguous responses (e.g. “Yes,” or “Ok, thanks.”) according to the context of the message. After each iteration of coding, all differences in how messages were coded were discussed, the definitions were refined, and categories that could not be coded reliably were dropped. After several iterations, we believed we could classify messages reliably, so we independently categorized another sample of 200 messages, and calculated a Cohen’s kappa [9] of .88, indicating near-perfect agreement [22]. The final set of categories and (summarized) definitions are shown in Table 1. Based on these categories, we coded all messages in the entire corpus of messages from the groups (4242 messages), which forms the basis of all analyses that follow.

Availability	Negotiating availability, either now or in the very near future (e.g., same day, maybe next day).
Non-work topics	Talk that has specific non-work content, e.g. cars, fishing, sports, etc, that are discussed for their own intrinsic interest.
Work	Anything that relates to specific work tasks, the process for those tasks, planning for them, general discussions of business or economics related to work, and discussions about the use of RVM itself.
Greeting	Hello, etc., also references to weather (e.g., “Hi, how’s the weather there”) or health (“Hi, how are you”) and so forth that are intended primarily as a greeting, not a question requiring an answer. “Closings” such as “Bye!” were also categorized as greetings.
Humor	Comments clearly intended to be primarily humorous, even if they have some specific work or non-work content.
Other	Anything that cannot be categorized elsewhere.

Table 1. Chat Content Categories

RESULTS

Using group chat. *To what extent did users create groups, join groups and use group chat? To what extent was group chat used for synchronous and for asynchronous communication?*

Users tended to join groups. For most users, the research team created a group for them before their first login. However, of active RVM users (i.e., those who logged in at least three times), 57% were members in at least two groups—joining a group other than their starting group, and a quarter of them were members in three or more groups. Among the top quarter of users by number of logins, group membership is even more pervasive, with 77% of these users in two or more groups, and 23% of these users in five or more groups. There were 30 groups we considered “active” that is, having at least one login every other day (i.e., at least one person who is a member of the group logged in to RVM). These groups had an average of 8 members, with the largest having 30 members. They tended to be long-lasting, with the average lifespan of a group being 199 days, with some groups lasting the entire length of the 17 month study period.

Extent of group chat use. Of the 30 active groups, at least one chat message was entered in the group chat window of 21 groups. This number, however, includes several groups that had only one message over their entire lifetime, and others that had only a few messages over a few days. We wanted to restrict our analysis to groups that appeared to make actual use of chat over some extended period of time. For that reason, we selected only groups that had at least 30 days of chat (the first recorded message and the last chat recorded message occurred at least 30 days apart), and had exchanged at least 100 total messages.

Group	Number of Locations	# Chatters (Total Members)	Total messages	Time period (days)
Quality	2	5 (5)	175	235
Research	6	17 (21)	2908	239
Test 1	3	5 (8)	146	224
Kn. Mgt.	3	7 (7)	124	46
Test 2	2	4 (4)	111	116
Arch.	4	4 (8)	104	113

Table 2. Chat activity in groups

Seven groups met these criteria. Six of them had membership that reflected all or part of a work group. The other group was created purely for recreational purposes, had only two members, and had no work-related content. Since we were primarily interested in studying group chat as an adjunct to work in distributed teams, we excluded the seventh group. The remaining six groups all had members from at least two sites, and

ranged in size from 8 (Arch.) to 21 (Research), with an average of 9. The chat from these 6 groups became our corpus of chat content. The daily numbers of messages in each group are shown in table 2. The research group is clearly an outlier.

Number of Messages and Logins per Day




Figure 2. Messages per day and logins per day in the Quality group

Patterns of use. By looking at the daily number of logins and chat messages for two typical groups, we can get a fairly clear view of the “bursty” nature of group chat. Figure 2 shows the daily logins and the daily number of chat messages for the Quality group. The number of logins is relatively constant, with periodic dips representing the weekends. On a typical day, 3 people log in. By comparison, the chat traffic is very bursty, showing several large spikes, and several periods with no traffic at all. The spikes for the Quality group range from one message up to 25.

