

Hubbub: A sound-enhanced mobile instant messenger that supports awareness and opportunistic interactions

Ellen Isaacs, Alan Walendowski, & Dipti Ranganthan

AT&T Labs
75 Willow Road
Menlo Park, CA 94002
{ellen, walendo, dipti@research.att.com}

ABSTRACT

There have been many attempts to support awareness and lightweight interactions using video and audio, but few have been built on widely available infrastructure. Text-based systems have become more popular, but few support awareness, opportunistic conversations, and mobility, three important elements of helping distributed groups coordinate. We built on the popularity of text-based Instant Messengers (IM) by building a mobile IM called Hubbub that tries to provide all three, notably through the use of musical sounds. In a 5.5-month use study, we found that Hubbub helped people feel connected to others in different locations and supported opportunistic interactions. Sound was a useful cue for helping people feel connected, although some found it annoying. It was more important to support graceful transitions between multiple fixed locations than to support wireless access, although both were useful.

Keywords:

Instant messaging, awareness, sound instant messages, sound IDs, earcons, wireless, Palm, mobile computing.

INTRODUCTION

Over the past fifteen years, researchers have been experimenting with techniques that enable distributed work groups to maintain a background awareness of one another. This work is based on evidence that awareness plays a key role in keeping co-located groups coordinated and motivated. One of the ways it does so is by enabling opportunistic interactions, those that occur by chance when people happen to see one another. Those interactions are often the mechanism by which people learn relevant news and get updates on current activities, and by which people develop personal relationships with their colleagues, which in turn motivate them in their work [11, 12, 13, 20]. Kraut et al. [12] found that teams that cannot make use of opportunistic and spontaneous interactions take longer and produce lower-quality results, even if they have just as many or planned or intended interactions (email, phone calls, meetings).

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 2001 ACM 1-58113-453-3/02/0004.. \$5.00

Based on these studies and others, many researchers have built systems to support awareness and opportunistic virtual interactions among distributed groups. The presumption behind the early systems was that high-bandwidth connections are needed to support awareness, typically through video and high-quality audio [2, 6, 9, 15, 18]. Kraut et al. [12] explained, "a technology for informal interaction must support both audio and video communication if we are to successfully simulate chance encounters." Unfortunately, those who have worked with such media find it non-trivial to overcome the associated installation, interoperability, and network problems, even when supporting an internal group with a homogenous computer environment. And yet, as Kraut et al. explain, "If the costs of using a communication source are too high... the user will be either unable or unwilling to use that system for the brief, frequent, spontaneous conversations that are characteristic of informal communication."

Partly for that reason, more recent research has focused on text-based systems, which are easier to set up. For example, a chat-based system called Babble presented abstract representations of users' availability based on computer activity, which supported awareness and triggered opportunistic interactions [5]. Churchill & Bly reported on a workplace MUD-based system that also supported opportunistic exchanges among a distributed group [4]. Meanwhile, some researchers have pointed out that people move about during the day and work from multiple fixed locations (e.g. home, work, school), so awareness systems need to allow people to stay connected as they move around [1, 14]. A text-based system called ConNexus nicely integrates multiple mobile devices and provides subtle awareness cues. This system mainly supports intended rather than opportunistic interactions, especially on the mobile devices [19].

In the marketplace, a few related products have become popular. Instant Messaging (IM), a text-based system for lightweight messages, took off a few years ago, initially among teenagers and increasingly among business users. Some IM systems are starting to offer versions that run on wireless devices, but their uptake has been slow. More popular, at least in Europe and Japan and especially among teenagers, is text-based messaging on cell phones (SMS) [8]. One reason IM and SMS became so popular is that they are trivial to install and set up, and they are lightweight to use, i.e. the effort to initiate and retrieve a message and the time to transfer it are minimal.

These text-based systems enable people to have lightweight interactions while mobile, but IM has only primitive support for awareness and SMS has none. IMs allow users to create “buddy lists” of contacts. Because messages can’t be sent to those who are offline, IMs indicate which buddies are online. Some IMs play a sound when someone goes online or offline, but you can tell only that *someone* logged on or off, but not who unless you look and figure out who is new or missing from the list. Also, many people stay logged in for days or weeks at a time, so seeing that someone is online indicates little about their availability. Some IMs indicate that someone is “active,” but the grace period before becoming idle is often long, so seeing that someone is active does not mean they are there. Nardi et al. found that IMs frequently start with “Are you there?”, indicating that the cues were not sufficient [16].

Also, existing IMs are not well tuned to people who work from multiple locations. Although you can run them from many locations, you can stay logged into only one at a time, and it’s easy to miss messages sent to one location while you are in transit to another.

