

Presence versus Availability: The Design and Evaluation of a Context-Aware Communication Client

James Fogarty
*Human Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213
jfogarty@cs.cmu.edu*

Jennifer Lai, Jim Christensen
*IBM T.J. Watson Research Center
Hawthorne, NY 10532
{jlai, ibmjim}@us.ibm.com*

Although electronic communication plays an important role in the modern workplace, the interruptions created by poorly-timed attempts to communicate are disruptive. Prior work suggests that sharing an indication that a person is currently busy might help to prevent such interruptions, because people could wait for a person to become available before attempting to initiate communication. We present a context-aware communication client that uses the built-in microphones of laptop computers to sense nearby speech. Combining this speech detection sensor data with location, computer, and calendar information, our system models availability for communication, a concept that is distinct from the notion of presence found in widely-used systems. In a four week study of the system with 26 people, we examined the use of this additional context. To our knowledge, this is the first field study to quantitatively examine how people use automatically sensed context and availability information to make decisions about when and how to communicate with colleagues. Participants appear to have used the provided context to as an indication of presence, rather than considering availability. Our results raise the interesting question of whether sharing an indication that a person is currently unavailable will actually reduce inappropriate interruptions.

INTRODUCTION AND MOTIVATION

Electronic communication has become a critical part of many workplaces, with workers relying on email, instant messaging, and phone calls to communicate with colleagues across the hall or across the world. Whether chosen out of necessity or for convenience, these tools provide important complements to face-to-face meetings. However, current widely-used systems generally provide inadequate support for estimating the availability of colleagues.

Consider approaching a colleague's office to initiate a face-to-face meeting. People are usually able to quickly assess the presence and availability of a colleague. A colleague who is on the phone is present, but probably not available. An empty office, with the

door left open and a laptop computer still open on the desk, probably indicates that the colleague has just stepped out. While not present, it seems that the colleague is more likely to be available soon than if the door were closed and the colleague gone for the day. A colleague who is quietly browsing through email would seem to be relatively available. People regularly use these types of real-world cues to make such assessments of the presence and availability of another person.

In contrast, widely used electronic communication tools generally support, at most, only an indication of presence. Phones generally provide no information about the context of colleagues, and so people may unwittingly call colleagues at awkward or inappropriate times, such as when a person is in a meeting (Schmidt, 2000). Instant messaging clients often use recent computer activity to estimate presence, but many instant messaging sessions start with queries to determine if a colleague is actually present and to ask whether the colleague is available (Voids, 2002). This behavior seems to indicate that information available in an instant messaging client is insufficient for making an informed decision about the availability of a colleague.

In a study of interruptions and corporate research managers (Hudson, 2002), Hudson *et al* found that some people consider electronic interruptions so distracting that they physically move, either to a part of the office away from the computer or to somewhere outside the office. Hudson *et al* also suggest that the work required to configure electronic communication tools to filter interruptions might be more than can be reasonably expected of an already busy person. One suggestion of Hudson *et al* is the consideration of socially translucent systems, as discussed in (Erickson, 2000). They suggest that a system automatically collecting and sharing information about the availability of colleagues could allow a person to use that information to decide whether it is appropriate to initiate communication.

This paper presents MyVine, a client that integrates phone, instant messaging, and email communication channels with automatically sensed information about the context and availability of colleagues. MyVine analyzes audio from the built-in microphone of a laptop computer to detect speech. Speech has been shown to be a very good indicator that an office worker is busy and should probably not be interrupted (Fogarty, TOCHI 2004; Hudson, 2003; Fogarty, CHI 2004). MyVine uses speech detection, location information, computer activity, and calendar information to model a person's availability. We chose these sensors because of their relation to availability and because we feel that they can reasonably be considered for large-scale deployment.

Prior work on presence and availability has primarily examined cases where a person does not appear to be present, but might be available. The Priorities system examines incoming messages and, if a colleague is not present, forwards urgent messages to a mobile device (Horvitz, 1999). The Coordinate system uses calendar information and previous computer activity to model the likelihood of a person being available on a particular computing device, such as a desktop computer or a PDA (Horvitz, 2002). Begole *et al* have examined temporal patterns in a person's presence to estimate when that person might be likely to return (Begole, 2002; Begole, 2003). They point out that a person who is usually present at a given time of day might be available at that time even

if the person's computer has been inactive for several minutes, as in the case where a person is reading near the computer. MyVine focuses on supporting the inverse, situations where a person might be present but not available.

MyVine's use of a laptop computer's built-in microphone has some important properties. First, laptop computer sales have recently surpassed desktop computer sales ("Laptop sales beat desktops for first time," 2003), making computers with these microphones increasingly ubiquitous. Second, there is no additional cost associated with this sensing capability, making it possible to consider large-scale deployment. In contrast, even individually inexpensive sensors quickly become expensive if they are to be deployed throughout a large organization. Finally, it is a part of many corporate cultures to bring laptop computers to meetings, as they are useful even when people are away from their offices. This allows MyVine to detect speech anytime a person's laptop is nearby and on, not just inside a person's office.