Number of Messages and Logins per Day




Figure 3. Messages per day and logins per day in the Research group

Another example, the Research group, is presented in figure 3. In some respects, this group is rather different,

since it has many more logins and many more chat messages, ranging up to nearly 200 per day (we truncated the figure vertically). The same “burstiness” of chat is evident, however, in contrast to the relatively constant number of logins.

The overall pattern is one of short bursts of synchronous activity, as indicated by the analysis of intervals above. Most conversations are synchronous, i.e., do not span periods longer than a few minutes, but this is not always the case. For example:

Roger: Folks, I need to leave early today

Roger: My son had an accident and needs to go to the doctors to get sewed at the head.

Roger: I have to look after my daughter, so I gotta leave now.

Roger: Talk to you tomorrow

Frederick: Ok Roger, I really hope that it is not serious

[gap here of about 18 hours]

Roger: Hi Frederick, no, it's nothing serious. While playing in the garden he fell from the slide and [hit his head]

Frederick: Hi Roger, thank you for these good news

Persistence of chat allows for such follow-ups, so long as they do not exceed the persistence parameter.

Content of Chat. *What did users talk about in group chat? To what extent did users “flame”?*

In this section, we look in more detail at the content of chat activity. Table 3 shows the relative frequencies of the various content types. Overall, several characteristics stand out. First, like previous studies of individual chat, we found a portion (13%) of the messages concerned negotiating availability. This reinforces the view that chat is sometime used as a lightweight tool for seeking out and arranging heavier-weight interactions such as telephone calls and meetings. This was not, however, a very frequent use of the tool in our data.

There was a relatively small amount of conversation on non-work topics, which ranged over a variety of themes such as cars, weekend plans, and so on. Chat was also laced with occasional humor. Nearly all of the humor was job-related, and used exaggeration or other sorts of playfulness in the midst of working conversations. In one example, two people in a group were playfully vying for the top entry in a log of use for another tool. Whichever person logged in first got the “honor” of the top entry.

John: “Yes! By hook or by crook i’ll be first on this page heh heh!”

Dave: “Damn you Peterson, daaaaammnnn yoooooooouuu!”

Category	Num. Messages	Percent of Total
Work	2914	69%
Availability	533	13%
Greeting	288	7%
Non-work	143	3%
Humor	203	5%
Other	161	4%

Table 3. Chat Message content by category

Such “running gags” sometimes spanned many days, as when this follow-up occurred 11 days later:

Dave: “I’m the first! Eat my shorts, Peterson, eat my shooorts!!”

But non-work and humorous entries totaled only about 8% of the total.

The largest category of activity, however, was actually doing work. Based on this, we decided to further divide the “work” category. We started out by using Olson, et al.’s [27] categories, which were originally intended for real-time, face-to-face design meetings. We found they had to be modified substantially in order to apply them specifically to work-related content of synchronous messaging. We collapsed all three of the Olson categories representing the immediate task of the meeting (which in their case was design work, and included “issues,” “alternatives,” and “criteria”) into a single category representing the “technical” work the group performed (e.g., designing tests, gathering information about process exceptions, etc.), which we called “technical work.”

We dropped “walkthrough,” “goal,” “digression,” and “clarification,” since we never observed them within the “work” portion of our protocol. We retained project management and meeting management, and added “company,” since there was considerable discussion of company-wide news and issues, not directly related to the current project. The result was a set of five categories. As with the initial categorization, the authors independently coded random sets of messages, obtaining kappas of .81 to .85, indicating very high agreement. We then coded the entire corpus of 2914 messages (all of the messages categorized above as “work”). Summaries of our definitions, along with the results are in table 4.

Interestingly, we did not identify a single incident of “flaming,” or even any clear expressions of anger. There were some messages that feigned anger for humorous ends (as seen above), but we saw no examples of anything resembling flaming.

Category	% of Msgs	Description
Technical work	66%	Discussion of actual work at hand, e.g., carrying out tests; selecting and defining approaches to the work
Project management	21%	Planning the current project, project process issues, project status
Meeting management	8%	Planning and running meetings, including the current chat session, as well as future chat and face-to-face meetings, e.g. agenda topics, locations, technology support.
Company	4%	Issues affecting the entire company, rather than the current project, e.g., "You certainly know that discussions [concerning a merger] with [another company] are definitely stopped?"
Other	1%	Messages that fit in no other category.