We set out to build a text-based system that would (a) provide awareness information among distributed groups, (b) encourage opportunistic conversations, (c) allow people to stay connected as they move among multiple fixed locations and as they move about, and (d) be readily available and easily installed. That system is called Hubbub (available from www.HubbubMe.com). It is an instant messenger that runs on a wireless Palm and a PC, and it makes extensive use of sound as well as visual cues to provide background awareness information without requiring the user’s explicit attention.

Others have explored the use of sounds for awareness because it is effective at calling one’s attention to an event or a change in state [7, 15]. Visual cues are better for providing an overview of information but only when a person chooses to attend to them. We tried to combine these advantages to give people a snapshot of activity while also providing a sense of the comings and goings of their buddies. We hoped the auditory cues would support the opportunistic exchanges that are missing among people who don’t work near one another.

We also wanted to see if we could use sounds for an even lighter-weight method of making contact than text messaging. Our intuition was that remote colleagues, family members, or friends might like to send each other quick “hello’s” with little effort, and that sounds might be effective for enabling people to request text interactions unobtrusively. We also thought a sound-based message system might be more fun.

This paper first describes Hubbub’s user interface, and especially the sound UI. Then we describe a long-term use study and report our findings about the awareness cues, how well they led to opportunistic conversations, the use of sounds, and the value of supporting mobility. Finally, we discuss lessons learned about supporting awareness with a text-based tool using non-speech audio.

HUBBUB

Figures 1 and 2 show Hubbub on the Palm and the PC.

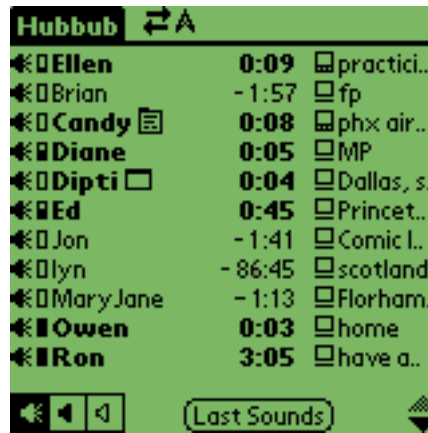


Figure 1. Hubbub on the Palm.

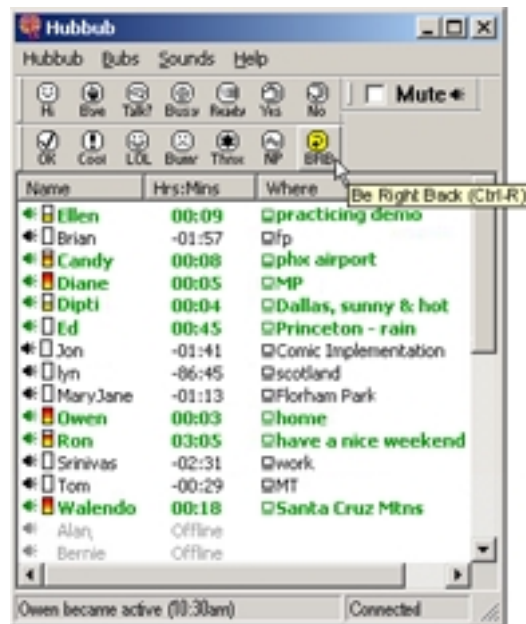


Figure 2. Hubbub on the PC.

Each user’s “bubs” (or buddies) is listed in bold (and green on the PC) if they are “active,” or black and non-bold if they are “idle.” Bubs who are offline are shown at the bottom of the list (and in gray text on the PC), with the label Offline as their location. Active is defined as having tapped, typed, or clicked in any application within the last five minutes. In addition to having a Hubbub name, each user has a “Sound ID,” a short string of notes from a song, that represents them. Each time a bub becomes active after being idle or offline, a sound plays (two notes, rising pitch) indicating that someone became active, followed by that bub’s Sound ID, so that people can tell who became active without looking. Users can turn on or off each bub’s “activity sounds” by toggling the sound icon to the far left of each bub’s name. In addition, users can easily Mute all Hubbub sounds.

To the left of each name is an “activity meter,” indicating each bub’s level of activity within the last 15 seconds. It is empty if the person has not used an input device within that time, or it may be low, medium, or high, depending on how actively they have been tapping, clicking, or typing. If any bub’s meter is

not empty, you can be reasonably sure that bub is at their computer. The column to the right of the bub's name indicates how long they have been either active or idle. The next column indicates what type of device the bub is on (Palm or PC) and a short message entered by that bub, usually indicating their location or status.

Hubbub includes the basic text message features of all IMs. Figure 3 shows a message on the PC. On the Palm, the user taps a bub's name to see the Bub Screen. They type a message at the bottom of that screen, and when they send or receive a message, the Text Screen appears (Figure 4).




Figure 3. Text Message window on PC Hubbub.