We studied MyVine's use for four weeks by a total of 26 people. To our knowledge, this is the first field study to quantitatively examine how people use automatically sensed context and availability information to decide when and how to communicate with colleagues. Participants appear to have used the provided context as an indication of presence, rather than considering availability. Our results raise interesting questions over whether sharing an indication that a person is currently not available will reduce inappropriate interruptions.

In the next section, we briefly review related work on media spaces. This is followed by a discussion of the functionality of MyVine. Next is a presentation of some rationale for the design of MyVine. We then discuss our deployment of MyVine and present the results of this deployment. Our final section is a discussion and a consideration of approaches to future work.

RELATED WORK ON MEDIA SPACES

Extensive work on media spaces has examined the possibility of audio and/or video links between physical spaces, often with the goal of encouraging casual interactions among distributed colleagues or improving awareness of the activities of a group (Gaver, 1992; Bly, 1993; Fish, 1993; Adler, 1994; Tang, 1994). One important issue with media spaces is the privacy implications of somebody having access to live audio and/or video of a person. Prior work has sought to address privacy concerns by filtering the audio and/or video to remove sensitive information (Smith, 1995; Hudson, 1996). There are also problems with the distraction created by a media space. Prior work in media spaces has argued a fundamental tradeoff between information and distraction, suggesting that important information could be distilled and presented (Smith, 1995; Hudson, 1996).

MyVine is not a media space, and its primary purpose is not the encouragement of casual interaction among distributed workers. Rather than creating a shared space and filtering sensitive information, MyVine is a context-aware communication system that analyzes audio to detect speech, but never allows a person to listen to the audio. This distinction is important, because some people who might never allow colleagues to listen to audio could be more willing to accept a sensor-based approach. MyVine also serves a different

purpose than a media space, sharing information about context and availability to help prevent inappropriate interruptions.

(Figure 1 goes here or near)

MYVINE FUNCTIONALITY

Figure 1 shows MyVine being used within a group of five colleagues. Each person's availability is presented in three levels of detail, starting with the person's image. When highly available, a person's image is shown in full color. Otherwise, a grayscale filter is applied to the person's image such that highly unavailable people are shown completely in gray and people of moderate availability are shown in a combination of color and gray. In Figure 1, the third person and the fourth person are unavailable, while the other three people are currently available. A second level of detail can be obtained by placing the cursor near a person's image. In Figure 1, the third person's context is being investigated. A context tray has appeared over the person's image, and four icons are visible in the tray. From the top to the bottom, these icons are an open door, indicating that the person is probably in their office, a speech bubble, indicating that the system has detected speech near the person's computer, a picture of a computer, indicating that the person is active on the computer, and a picture of a small flip calendar, indicating that the person's calendar shows them as busy. The cursor is also close enough to the second person that her details are partially visible. A third level of detail, that MyVine has detected speech for the previous 13 minutes, has been obtained by hovering the mouse over the appropriate icon. Even though the middle person is present, the details of his context indicate that he is probably not currently available for communication, so his image is shown in gray. MyVine has also suggested email as the most appropriate way to reach the middle person, and has indicated this by making the button for email appear larger and centered above the buttons for instant messaging and the phone.

Context Used by MyVine

MyVine uses speech detection, location, computer activity, and calendar entries to model availability. These pieces of context were chosen because of their relation to availability and because they can be obtained without installing any additional hardware or sensing infrastructure, making it reasonable to consider large-scale deployment. Context is shared by a background process on a person's computer, so it is available even if a person is not using their client to view the context of colleagues. The server infrastructure used by MyVine also supports global or individual permissions on the context that is shared.

Previous work examined a variety of potential sensors for office environments and found that a speech detector is an excellent indicator that an office worker is currently busy and should not be interrupted (Fogarty, TOCHI 2004; Hudson, 2003; Fogarty, CHI 2004). Our speech detector is designed to detect extended conversations, as opposed to short utterances. We currently analyze audio using a combination of an adaptive energy function implemented by the Sphinx silence detector ("CMU Sphinx," 2003), a high zero crossing rate ratio (HZCRR) feature that has been previously shown to be a good filter of non-speech noise (Lu, 2002), and a filter for noise created by typing on a laptop computer's keyboard. Sphinx's silence detector adapts to the energy level of audio

frames read from the microphone. In our experience, this adapts well to the laptop computer's fan noise, the sound of air conditioning, or other relatively constant background noises. HZCRR detects the large differences in the zero crossing rates of audio frames characteristic of speech. In our experience, this works well for filtering out many noises that are loud enough to activate the Sphinx silence detector, but are not speech. Finally, the extreme proximity of a laptop computer's keyboard to its microphone means that typing is very loud, leading us to filter the output of the silence detector to suppress activations that occur when a person is typing. MyVine analyzes the filtered output of the silence detector and the HZCRR feature in a sliding 15-second window, applying a threshold to determine if it detected conversation during the window. MyVine reports a conversation when the threshold is met for 30 of the past 60 seconds. This allows it to ignore brief utterances and account for natural pauses in conversation. When speech is detected, MyVine includes a speech bubble in a person's context tray with a tooltip showing for how long speech has been detected.