Table 4. Work content sub-divisions

Daily Rhythms of Chat. *Is there evidence that chat was used in different ways, or for different purposes at different times of day?*

In order to determine how the content of chat evolved

Percent of each message type




Figure 4. Percentage of Availability and Greeting messages that occur, on average, in each 30-minute period of a weekday.

over the course of an average day, we identified the time zone for each user and aligned all messages according to the sender's local time. Based on this idealized day, we plotted the percentage of messages for each content category over a 30 minute window for the entire work day (4:00 to 20:00, i.e., 4:00 AM to 8:00 PM). Figures 4 and 5 shows the result. As can be seen from the graphs, there

are fairly clear trends over the course of a day. In the morning, most of the greetings and availability messages are exchanged, finding out who has gotten into the office, and how they are doing. Throughout the day, work messages stay relatively constant, then towards the end of the day, we see a sharp increase in humor and non-work related messages. In addition, in all categories, there is a slump in activity around the lunch hour.

Percent of each message type

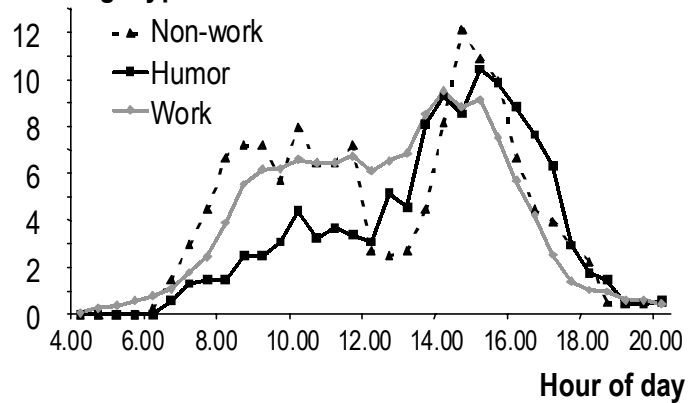


Figure 5. Percentage of Non-work, Humor, and Work messages that occur, on average, in each 30-minute period of a weekday.

These findings seem to correlate with those reported in [16], where managers reported being less interruptible in the early morning and late afternoon. If group chat can be used as a lightweight, non-intrusive way to see if someone is in fact interruptible, then one would expect heaviest use at those times when more intrusive means of signaling, such as phone calls or stopping by in person, are least appropriate. In fact, we do observe an increase in availability and greeting messages at those times. In the late afternoon, one begins to see a much larger proportion of work, non-work, and humor messages, as the signalling function seems to give way to more substantive (if not always work-related) content.

Some users made use of RVM group chat during teleconferences, but typically not as part of the teleconference, and typically not with the other people involved in the teleconference. Rather, group chat allowed them to "multitask," i.e., chat with others, typically about other things, during teleconferences. There were also instances where users were able to get information needed in a teleconference from team members not involved in the call, by using group chat to pose a question. Chat was also very often used as an adjunct to teleconferences, but NetMeeting chat was most often used for these sessions rather than RVM.

Differences Between Groups. *Did different groups use group chat in different ways? Did the research team that*