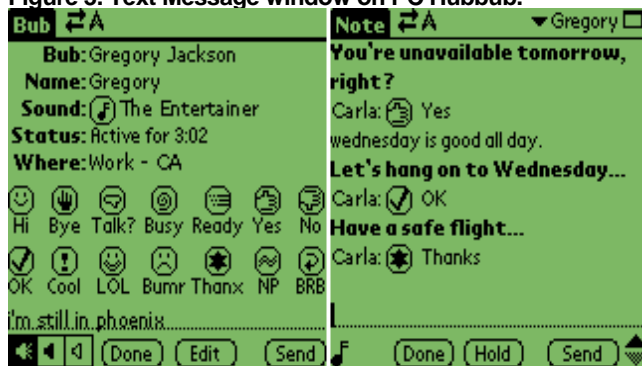


Figure 4. Bub and Text Message screens on the Palm.

When a message arrives, the recipient hears the Sound ID of the sender, followed by a sound associated with incoming messages. Again, hearing the Sound ID lets people identify the sender without looking. Hubbub also visualizes the other person's status during a conversation. An icon shows whether the other person is typing in the window, focused in the window, or not focused in the window, which roughly corresponds to speaking, listening or not attending. The icon also shows the other person's device, which is especially important if they are on a Palm, since the connection is slower and they are likely to type more slowly and with more typos.

In addition to text messages, users can also send each other *Sound Instant Messages* (SIMs). This is a novel concept that is essentially an earcon with an associated meaning, e.g. "Hi," "Thanks," "Cool." Figures 3 & 4 show Hubbub's sound messages. (LOL stands for Laughing Out Loud, BRB is Be Right Back, and NP is No Problem.) When a user sends a sound message, their Sound ID plays, followed by that sound message, so the recipient can hear the message and interpret it. For example, if Dipti sends a "hi" sound message to Alan, her Sound ID plays, followed by the three notes that mean "hi." When the sound arrives, a footer (PC) or header (Palm) appears indicating who sent which message, to help users learn the sounds. Sound messages can also be incorporated into text exchanges. Figures 3 and 4 show several sound messages integrated into a real conversation. Each time one is sent, the associated sound plays and the icon and label appear in the conversation.

To preserve privacy, Hubbub membership is reciprocal and requires permission. If Dipti wants to add Alan to her list, she needs his permission and, if granted, she is also added to his list. If Alan wants to renege his permission, he can remove Dipti from his list and he is removed from hers. This policy prevents someone from spamming a Hubbub user without the recipient knowing who sent the messages or being able to prevent them from sending more. Users can also block each of their bubs' from seeing their activity (time active and activity meter) and/or location text. We believe this approach provides better security than existing IMs because a pest cannot simply create a new account if they are blocked from sending messages to someone. Furthermore, giving each user the ability to block information about activity or location provides finer-grained privacy protection.

Sounds

We used earcons (musical sounds) for the SIMs and Sound IDs. We chose 14 messages and designed sounds that conveyed the attitude underlying each message. We tested the sounds with volunteers in our office unrelated to the project, iterating on them until they seemed appropriate and learnable. The Sound IDs, as mentioned, are short riffs from songs, usually ones with strong melodies that can be identified in a few notes. To start, we looked for songs that were well known and distinctive, and then created new Sound IDs as requested by study participants.

To make the Sound IDs distinguishable from the SIMs, we made all SIMs shorter and two octaves higher than the Sound IDs [3]. SIMs are typically a half second and Sound IDs are between 2 and 3 seconds. Although both these cues may be helpful in telling them apart, in practice the main distinguishing factor seems to be that Sound IDs are more melodic than SIMs. We used MIDI-0 sounds, which are quite primitive, because they are the only format available on the Palm. When played on a PC, these sounds are played on a piano with only one note at a time (no chords).

Mobility

Hubbub supports mobility in two ways. First, it supports use while on the road by running on a wireless Palm (a Palm V or

III with a Minstrel modem using CDPD) in addition to any Windows machine. The architecture is designed to handle low-bandwidth clients that may move in and out of coverage. Specifically, the message protocol is based on UDP to keep the connection overhead small, and it includes a mechanism to make sure messages arrive reliably. Text messages are sent through the server to the partner's client, which replies with an acknowledgement (ack). Messages that are not acknowledged are resent (every 3 seconds, up to 15 times). If a message had already been received, the receiving client resends an ack without redisplaying the message. The user interface indicates to users when/if the message arrives, so they don't have false impressions of who has seen which messages.

Second, Hubbub allows people to stay logged in from as many locations as they like. As they move about, messages "find them" wherever they are. Messages are sent to whichever client is active; if all clients are idle, they are sent to all clients. If someone just misses a message on an active client and then becomes active on another one, it is resent to the newly active client. In practice we found that messages are rarely missed with this system. See [10] for more information about Hubbub's design and architecture.