MyVine uses wired and wireless network connectivity to estimate location. One of four location icons is always in each person's context tray. If a person is currently on a virtual private network or dial-up connection, a "remotely connected" icon is used. A "probably in office" icon is used if a person is connected to the local network from the access point that the person uses most often. If a person is on the local network but at a different access point, a "probably in the building" icon is used. If not connected, or if a person's computer has been idle for a long period of time, a "probably not available" icon is shown.

A computer icon is used to indicate that a person has been active on their computer in the last three minutes. If a person's computer has been idle for an extended period of time, an idle icon is shown with a tooltip indicating how long the computer has been idle. Similarly, a calendar icon is shown when a person's electronic calendar contains an event scheduled for the current time.

Communication Channels

MyVine provides access to colleagues via phone, instant messaging, and email. It also suggests a communication channel for each colleague, using a set of rules based on the colleague's context and availability. Email support is provided by spawning a mail client, and instant messaging support is provided by programmatically starting a chat session with the selected person. Phone support is provided by a server that stores phone numbers for people and locations. When a caller initiates a phone call to a callee, the server uses the location of the caller to determine an appropriate phone number, calls the caller, and waits for the caller to answer the call. The server then uses the callee's location to call the callee at an appropriate phone number and finally connects the two calls. Phone numbers are initially obtained from a database containing the office phone numbers for employees, but the server stores numbers that a person enters for a particular location. Thus, a person could provide a mobile phone number the first time the person made a call from a remote location, and the system would use the number for future calls made when that person was remotely connected. The system can similarly be told what phone number to use when a person is in a lab or an office.

(Figure 2 goes here or near)

Figure 2 shows how MyVine suggests a communication channel. Because the person on the left is currently in his office and active on his computer, MyVine has suggested the phone by making its button larger and centering it above the buttons for email and instant messaging. But MyVine leaves the choice to the person, and a person could still choose to send an email or an instant message. Because the middle person has been idle for a long time and is probably unavailable, MyVine suggests email. Instant messaging is unavailable because she is not running a client, but a person could still try to call her if the person knew her computer was broken or if they just wanted to leave a voicemail. The person on the right is currently in his office and active on his computer, but MyVine has detected that he is engaged in a conversation and suggests instant messaging. If his conversation continues, his availability will be further reduced and MyVine will suggest email.

DESIGN RATIONALE

This section presents two aspects of the design rationale for MyVine. First is our decision to design for small groups of work colleagues. Second is the decision to use multiple levels of detail, as opposed to a design in which the details of each person's context were always visible.

Intended Use of MyVine

MyVine is intended for use by a small group of colleagues in a workplace, in conjunction with existing communication tools. This is important for the obvious reason that the current interface could become unwieldy when used with more than about ten people. While a person may have fifty or more people in the buddy list of their instant messaging client, not all of those people should be added to MyVine. Instead, a person should use MyVine for the colleagues that the person communicates with most often, occasionally adding or removing people who become more or less important to everyday communication. People should continue to use existing mechanisms, such as an instant messaging client's buddy list, for communication with people outside their core group.

The intended use of MyVine is also important for another reason that is more subtle, but critical to successful use. If a person agrees to make information available to a socially translucent system like MyVine, they will expect some benefit. In the case of MyVine, that person expects that colleagues will use the person's context to make more appropriate decisions about the person's availability for communication. The person expects to be subjected to fewer socially inappropriate interruptions, thus justifying the decision to share the context. However, this expectation will probably not be met without a motivation for other people to respect a person's availability. Consider that spammers typically ignore requests to be removed from their lists because there is no viable mechanism for technically, legally, or socially applying a cost for not complying with the request. Once a spammer has your information, there is no reason to expect that the spammer will use the information appropriately. In contrast, workers generally expect their colleagues to treat them with a certain degree of politeness and respect. If a person shares context indicating that they are not available, and a colleague interrupts the person,

the interrupted person knows that the colleague knew the person was not available. The colleague can then be held accountable if the interruption is inappropriate. This notion of accountability is critical to socially translucent systems (Erickson, 2000). The results section of this paper presents quantitative data from a pilot study examining how groups of colleagues used the presence and availability context shared by MyVine and discusses those results in the context of accountability and socially translucent systems.

Multiple Levels of Detail in Availability Presentation

Because traditional messaging clients generally support only presence, they convey just one or two bits of information about each person on a buddy list. They might use the font in a person's name, as when a bold font is used for people who are currently online. Alternatively, they might show an icon next to the person's name, which could be green when the person is online and active, or red otherwise.

The context and availability information presented by MyVine seem to require a different approach. While any one of the pieces of context might be integrated into an existing client by adding an icon or a bit of descriptive text, the approach would seem to quickly break down as additional forms of context become available. The result could be many confusing and distracting icons, or a large amount of text associated with each person. In either case, it could be difficult to make a quick and accurate judgment of whether or not a person is currently available.