Type of Content

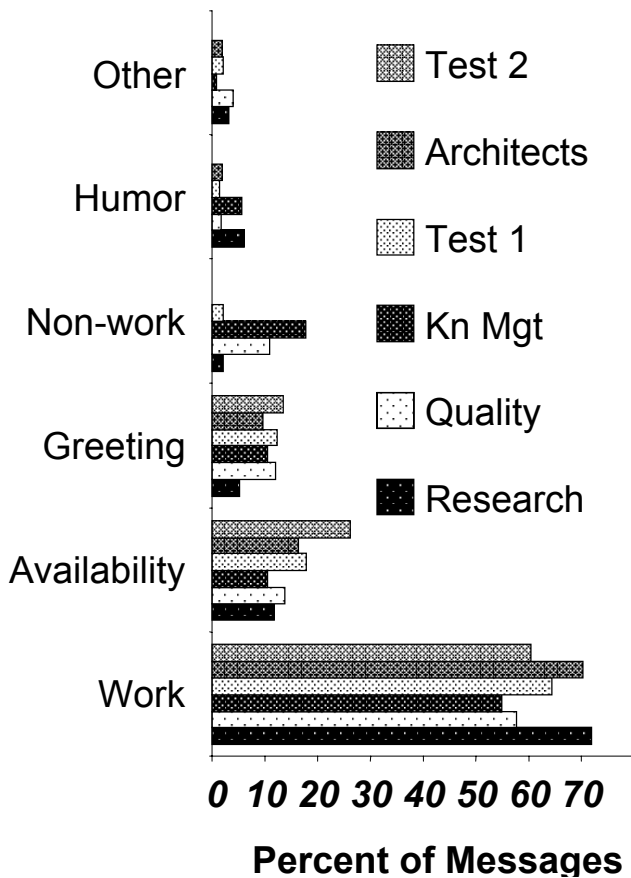


Figure 6. Percentage of messages in categories for each group.

developed the tool use it very differently from other groups?

Finally, we wished to compare the use of group chat by all the groups. These data are shown in Figure 6.

It is clear that these data, coded in this way, show very little difference between the research group and the other groups. In fact, the pairwise correlations between the values for the research group and the other groups range from .94 to .99. All the groups are remarkably similar under this coding scheme, with all pairwise correlations .88 or higher. The distribution of the research team's chat content was indistinguishable from the content of any of the other groups. Under different coding schemes, of course, the results might have been different.

In looking at the data for the sub-division of work (figure 7), there are still substantial similarities between all the groups, with pairwise correlations ranging between .70 and .99. For every group, talking directly about the technical content of their work was the most frequent type of message, although there was also considerable discussion about managing the project, managing meetings, and general company news.

Type of Content




Figure 7. Percentage of messages in each of the "work" sub-categories for each group.

DISCUSSION

The role of synchronous messaging in the workplace. For those who may still be skeptical about the potential value of synchronous messaging in the workplace [30, also 12], who see it as a way of gossiping or goofing off, these results should provide some ease from that worry. The content analysis shows that chat was used overwhelmingly for work discussions or for articulation work to coordinate projects and meetings, and to negotiate availability. The smattering of non-work topics and humor could be of concern only to the most draconian of employers.

It is not clear if this pattern is typical of all synchronous messaging, however. Compared to IM, group chat is relatively public. When conversing one on one, there may be a greater temptation to gossip, flame, etc., since there is a much higher expectation of privacy. In addition, RVM associated each message with the user's identity in unmistakable fashion, and there was no straightforward way to hide one's identity, or create a false identity. Furthermore, since each conversation took

place in a chat window associated specifically with some work group, continuity of identity makes each user accountable for his/her remarks [23, 28]. Particularly where anonymity is possible and is permitted, we suspect that the findings might well be different. It would be most interesting to conduct experiments to determine how various types of synchronous messaging are used in similar settings, to begin to sort out the effects of the tool features from the effects of settings [18].

There is another potential advantage of group chat over IM – some users who were familiar with IM reported that group chat tended to be less intrusive than IM. When an IM window pops up on one's screen, it draws one's attention, and people often feel obliged to respond. Group chat is different, in that the window is always there when one is logged on, but since many messages appear that are directed toward a group or other individuals, no one feels compelled to look at new messages immediately if it is inconvenient to do so. This provides another level of "plausible deniability" [23] or perhaps diffusion of responsibility, and allows users to free themselves from the obligation to interrupt other tasks and respond immediately. It carries the potential cost, however, of making group chat somewhat less synchronous than IM.

Group similarity. We found the similarity in message content among the various groups striking. We must be cautious in interpreting this finding, since there are many dimensions on which content could be sorted, and we looked only at one particular categorization. Further, all groups were involved in some way in software development, and hence shared some overall task similarity. Nevertheless, groups composed entirely of US residents were indistinguishable from mixed US-European groups, which looked just like all-European groups. Given the considerable cultural differences, we did not expect such similarities.

It would be most interesting to know how far this similarity extends, e.g., to IM, to e-mail, phone conversations, face-to-face conversations? Such analyses would help us to understand the role of synchronous messaging in the overall communication ecology of the workplace.