USE STUDY

Method

We conducted a use study of Hubbub over 5.5 months as part of our iterative design process. When Hubbub had enough functionality to be useful, we gave it to 13 people at three sites within AT&T Labs, two in New Jersey and one in California. Over the next few months, we released five versions, each time adding new features and refining existing ones, and each time adding users. By the end, the study included 28 regular users, 9 of whom worked from home frequently or exclusively or traveled often for business. We logged all Hubbub activity (with permission) and surveyed users after one week of use and again after two or four months. We interviewed many of the users after either three weeks or two months of use, and we had many informal conversations with them. Three months into the study, we videotaped one person in her office for two days to capture realistic, spontaneous use. To gather further feedback, we included a "Send Feedback" menu item in the PC client that started an email to the Hubbub team.

Usage

During the study, 32 people started Hubbub at least once.¹ Four people never ran it after they tried it for a day or so, so most of our data come from 28 people. Our users averaged 15.0 bubs in their list, but this number is inflated because we initially added everyone in the study to everyone's bub list. A better measure is the average number of bubs with which each person had at least one conversation, which was 7.0.

We recorded 1,726 conversations among 135 different pairs of participants. (Since people had ongoing exchanges during the day, we defined a conversation as a sequence of messages with no two messages separated by more than five minutes.) Across both devices, the participants averaged 1.6 conversations per day for those days in which they ran Hubbub and used their computer at least once. Good data about usage of other IMs is hard to find, but a Forrester report based on self-report data indicated that IM users average 3.2 conversations per day and have an average of 14.5 buddies [17]. Apparently, our participants were less regular users than the IM population overall, perhaps because the set of users was smaller, or because they were less inclined to use it. Few of our users had used one of the free IM products before participating in the study.

Awareness Cues

We evaluate how effectively Hubbub supported awareness by looking at our quantitative log and survey data combined with our qualitative interviews. The logs show that, on average, each person listened to the activity sounds for 61.4% of their bubs', and in most cases all of those with whom they had interactions. Five people muted all their bubs, nine muted none, and 14 listened to only some. Since people easily could have turned everyone's sounds off, the fact that most people did not tells us that they found them useful enough to keep listening to them.

However, the survey data are less conclusive. When asked how "useful" the activity sounds were after one week of use, our users rated them 2.5 (1=not at all useful, 5=very useful). After using Hubbub for two to four months, that rating increased to 2.9. When asked how annoying the sounds were, the participants rated them 3.0 after one week, and then 2.5 after several months. So while the sounds at first seemed more annoying than useful, that tradeoff changed over time as people got used to them and learned their value. Still, the usefulness ratings were unimpressive, even after months of use, especially considering that most people did not turn them off. Given that people tend to undervalue informal interactions as useful for getting work done [12, 13], people are probably even less likely to consider awareness of colleagues "useful."

The interviews further support this conclusion. People told us they did enjoy the sense of awareness provided by the sounds, but that it was subtle and they couldn't attribute a concrete benefit to that feeling. For example, one person who worked exclusively from home said, "Sometimes I just like the sounds, just hearing it. It gives me kind of this state of feeling that there's this group and they're working together and you know things are happening. It's almost like a sound screen saver."

People mentioned feeling closer to those who were on the opposite coast. One manager in New Jersey said about a colleague in California that by seeing her in Hubbub, "I know she's there. In some respects if she's not up here [in his bub list] and I don't see her, I don't think about her, I mean, for better or worse. But when she's there, and you can see she's doing stuff, she's sort of closer to the front of your consciousness." Interestingly, he never exchanged a

¹ The authors used Hubbub to interact with some participants during the study. Unless noted, we exclude our data in usage measures. For conversation data, we include conversations between participants and one of us, but not those between pairs of us.

conversation with her throughout the study, and yet he clearly liked seeing her in his Hubbub list.

One New Jersey manager said she muted the activity sounds of people in nearby offices, but kept them on for colleagues in another part of the building and other sites, even those she didn't work with, just to feel more connected to them. Another person found the sounds useful for connecting her to people in another wing of the building. One person in California said "I think I glance at Hubbub about 5 or 6 times a day just to see who is on... Friday is often a slow day around here and it's nice to get a sense of virtual companionship." Another person said simply, "I miss people when they are idle. :-)"

Several people complained that the sounds were annoying. One person found them grating because of their low quality. Others complained that the Sound IDs didn't do justice to the real songs and sometimes were hard to recognize. We used low quality sounds so that the same sound could play on both devices – and so that eventually people could generate their own sounds that would play on both devices. Perhaps we made the wrong tradeoff. Even with higher quality sounds, though, we think some people still would have found them irritating. We heard from a non-user that "the thing that makes Hubbub cool is the thing that makes it annoying."

Several other awareness cues were widely appreciated. One was the location feature. We had started by allowing only four characters for the location so that it would fit on the small Palm screen. However, people quickly started to use it in creative ways and asked for more space, so we expanded it (and modified the Palm UD). It became a sort of community bulletin board in which people not only posted locations but also notified others of the weather in their area ("19 in. of snow so far"), events ("UCLA by 8 at the half"), current status ("conf call – mute"), unusual locations ("København," "phx airport"), or just to amuse ("in my coding jimmies" from someone who worked from home).