MyVine manages the complexity of presenting context and availability by using multiple levels of detail, an approach commonly known as progressive disclosure. At the first level of detail, a person's image fades gradually, conveying the person's availability while minimizing distraction. The next two levels of detail allow a person to investigate the details of a colleague's context. This approach is enabled by our model of the availability of a person and our model for selecting an appropriate communication channel. We believe this strategy to be good approach to handling the fundamental tradeoff between providing information to make a decision and providing too much information such that it is distracting (Smith, 1995; Hudson, 1996). Tang *et al* have also examined interfaces to present similar pieces of context for a colleague, but did not use a speech detection sensor and did not include a model of availability (Tang, 2001).

DEPLOYMENT AND STUDY

In order to gain insight into how people use the context and availability indications provided by MyVine, we deployed MyVine for four weeks with a total of 26 people in four groups not directly related to this work. This section introduces the groups and discusses the deployment. The next section will then discuss some results of the deployment and implications for sensing and sharing context and availability information.

In our study, groups A, B, and C are actively collaborating groups of research colleagues. Group A has seven members, Group B has four members, and Group C has five members. Although there is some communication across group boundaries, each group communicates heavily within the group. Group D is very different from the first three groups. The ten people in this group host customers visiting the research lab,

demonstrate technology to customers, and coordinate interactions between researchers and customers. The presence of customers would seem to make their communication needs more immediate.

MyVine was installed on each person's computer and each person was given a short introduction. Participants were encouraged to use MyVine for the colleagues they communicated with most often. Participants were also advised that they should feel free to add or remove people throughout the study. We asked participants to use the MyVine client to initiate communication with people they had added to the client, rather than manually making a phone call or initiating a chat session with the chat client interface. We did not, however, take away or otherwise disable their existing communication mechanisms, for two reasons. First, we wanted to minimize the intrusion created by participating in our deployment, as the participants were all busy with their normal work. Second, their existing tools were still necessary for communicating with people not participating in our deployment. We also told participants that they should not feel they needed to add a person to the client in order to communicate with them once, but rather should focus on using it with the people they communicate with most often.

In order to study this deployment, we automatically logged the state of each person's client. Anytime a person initiated communication with a colleague, we logged the context and availability information being shown for that colleague. We also logged the information that was shown whenever a person hovered the mouse over a colleague to examine the details of that colleague's availability. Finally, every 30 seconds we logged the information for each colleague in a person's client. As we will show in this section, these logs allow us to quantitatively examine how people used the context and availability information provided by MyVine.

RESULTS

This section presents results from logs collected during our pilot study and from a questionnaire given at the end of the study. To our knowledge, this is the first field study to quantitatively examine how people use automatically sensed context and availability information to make decisions about when and how to communicate with colleagues.

During our deployment, we logged approximately 900 person-hours of usage, or an average of approximately 35 hours per participant. This figure represents the total time that participants had the client running and were active on their computers. This is lower than the 160 hours that might have been expected from four weeks with each participant, due to a variety of reasons. Some participants were on vacation or otherwise away from the office during the deployment. Because they were busy with their normal jobs, participants were often away from their computers, such as when they were on the demonstration floor with a customer. We did not configure participant's computers to automatically start MyVine, and some participants reported that they regularly forgot to start the client. Finally, occasional problems with servers meant that MyVine was sometimes unavailable.

During the approximately 900 person-hours of use, participant had an average of approximately 6 colleagues visible in their MyVine clients. There were 2560 instances of

a participant hovering the mouse over an image of a colleague to reveal the details of that colleague's context, or an average of approximately one hover for every 20 minutes of use. Communication was initiated via the MyVine client for 180 of these hovers, or for approximately one of every 14 investigations of a colleague's context. These 180 communications consisted of 42 phone calls, 117 instant message sessions, and 21 email messages.

The level of use found during our deployment leads us to believe that participants did not initiate all of their communication with MyVine. This is confirmed by responses on our post-deployment questionnaire indicating that most participants used the client for less than 20 percent of their communication. Instead, we believe that participants used MyVine only when they were initiating a new communication and the information available in MyVine might affect their decision to initiate communication. When participants had already decided how to communicate with a colleague, they used their existing communication tools. For example, participants reading email used the reply functionality of their email client, rather than consulting MyVine or using the communication channel buttons provided by MyVine. Similarly, participants with an especially urgent need for a colleague probably called the colleague using their normal tools. In future work, it would be helpful to develop a way to log interactions with existing communication tools. We also plan to examine including availability and context information in the interfaces of existing tools.

In discussing the results of this deployment, we limit ourselves to analyses we feel are appropriate for the data we collected. For example, our analyses combine the data from all four groups because we did not intend to focus on differences between groups and do not really have enough data, or the right data, to carefully examine possible differences. Similarly, we believe it may be the case that some of our participants, perhaps the managers of the groups that participated, are treated differently than other participants. But these types of questions are not the focus of this deployment, and we would have designed our deployment differently if we had intended to pursue such issues. For example, we might have instrumented MyVine such that it automatically administered a short questionnaire to participants when they did something we were interested in, such as examining a colleague's detailed context and then deciding whether to initiate communication with that colleague. Instead, we focus on the relationship between the context and availability information provided by MyVine and the actions taken by participants, using data collected in a minimally-intrusive manner from our four groups of participants during their normal work.