One interesting exception to this similarity is the sheer number of messages in the research group's chat, which was easily an order of magnitude greater than any other group. One possible interpretation is that given a particular level of use, a group with a stake in a tool is likely to use it in ways similar to how general users will use it. Their enthusiasm and commitment, however, lead them to make use of it much, much more. We think of this as the "more of the same" hypothesis, and if true, it indicates that modulo the general enthusiasm level, results from research groups may generalize well.

Temporal patterns. The extremely "bursty" nature of synchronous messaging that we observed may suggest that in the setting we studied, it takes some unusual event

to motivate a user to use the chat room. If the purpose of a synchronous messaging system is to increase overall communication, for example among distributed team members, RVM was only modestly successful. While six groups exchanged more than 100 messages each, we must assume that co-located groups talk informally much more frequently than that in a 17-month period. It may be that the usual indication of presence, i.e., that someone is logged on, does not itself provide sufficient conversation-starting material. We believe it is worth experimenting with additional cues, such as whether one has a meeting scheduled right now [31], or other automatically-displayed indicators of activity at other sites, in order to provide an idea or a question around which one might instigate a conversation.

Our results showing different trends for different types of content shed a bit of additional light on patterns that have been recognized in research on other synchronous messaging tools. Begole, et. al. [3] have convincingly suggested that various visualizations of activity in an IM system may provide users with helpful cues about other users' activities and interruptibility. The current work suggests that "busy," in the sense of system activity, may mean different things at different times of day. Against the backdrop of a fairly constant level of work, activity in the morning may indicate that people are arriving, seeing who is around, and negotiating availability. Busy in the late afternoon, however, may mean we're talking about movies, or interlacing our discussion with some joking around. Perhaps the late afternoon "busy" is more interruptible than the morning "busy" (but see [16], indicating managers don't like to be interrupted in late afternoon).

User interface for synchronous messaging. On the level of the user interface, one of the design trade-offs is how and how forcefully the application itself should interrupt the user, and in response to what events. Dismissing pop-up windows or recognizing alerting sounds all have the potential to unduly distract the user. On the other hand, without any alerts at all, the user may miss a conversation. Recent research on managers [16], for example, indicates that interruptions are not necessarily bad, and in fact constitute the very nature of management work. There does not seem to be a single right answer to these questions. Our user interviews revealed that preferences as to the degree of interruption depended greatly on the user; this was reflected in RVM where much of the behavior was user-selectable. Future work should look at new methods of alerting, as well as different behaviors based on the content and sender of messages.

REFERENCES

1. Abel, M. J., Experiences in an Exploratory Distributed Organization, in *Intellectual Teamwork: Social Foundations of Cooperative Work*, J. Galegher, R.E. Kraut, and C. Egidio, Editors, 1990, Hillsdale, NJ: Lawrence Erlbaum Associates, 489-510.