People also especially appreciated the ability to tell whether the other person was typing, focused, or not focused in their message window. Many people find it difficult to coordinate conversations in IM, and this feature reduced that problem. As one person explained, "If I was about to type something and I saw that you were typing, I would stop and wait, because I assume that meant you were going to send me something." Also, a few people came to rely on the activity meter, which indicates whether the other person is really at their computer and how busy they are, but others rarely used it. Those who did said it was critical for deciding when to contact someone and how to "approach" them.

The cost of all these awareness cues is that they reveal information about each user's activity. It wasn't until our fourth iteration that we implemented the ability to block others' access to one's own activity and location information, although from the beginning, membership was required to be symmetric, so everyone knew who could monitor them by looking at their bub list. Our interviews turned up little evidence of privacy concerns, even though we focused on the issue. One person said this when asked if he felt his

management was evaluating his Hubbub use: "Yeah, it's an interesting question, right? And knowing the two of them [his boss and boss' boss], number one, I don't think they would. Number two, if they would, then I don't want to work here."

No doubt, our finding is related to the fact that our study was conducted in a collegial environment where trust is high and people are not expected to be constantly available on their computers (or even mobile devices). However, this organization is also very hierarchical and level conscious. It is possible that the three administrative assistants would feel evaluated by their managers. Instead, when asked if they had any concerns, one said, "We have so many responsibilities other than sitting at a keyboard. There's so much going on that if I'm idle from my computer, I don't think anybody's making a judgment. I'm not too concerned about that... If [my boss] felt like I wasn't doing what I was supposed to, I don't think he would need to monitor me in that way. He would know from other things." Another said, "If anyone said anything to me [about not working hard enough], I would laugh." The third said she felt comfortable that her boss knew how hard she worked.

Perhaps not surprisingly, none of the managers said they used Hubbub to check up on people or that they made judgments about others' productivity based on it. Of course, they may have made subtle judgments without realizing it or admitting it to themselves. We found the comments of one person particularly apt. He said that Hubbub reflects whatever level of trust already exists between co-workers. If a manager is suspicious that an employee is not working hard, they can find evidence in Hubbub. If they believe the person is a stellar employee, they interpret the information favorably.

Interestingly, several people quickly added that they thought that *other* people in *other* workplaces would feel monitored. Some mentioned the administrative assistants in their workplace, who themselves expressed no concerns. People seem to believe that increased awareness is a risk for others, even though they found the lightweight access to others worth the tradeoff themselves, as long as their bub list included people they knew and trusted.

Once we implemented the administration and privacy features, some people removed bubs they never interacted with, but no one made use of the privacy controls. We even blocked access to our information to model the feature to others, and yet still no one else did so. We have found that privacy features are critical for giving people peace of mind that they *can* control access, even though they rarely do [18].

Opportunistic Interactions

An important purpose of the awareness sounds was to afford opportunistic conversations. In the interviews, many people said they occasionally contacted people after hearing them become active. One person explained that activity sounds "are a trigger to tell someone something you've been meaning to ask." The logs supported this statement.

Figure 5 shows the percentage of the conversations initiated within the first 2.5 minutes after the recipient went active (for

the 1,349 conversations for which we captured the active status of the recipient). The largest percentage of conversations happened within a minute after the recipient became active, just after that person's Sound ID played.

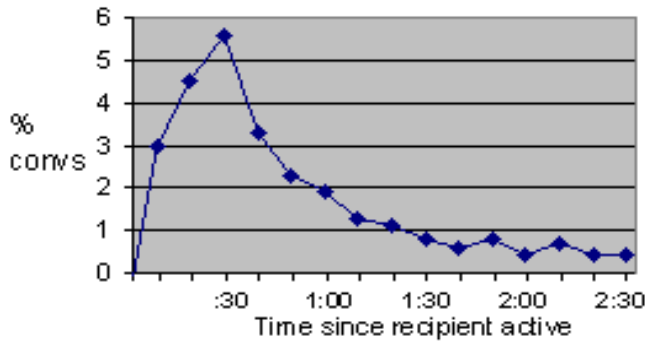


Figure 5. Percentage of conversations started relative to the time the recipient became active.

This graph indicates that Hubbub did afford opportunistic conversations, those triggered by one person becoming aware of another at a time when they had something to discuss.

We also saw a small number of cases (22) where two people exchanged “Hi” sound messages with no other message. Certainly this is a tiny amount, but it gives us some indication that people did use Hubbub simply to make a connection. Other interactions appeared to be triggered by someone’s unusual location. A colleague would send a message to inquire, similar to seeing someone wearing an unusual outfit and commenting on it.