In the remainder of this section, we discuss two primary results. We first consider how participants used our grayed images of colleagues as an abstract presentation of availability. This approach was chosen because it hides icons that might be distracting if always visible, but it was unclear whether participants would trust our models of a colleague's availability. If participants felt they needed to constantly investigate the details of a colleague's context, it might indicate that our model of availability is inappropriate. We then examine how the transition from hovering over a colleague's image to initiating communication was affected by the detailed context. Specifically, we

explore whether participants seem to have used our speech detector as an indication that a colleague was not available, or only as an indication that a colleague was present.

Availability Presentation and Hovering for Details

Our model of availability has four levels, with a 1 representing highly unavailable and a 4 representing highly available. This subsection examines these four levels of availability and how their visualization as different levels of gray in the image of a colleague affected participant decisions to investigate the detailed context of a colleague.

(Figure 3 goes here or near)

The left side of Figure 3 shows the distribution of availability seen by participants during the MyVine deployment. Colleagues were much more likely to be shown as highly available or highly unavailable than they were to be shown in state of moderate availability. This is a result of how our model of availability is implemented. If speech is detected for an extended period of time or a person has been away from their computer for an extended period of time, the person will be shown as highly unavailable. But when speech is first detected, it might be premature to indicate that the person is highly unavailable. In this case, MyVine shows the colleague in one of the states of moderate availability. The colleague will move to being highly unavailable if the speech continues, or back to being highly available if the speech ceases. The distribution of availability obtained with our participants indicates that our participants spent relatively little time in these ambiguous transition states.

The right side of Figure 3 shows the relationship between a colleague's availability and how often our participants hovered the mouse over that colleague to obtain detailed context. These numbers are based on the total time that colleagues were shown in these availability levels and the number of investigations of the detailed context of colleagues in these availability levels, but do not account for the fact that participants had an average of 6 colleagues visible at a time (which is why they are higher than the hover frequency of once per 20 minutes that we reported in the very beginning of this section). Participants were significantly more likely to investigate the detailed context of a colleague in a moderate state of availability (once per 67 minutes in an availability of 2 and once per 70 minutes in an availability of 3) than they were to investigate the detailed context of a colleague who was shown as highly available or highly unavailable (once per 138 minutes in an availability of 1 and once per 104 minutes in an availability of 4) ($\chi^2(1, 539795) = 161, p < .0001$). This seems to indicate that participants generally trusted an indication by MyVine that a colleague was highly available or highly unavailable. When MyVine's indication of availability was more ambiguous, participants were more likely to investigate the details of a colleague's context. It is also interesting to note that participants were significantly less likely to investigate the context of a colleague shown as highly unavailable (1) than they were to investigate the context of a colleague shown as highly available (4) ($\chi^2(1, 439862) = 36.6, p < .0001$). This is consistent with participants being accustomed to clients that convey only presence, because widely-used commercial clients generally provide very little information about a colleague who is not present.

The findings presented in this subsection indicate decisions by participants to investigate the details of a colleague's context were not random. Participants appear to have believed MyVine when our model indicated that a person was unavailable, and they appear to have been willing to examine the details of a colleague's context when our model provided an ambiguous indication of availability.

Context Details and Initiating Communication

While the previous subsection considered when participants were likely to investigate the details of a colleague's context, this subsection investigates whether they were likely to initiate communication after viewing the details of a colleague's context. We will focus on participant use of instant messaging, for three reasons. First, participant use of the telephone functionality was much less affected by the available context than participant use of instant messaging, a point we will briefly illustrate in this subsection. Second, we believe that participant use of context and availability information with instant messaging is most relevant to the research community, as most communication clients do not support the additional channels included in MyVine. Finally, we logged relatively few instances of participants using the email functionality of MyVine, and it would seem to be inappropriate to pursue an analysis based on that data.

(Figure 4 goes hear or near)

Figure 4 shows participant use of MyVine's phone and instant message functionality by the availability of the person being communicated with. Most instant message communication took place with a colleague who was highly available. This might indicate people generally waited for a colleague to be available before sending an instant message, or that they messaged whoever seemed to be available when a need arose. In contrast, phone calls appear to have been made with less concern for a colleague's availability. This seems to indicate that participants used phone calls when there was a sense of urgency or some other force motivating them to disregard the availability of a colleague. Given this difference between how participants seem to have used the phone and instant messaging, we will focus on how the context provided by MyVine affected the decision to send an instant message to a colleague.

In examining the detailed context that participants used to decide whether to send an instant message, it is not surprising that participants were much less likely so send an instant message to a colleague who was not actively using their computer. Instant messaging clients commonly describe such a person as "idle" or "away," and people are accustomed to this indication that a colleague is not present. In the remainder of this section, we will use "idle" to describe a participant who was not active on their computer and "non-idle" to describe a participant who was actively using their computer.

(Figure 5 goes here or near)

The top portion of Figure 5 shows that hovers over the image of a colleague were much more likely to result in an instant message if the colleague was not idle. To clarify the information presented in the top part of the figure, there were 1325 instances of a participant hovering the mouse over non-idle colleagues, and 86 instant messages were sent to non-idle colleagues. In contrast, 1240 hovers over idle colleagues resulted in only

31 instant messages. Together, this shows that participants were 2.6 times more likely to initiate an instant message session with a colleague who was not idle, a significant difference ($\chi^2(1, 2682) = 21.4, p < .0001$). While widely-used communication systems often convey a notion of presence based on whether or not a person's computer is idle, we are also interested in how participants used the additional context provided by MyVine. We first examine how additional context affected communication with colleagues who were not idle, and then consider participant use of additional information about idle colleagues.