2. T. J. Allen, *Managing the Flow of Technology*. 1977, Cambridge, MA: MIT Press.
3. Begole, J., Tang, J.C., Smith, R.E., & Yankelovich, N. Work rhythms: Analyzing visualizations of awareness histories of distributed groups. In *CSCW'02*, 2002, New Orleans, LA.
4. Boyer, David, Cortes, M., & Handel, M. Presence Awareness Tools for Virtual Enterprises. In *OOPSLA 1998, Virtual Enterprise Workshop*. Available at: http://www.research.avayalabs.com/project/RVM/OOPSLA_workshop-final1.doc
5. Bradner, E., Kellogg, W.A., & Erickson, T. The adoption and use of 'Babble': A field study of chat in the workplace, *ECSCW'99*, 1999, Copenhagen, 139-158.
6. Cairncross, F. *The Death of Distance*: Boston, MA: Harvard Business School Press, 1997.
7. Carmel, E. *Global Software Teams*. Upper Saddle River, NJ: Prentice Hall, 1999.
8. 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.
9. Cohen, J. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1), 1960, 37- 46.
10. Dourish, P. and S. Bly. Portholes: Supporting Awareness in a Distributed Work Group. In *CHI'92*, 1992. Monterey, CA, 541-547.
11. Fish, R. S., et al. Evaluating Video as a Technology for Informal Communication. In *CHI'92*, 1992. Monterey, CA, 37-48.
12. Herbsleb, J.D., Atkins, D., Boyer, D.G., Handel, M., & Finholt, T.A., Introducing Instant Messaging and Chat into the Workplace. In *CHI'02*, 2002, Minneapolis, MN, 171-178.
13. Herbsleb, J. D., Klein, H., Olson, G. M., Brunner, H., Olson, J. S., & Harding, J. Object-oriented analysis and design in software project teams. *Human Computer Interaction*, 10(2-3), 1995, 249-292.
14. 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*, 2001, Toronto, Canada, May 15-18, 81-90.
15. Hindus, D., et al. Thunderwire: A field study of an audio-only media space. In *CSCW'96*. 1996, Boston, MA, 238-247.
16. Hudson, J., Christensen, J., Kellogg, W., & Erickson, T. 'I'd Be Overwhelmed, But it's Just One More Thing to Do:.' Availability and Interruption in Research Management, In *CHI'02*. 2002 Minneapolis, MN, 97-104.
17. Isaacs, E., Walendowski, A., & Ranganthan, D. Hubbub: A sound-enhanced mobile instant messenger that supports awareness and opportunistic interactions. In *CHI'02*, 2002 Minneapolis, MN, 179-186.
18. Kaynay, J. Contexts of Uninhibited Online Behavior: Flaming in Social Newsgroups on Usenet in *Journal of the American Society for Information Science* 49(12), 1998, 1135-1141.
19. Kiesler, S. and Sproull, L. Group Decision Making and Communication Technology in *Organizational Behavior and Human Decision Processes*, 52(1), 1992, 96-123.
20. Kraut, R. E., C. Egidio, & J. Galegher, Patterns of contact and communication in scientific research collaborations, in *Intellectual Teamwork: Social Foundations of Cooperative Work*, J. Galegher, R.E. Kraut, and C. Egidio, Editors, 1990, Hillsdale, NJ: Lawrence Erlbaum Associates, 149-172.
21. Kraut, R.E., Fish, R.S., Root, R.W., & Chalfonte, B.L. Informal communication in organizations: Form, function, and technology. In S. Oskamp & S. Spacapan (Eds.), *People's Reactions to Technology*, 1990, Newberry Park: Sage, 145-199.
22. Landis, J.R., & Koch, G.G. The measurement of observer agreement for categorical data. *Biometrics*, 33, 1977, 159-174.
23. Nardi, B., Whittaker and S., Bradner, E. Interaction and Outercation: Instant Messaging in Action. *CSCW'00*, 2000, Philadelphia, PA, 79-88.
24. Obata, A. and K. Sasaki. OfficeWalker: A virtual visiting system based on proxemics. *CSCW'98*. 1998, Seattle, WA, 1-10.
25. O'Hara-Devereaux, M. and Johansen, R. *Globalwork: Bridging Distance, Culture, and Time*. 1994, San Francisco: Jossey-Bass.
26. Olson, G. M. and J. S. Olson. Distance Matters. *Human-Computer Interaction*, 15(2/3), 2000, 139-179.
27. Olson, J.S., Olson, G.M., Storrorsten, M., & Carter, M. Groupwork up close: A comparison of the group design process with and without a simple group editor *ACM Transactions on Information Systems*, 11(4), 1993, 321-348.
28. Putnam, Robert D., *Bowling Alone: The Collapse and Revival of American Community*. 2000, New York: Touchstone Books.
29. Reinig, B, Briggs, R, & Nunamaker, J., "Flaming in the Electronic Classroom" in *Journal of Management Information Systems* 14(3), 1997, 45-59.
30. Slatalla, M., The office meeting that never ends. *New York Times*, Sept. 23, 1999.
31. Tang, John C., Yankelovich, N., Begole, James, et. al. ConNexus to Awarenex: Extending awareness to mobile users. *CHI'01*, 2001, Seattle, WA, 221-228.
32. Whittaker, S., Frohlich, D., & Daly-Jones, O. Informal workplace communication: What is it like and how might we support it? *CHI'94*, 1994, Boston, MA, 131-137.