Role of Sound

In addition to the activity sounds, Hubbub included two other types of sounds: Sound Instant Messages and Sound IDs. Both seemed to be well received.

Sound Instant Messages. The logs show that SIMs were actively used, with 6,731 sent during the study. Of the 1,736 conversations, 68.8% included at least one SIM. The average number per conversation was 2.3, making up 15% of the conversational turns. To determine how broadly SIMs were adopted, we looked at the percentage of each person’s conversations in which they *sent* a sound message. This ratio averaged 39.4% with a bell-curved distribution ranging from 4.3% to 84.0%. So on average, people sent at least one SIM in about two-fifths of their interactions.

Message	%	Message	%	Message	%
Hi	22.3	LOL	6.2	No	1.5
Bye	22.0	Cool	5.7	BRB	1.2
OK	14.1	No problem	3.9	Busy	1.1
Thanks	8.2	Bummer	4.2	Ready	0.5
Yes	7.7	Talk?	1.5		

Table 1. Usage of SIMs, out of 6,731 instances.

Table 1 shows how frequently each SIM was used. “Hi” and “Bye” were by far the most commonly used, followed by “OK.” In the interviews, people said they liked that “Bye” gave them an easy way to end conversations, which can be

difficult to coordinate in other instant messengers. We were surprised that other possible openings (Talk?, Ready) were used much less often. It seems that the most popular SIMs are part of ritualized exchanges (Thanks-No problem) or common responses (OK, Cool, Bummer). We were especially surprised that “Busy” was not used more, since it is easier to click a button than to type a message when you are focused elsewhere. It’s also interesting that “No” was used much less than “Yes.”

Again, despite the relatively high usage rate, people rated the SIMs only moderately useful. After one week, people rated their “usefulness” as 2.4 (out of 5), which increased to 2.9 after 2-4 months. When asked how annoying they were, the participants rated them at 1.7, increasing to 2.5 after a few months. So even though people came to like them better, they also found them more annoying. However, several people made a point of telling us that the SIMs were more “fun” than “useful,” so the “useful” question may have underestimated people’s appreciation of the sound messages. People reported learning to identify about 3-5 of the sound messages; we did some spot-checks during the interviews, which confirmed this. Usually people learned Hi, Bye, and OK, plus one or two of the more distinctive sounds.

Sound IDs. The Sound IDs accompanied both the activity sounds and the SIMs. On the whole, people learned the Sound IDs of the people they interacted with, but not peripheral contacts. (Remember that everyone was given a list of bubs rather than choosing their own.) After two months of use, people reported learning an average of 6 people’s Sound IDs with at least 75% accuracy, and 5 with 100% accuracy. (Everyone said they had learned their own Sound ID with 100% accuracy, and everyone we tested had done so.) Given that the participants listened to an average of 9.3 people’s sounds, they report learning about 2/3 of those they listened to. Again, our spot-checks found this estimate to be accurate.

People liked that Sound IDs gave Hubbub and its users a personality. They were interested in other people’s choices of Sound IDs and cared about their own. People were creative in expressing themselves through their Sound ID, and several people asked to change theirs after using Hubbub a few weeks. For example, two people requested songs with their name in the lyrics. One administrative assistant expressed her feelings about her job by choosing “Mission Impossible.” Another admin chose her song because it sounded happy and upbeat. One of the early users had been given a song at random before we started the study, but she decided to keep it. The song was “Aqualung” by Jethro Tull, a rock-and-roll song that didn’t match many people’s notions of her personality. She said she enjoyed making people question their assumptions about her!

While developing the concept for Hubbub, we were asked why we didn’t use voice for Sound IDs or for sound messages. Aside from the Palm’s limited sound capability, we chose earcons because they are less distracting than disembodied voice messages. Furthermore, we thought earcons would offer more privacy, since they are not immediately interpretable by bystanders. Only those who are part of the community can

interpret who had just sent someone a message. The tradeoff, of course, is that our sounds had to be learned.

Mobile Use

Hubbub supported mobility among multiple fixed locations and while in transit. The former was used more extensively because more people worked from several locations than had wireless Palms. Twelve (42.8%) of the 28 participants ran it from more than one location, whereas only three (10.7%) used the Palm regularly. Of the 185 Palm conversations (between 37 different pairs of people), 171 included at least one of those three active Palm users.

By allowing people to stay logged into multiple clients, Hubbub made it easy to move from place to place without having to remember to log in and update the location. Each time someone became active on a new device, their hubs automatically saw their new location and any incoming messages were sent there. This design enabled people to contact each other without keeping track of who was working where when. It also facilitated weekend, morning, and evening interactions among those who ran Hubbub at home, which helped bridge time zone differences.

For example, one person in New Jersey said she liked hearing activity in California when she got home at the end of the day, since most of the remote colleagues were still at work. She especially liked hearing when others were working on their computers over the weekend, a common activity for her. Another person said this after having used Hubbub for about two months at work: "I have used Hubbub at home recently for first time (new laptop), and it's great --there it gives me the sense that I am not alone."