The first line of the non-idle portion of Figure 5 shows the influence of the calendar icon. Calendar information is generally not considered very reliable, because not all calendar items imply unavailability and because many people have calendar items that they do not actually attend. This has prompted previous work to investigate automatically learning what calendar items indicate that a person is unavailable (Mynatt, 2001; Horvitz, 2002; Horvitz, 2003). In our field study, displaying a calendar item icon for a non-idle colleague whose calendar indicated they were busy reduced instant messages by 40%. It did so without exposing the length or content of calendar items, and was therefore rather sensitive to the legitimate privacy concerns associated with exposing calendar information. This result is not quite significant ($\chi^2(1, 1411) = 2.01, p \approx .16$), but it does seem to indicate a trend towards participants using the calendar as an indication that a non-idle colleague was not available.

The second line of the non-idle portion of Figure 5 shows the difference between the "Probably in Office" icon used when a person was connected to the access point that the person used most often and the "Probably in Building" icon used when a person was connected to the local network, but on a different access point. We had expected that people would consider their colleagues more available when they were in their office, and so it was surprising to see that the difference between these two conditions had no effect on the initiation of instant messaging sessions to non-idle colleagues ($\chi^2(1, 1278) = 0.01, p > .93$). On the other hand, the third line shows that participants were much less likely to send an instant message to a non-idle colleague who was remotely connected ($\chi^2(1, 1411) = 7.32, p < .01$). This seems to indicate that participants considered a colleague's presence in the building to be an important indication of the colleague's availability.

The last line of the non-idle portion of Figure 5 considers the impact of our speech detector. We expected it to be used as an indicator a colleague is not available, implying participants would be less likely to send an instant message. However, Figure 5 shows that participants were more likely to send an instant message to a non-idle colleague shown as speaking. Though not significant ($\chi^2(1, 1411) = 1.28, p \approx .26$), this indicates that our participants did not use our speech detector as an indication that a colleague was not available, and that they might have instead used it as an indication of a colleague's presence. We will revisit this result in our discussion, as it raises the question of whether or not sharing this sort of information reduces inappropriate interruptions.

Moving to consider what effect additional context had on sending instant messages to colleagues shown as idle, the first line of the idle portion of Figure 5 shows participants

were more likely to send an instant message to an idle colleague shown as in the building ($\chi^2(1, 1271) = 18.3, p < .0001$). This is consistent with participants being more likely to send an instant message to a non-idle colleague who was shown to be in their office or in the building. This indicates that participants considered physical presence important when sending an instant message, whether or not that colleague was idle.

Finally, the last line in Figure 5 shows that participants were significantly more likely to message an idle colleague whose calendar indicated that the colleague was busy ($\chi^2(1, 1271) = 17.3, p < .0001$). This is an odd result, but might be explained by the fact that actively collaborating colleagues, almost by definition, are likely to meet with each other. These messages might represent one participant trying to find out why another is not yet present at a scheduled meeting.

DISCUSSION

Our participants generally seem to have used the context shared by MyVine to determine whether or not a colleague was present. The difference between being in an office or in the building did not seem to be important, as both indicated that a colleague was present. While speech is a strong indicator a person is not available and should not be interrupted (Fogarty, TOCHI 2004; Hudson, 2003; Fogarty, CHI 2004), showing that a colleague was speaking does not seem to have discouraged instant messages.

The willingness of participants to send an instant message to a colleague who was present, but shown as not available, might be partially explained by a *perception* that instant messages are a non-intrusive way of communicating (Nardi, 2000). We find this explanation unsatisfying because it does not provide any insight into how we can help to prevent inappropriate interruptions. Although interruptions can help workers solve problems, there is a clear need to balance this potential gain with a need for uninterrupted time (Perlow, 1999). Given that workers have developed strategies to avoid interruptions, such as physically moving away from their computers or offices, there is clearly a problem with current systems (Hudson, 2002). Relying on manual do-not-disturb flags or other states that are manually set is inadequate, because people forget to set them when they are busy and forget to clear them when they become available (Milewski, 2000; Lai, 2003).

Part of the problem seems to be a lack of accountability. One of the critical ideas behind socially translucent systems is accountability for inappropriate behavior (Erickson, 2000). However, holding a work colleague accountable for sending an instant message at an inappropriate time, even by simply pointing it out to the colleague, might be considered more impolite than the original infraction. Because any single interruption does not seem to warrant serious concern, people might choose to be polite and not address an issue of inappropriate interruptions, even though the interruptions as a whole are a problem. Other factors related to the use of communication in a workplace, such as the organizational hierarchy or a sense of urgency due to an impending deadline, may mean that people focus more on their own needs, rather than the needs of a colleague. Perlow has shown that this behavior can have a very negative impact on productivity. Workers struggling to make a deadline might interrupt colleagues with an urgent request, but these interruptions delay the work of the person who was interrupted, and they might in turn

find it necessary to interrupt another person in order to make their deadline, creating a vicious cycle (Perlow, 1999).