	Palm	PC (Palm users)	PC (all)	Total
Conversations	185	969	1,541	1,726
Avg turns	19.6	17.7	15.5	15.9
Avg SIMs	4.1	2.3**	2.1**	2.3
Avg words/turn	6.4	10.0**	9.6**	9.2
Avg turn gap	0:23	0:22	0:22	0:18
Avg duration	5:44	5:54	3:23*	3:44

Table 2. Attributes of Hubbub conversations. Bold indicates a significant difference from Palm conversations (*= $p < .02$, **= $p < .001$)

As for Palm use, Table 2 shows characteristics of conversations when one or both people were on a Palm, PC-only conversations among those who had Palm conversations, all PC-only conversations, and all conversations. (The second column helps factor out effects attributable to the conversation styles of the Palm users.)

When one person was on a Palm, they exchanged no fewer messages (turns) and spent no less time conversing than when both were on PCs. Given that it is slower to enter text on a Palm than on a keyboard, it is surprising that Palm users did not keep their interactions short. The difference was that Palm users sent terser messages and made more use of the SIMs,

which are quick and easy to send. As a result, the gap between turns was no different when one person was on a Palm vs. when both were on PCs. Apparently, being on a Palm did not keep people from having substantial conversations. The fact that all PC conversations were shorter than Palm users' PC conversations indicates that people who had Palms had longer conversations overall. Perhaps having a Palm and chatting longer are both consequences of being a heavy adopter.

The few who got into the habit of using Hubbub on the Palm found it convenient and used it in ways they would not have used a phone or pager. One used it mainly when traveling, particularly when he reached his destination. The other two used it in meetings, while in transit, at travel destinations, and so on. One person used it to check in every time she arrived at an airport. For the three active Palm users, conversations on the Palm made up 20%, 10% and 8% of their conversations. Here are some examples of interesting Palm uses:

- While one person (Carol) was in a meeting, a question arose about a topic another Hubbub user (Lynn) was working on. Using her Palm, Carol sent a message to Lynn asking her the question. After a few messages to clarify the question, Lynn gave her answer, which Carol passed on to those in the meeting. She said they were pleasantly surprised that they could get the answer so quickly and without disrupting the meeting with a phone call.
- Lynn was stranded at the airport when her flight was delayed 4.5 hours. She used Hubbub to chat with her husband at home for about 3 hours, staying calm while other passengers stewed in frustration. They would not have spoken on the phone that long, but they liked "hanging out together" through Hubbub.
- The third person told this story. "I'd been using [Hubbub] reasonably often, and I went up to Toronto. The guy I was working with gave me an office, and I'm sitting in there and, you know, there's no music, there's no radio, there's no nothing. And then I put [Hubbub] on [on my Palm], you know it was like I missed everyone. And so I put this on and every once in a while there's some noise and you see what people are doing, you know they're working and feel connected."

The biggest problems with the wireless Palm were that the connection was often unstable and writing substantive messages took some patience. We made several design adjustments to help users know when they were connected and, especially important, when their messages had gotten through. We also modified the protocol to make messages more reliable. These changes made it much smoother to communicate on a wireless Palm. Still, it took some patience to use Graffiti to write messages. We found that Palm users were more willing to send messages with typos in them, and as we noted, they sent more SIMs.

CONCLUSION

Our interest in building Hubbub was to determine how effectively we could support background awareness and opportunistic exchanges among mobile groups through an IM on PCs and wireless Palms. Based on our usage study, we believe we achieved some success with this population. Our

log data indicate that people continued to listen to over half their hubs' activity throughout the study, and the interviews indicate that most people developed a sense of connection with their remote colleagues, and some came to rely on it. We saw an impressive number of opportunistic exchanges, triggered by one party becoming active on Hubbub, which would not have occurred through other IMs. Even though use of sounds was high, people reported finding them only moderately valuable, which we believe reflects people's tendency to undervalue awareness and informal interactions as "useful" for work.

Based on a preliminary analysis of message content, there seems to be two types of Hubbub users, both in their style of text interaction and in their attitude toward the sounds. Heavy adopters had many Hubbub conversations per day with multiple people and used it to *collaborate*, i.e. to accomplish work together and to keep up with each other's news. They frequently had what seemed like one ongoing but intermittent interaction with the same person throughout the day, many installments lacking explicit openings or closings, similar to face-to-face behavior [12]. These users became tuned to the sounds, and learned more Sound IDs and sound messages.

Other Hubbub users had infrequent conversations, perhaps 3 or 4 per week, mostly to *coordinate* meetings in another media or get a single piece of information. They learned only a few Sound IDs plus perhaps just the "Hi" and "Bye" sound. Some of these users liked the "feeling of being connected" even if they didn't interpret all the sounds. A few found the sounds unnecessary or annoying and used Hubbub as a standard IM.