One approach to future work suggested by our results is the integration of availability information directly into existing communication clients. Many of our participants used their existing clients, rather than our set of three buttons for initiating communication. It seems that integrating availability information directly into these interfaces is the best way to ensure that it is visible during the consideration and initiation of communication.

A second possible approach to future work stems from our finding that participants seem to have trusted the fully gray image as an indication that a colleague was highly unavailable, and were much less likely to investigate the detailed context of a colleague shown as a fully gray image. If a model were to be more aggressive in labeling a person as highly unavailable, fewer inappropriate interruptions might result. Designers might also consider only exposing the detailed context of a person who is available, hiding the details of colleagues who are highly unavailable. As participants primarily used detailed context as an indication of presence, such a design might have the potential for reducing inappropriate interruptions because people might defer to a colleague shown as highly unavailable, rather than seeing that the colleague is present and being tempted to try to communicate.

Finally, our results raise questions about the possibility of successfully applying the ideas of socially translucent systems to the problem of managing interruptions created by instant messaging and other synchronous communication tools. MyVine shares context and availability with the expectation that people will respect an indication that a colleague is unavailable. But the results of our deployment do not support this expectation, suggesting instead that participants primarily used the additional information as an indication of presence. If people do not benefit from sharing additional context information, they will probably not agree to share such information. Another possibility might be to instead consider a filtering approach, wherein a person's instant messaging client delays the arrival of new messages when a person appears to be otherwise busy. Such an approach might be successful if the delays were short, such as waiting for a person to finish a sentence in a document they are typing. On the other hand, longer delays might confuse the difference between email and instant messaging, possibly reducing the effectiveness of instant messaging in the workplace. While we remain interested in pursuing socially translucent systems as an approach to the problems of inappropriate interruptions in electronic communication, we argue that future work needs to carefully examine the use of additional context and availability information, in order to ensure that sharing such information provides people with the intended benefits.

CONCLUSION

We have presented a context-aware communication system that senses nearby speech, a novel piece of context previously shown to indicate that people are not available. Combining speech detection with location information, computer activity, and calendar information, our system models a person's availability for communication. In a four week deployment of the system with a total of 26 people, we examined how participants used this additional context. Finding that participants appear to have focused on

presence, rather than considering availability, we have raised the question of whether sharing an indication of availability will reduce inappropriate interruptions. Based on the results of our deployment, we have suggested possibilities for continued exploration of this problem.

ACKNOWLEDGMENTS

We would like to thank all of the participants in our deployments. We also thank all of the members of the Grapevine group, who have worked on the many pieces of infrastructure in MyVine. Finally, we would like to thank the many people who shared their comments and insights into MyVine and the problems it explores.

REFERENCES

- Adler, A. and A. Henderson (1994). "A Room of Our Own: Experiences from a Direct Office Share". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 1994)*.
- Begole, J. B., J. C. Tang, R. B. Smith and N. Yankelovich (2002). "Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups". *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2002)*.
- Begole, J. B., J. C. Tang and R. Hill (2003). "Rhythm Modeling, Visualizations, and Applications". *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2003)*.
- Bly, S. A., S. R. Harrison and S. Irwin (1993). "Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment." *Communications of the ACM* **36**(1): 28-46.
- CMU Sphinx: Open Source Speech Recognition. <http://www.speech.cs.cmu.edu/sphinx/>.
- Erickson, T. and W. A. Kellogg (2000). "Social Translucence: An Approach to Designing Systems that Support Social Processes." *ACM Transactions on Computer-Human Interaction (TOCHI)* **7**(1): 59-83.
- Fish, R. S., R. E. Kraut, R. W. Root and R. E. Rice (1993). "Video as a Technology for Informal Communication." *Communications of the ACM* **36**(1): 48-61.
- Fogarty, J., S. Hudson, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee and J. Yang (2004). "Predicting Human Interruptibility with Sensors." *To Appear, ACM Transactions on Computer-Human Interaction (TOCHI)*.
- Fogarty, J., S. Hudson and J. Lai (2004). "Examining the Robustness of Sensor-Based Statistical Models of Human Interruptibility". *To Appear, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2004)*.
- Gaver, W., T. Moran, A. MacLean, L. Löfstrand, P. Dourish, K. Carter and W. Buxton (1992). "Realizing a Video Environment: EuroPARC's RAVE System". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 1992)*.
- Horvitz, E., A. Jacobs and D. Hovel (1999). "Attention-sensitive alerting". *Proceeding of the Conference on Uncertainty and Artificial Intelligence (UAI 1999)*.
- Horvitz, E., P. Koch, C. M. Kadie and A. Jacobs (2002). "Coordinate: Probabilistic Forecasting of Presence and Availability". *Proceedings of the National Conference on Uncertainty and Artificial Intelligence (UAI 2002)*.
- Horvitz, E. and J. Apacible (2003). "Learning and Reasoning about Interruption". *Proceedings of the ACM International Conference on Multimodal Interfaces (ICMI 2003)*.
- Hudson, J. M., J. Christensen, W. A. Kellogg and T. Erickson (2002). "'I'd be overwhelmed, but it's just one more thing to do': Availability and Interruption in Research Management". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2002)*.
- Hudson, S., J. Fogarty, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee and J. Yang (2003). "Predicting Human Interruptibility with Sensors: A Wizard of Oz Feasibility Study". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2003)*.
- Hudson, S. E. and I. Smith (1996). "Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems". *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 1996)*.
- Lai, J., S. Yoshihama, T. Bridgman, M. Podlaseck, P. Chou and D. Wong (2003). "MyTeam: Availability Awareness through the Use of Sensor Data". *Proceedings of Interact 2003*.