We believe that a sound interface is a powerful mechanism for providing background awareness, which many people seem to appreciate, but only those who are tuned to auditory input are likely to learn the meaning of more than a few sounds. Some people commented that they "don't have the knack for it."

One of the more valuable attributes of Hubbub was that it allowed people to run multiple clients at the same time, displaying to others their most recently active location and automatically sending messages there. Also, those who ran Hubbub on the Palm used it in situations where no other communication medium would have been appropriate.

Finally, we were impressed at how effectively this text-based system afforded rich, complex, deeply engaging interactions. Although someday there will be a universal infrastructure to support lightweight video- and audio-based interactions, text will continue to be a valuable medium for collaboration.

ACKNOWLEDGMENTS

Thanks to Candace Kamm, Steve Whittaker, Diane Schiano, and Jonathan Helfman for helpful comments on this paper; Dave Kormann for his responsive support of the Hubbub server; and the Hubbub users who ran pre-alpha software and gave us valuable feedback throughout the study. Thanks especially to Candace Kamm and Ron Brachman for their enthusiastic support of the project.

REFERENCES

1. Bellotti, V. & Bly, S. (1996) Walking away from the desktop computer: distributed collaboration and mobility in a product design

team, *Proc. Computer-Supported Cooperative Work (CSCW)*, Cambridge, MA, 209-218.

2. Bly, S.A., Harrison, S.R. & Irwin, S. (1993) Media spaces: Bringing people together in a video, audio, and computing environment, *Communications of the ACM*, 36:1, 28-47.

3. Brewster, S., Wright, P., & Edwards, A. (1993) An evaluation of earcons for use in auditory human-computer interfaces. *Proc. InterCHI'93*, Amsterdam, Netherlands, 222-227.

4. Churchill, E.F. & Bly, S. (1999), Virtual environments at work: Ongoing use of MUDs in the workplace, *Proc. WACC*, 99-108.

5. Erickson, T., Smith, D., Kellogg, W., Laff, M. Richards, J. & Bradner, E. (1999) Socially translucent systems: social proxies, persistent conversation, and the design of Babble, *Proc. CHI*, Pittsburgh, PA, 72-79.

6. Fish, R., Kraut, R., Root, R. & Rice, R. (1993) Video as a technology for informal communication, *Communications of the ACM*, 36:1, 48-61.

7. Gaver, W. (1997) Auditory interfaces, in Helander, M., Landauer, T. & Prabhu, P. (Eds.), *Handbook of Human-Computer Interaction*, Amsterdam: Elsevier Science.

8. Grinter, R. & Eldridge, M. (2001) y do tngrs luv 2 txt msg?, *Proc. ECSCW*, Bonn, Germany.

9. Hindus, D., Ackerman, M., Mainwaring, S., & Starr, B. (1996) Thunderwire: a field study of an audio-only media space, *Proc. of CSCW*, Boston, MA, 238 - 247.

10. Isaacs, E. & Walendowski, A. (2001) *Designing From Both Sides of the Screen: How Designers and Engineers Can Collaborate to Build Cooperative Software*, New Riders.

11. Kraut R.E., Egidio C. & Galegher J. (1990) Patterns of contact and communication in scientific research collaboration. In J. Galegher & R. Kraut (Eds.) *Intellectual teamwork*, Hillsdale: Erlbaum, 149-171.

12. 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.), *People's Reactions to Technology*, Newbury Park: Sage, 145-199.

13. Kraut, R.E. & Streeter, L.A. (1995). Coordination in software development, *Communications of the ACM*, 38:3, 69-81.

14. Luff, P. & Heath, C (1998) Mobility in collaboration. *Proc. CSCW*, Seattle, WA, 305-314.

15. Mynatt, E., Back, M. & Want, R. (1998) Designing Audio Aura, *Proc. CHI*, Los Angeles, CA, 566-573.

16. Nardi, B., Whittaker, S. & Bradner, E. (2000) Interaction and interaction: instant messaging in action. *Proc. CSCW*. Philadelphia, PA, 79-88.

17. Rhineland, T. (2000) *Intense users will drive increased IM capabilities*, Forrester Technographics Brief.

18. Tang, J., Isaacs, E., & Rua, M. (1994) Supporting distributed groups with a Montage of lightweight interactions. *Proc. CSCW*, Chapel Hill, NC, 23-34.

19. Tang, J., Yankelovich, N., Begole, J., Van Kleek, M., Li, F. & Bhalodia, J. (2001) ConNexus to AwareNex: Extending awareness to mobile users. *Proc. CHI*, Seattle, WA, 121-128.

20. Whittaker, S., Frohlich, D. & Daly-Jones, W. (1994) Informal Workplace Communication: What is it Like and How Might We Support It? *Proc. CHI*, Boston, MA, 131-137.