- Laptop sales beat desktops for first time (2003). *Associated Press*, July 2, 2003.
- Lu, L., H. Zhang and H. Jiang (2002). "Content Analysis for Audio Classification and Segmentation." *IEEE Transactions on Speech and Audio Processing* **10**(7): 504-516.
- Milewski, A. E. and T. M. Smith (2000). "Providing Presence Cues to Telephone Users". *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2000)*.
- Mynatt, E. and J. Tullio (2001). "Inferring Calendar Event Attendance". *Proceedings of the International Conference on Intelligent User Interfaces (IUI 2001)*.
- Nardi, B. A., S. Whittaker and E. Bradner (2000). "Interaction and Outeraction: Instant Messaging in Action". *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2000)*.
- Perlow, L. A. (1999). "The Time Famine: Toward a Sociology of Work Time." *Administrative Science Quarterly* **44**(1): 57-81.
- Schmidt, A., A. Takaluoma and J. Mäntyjärvi (2000). "Context-Aware Telephony Over WAP." *Personal and Ubiquitous Computing* **4**(4): 225-229.
- Smith, I. and S. E. Hudson (1995). "Low Disturbance Audio for Awareness and Privacy in Media Space Applications". *Proceedings of the ACM International Conference on Multimedia*.
- Tang, J. C. and M. Rua (1994). "Montage: Providing Teleproximity for Distributed Groups". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 1994)*.
- Tang, J. C., N. Yankelovich, J. Begole, M. Van Kleek, F. Li and J. Bhalodia (2001). "ConNexus to Awarenex: Extending Awareness to Mobile Users". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2001)*.
- Voida, A., W. C. Newstetter and E. D. Mynatt (2002). "When Conventions Collide: The Tensions of Instant Messaging Attributed". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2002)*.

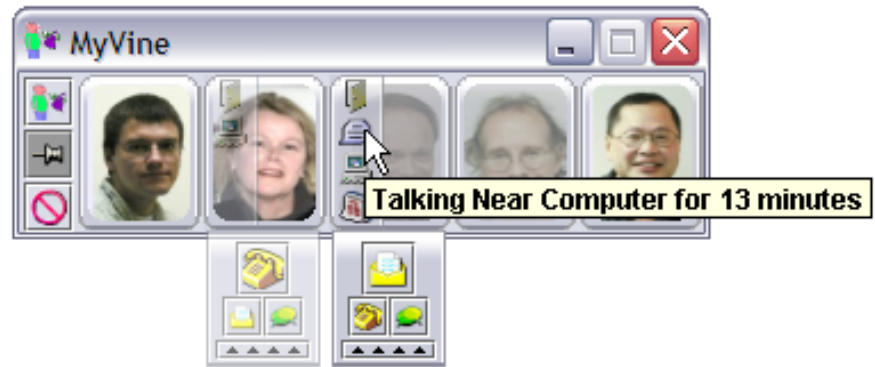


Figure 1. The MyVine client, with the mouse hovering over the details of person's context.

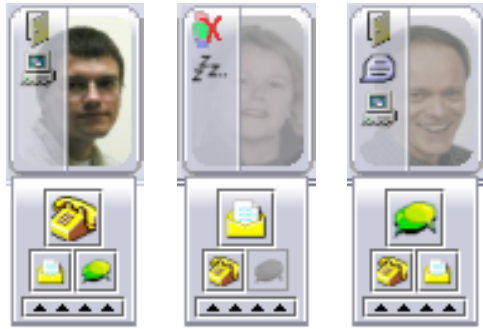


Figure 2. MyVine suggesting communication channels.






Figure 3. Distribution of availability shown during MyVine usage and the likelihood of examining a colleague's context by availability.



Phone use by availability



Instant messaging use by availability

Figure 4. Phone and instant message use by availability.

		Hovers	Instant Messages	Percentage	χ^2	<i>p</i>
Overall	Non-Idle	1325	86	6.50%	20.03	<i>p</i> < .0001
	Idle	1240	31	2.50%		
Non-Idle	Calendar Free	1110	77	6.94%	1.97	<i>p</i> ≈ .16
	Calendar Busy	215	9	4.19%	0.01	<i>p</i> > .93
	In Office	679	48	7.07%		
	In Building	514	37	7.20%		
	In Office or Building	1193	85	7.12%	7.32	<i>p</i> < .01
	Remotely Connected	132	1	0.76%	1.27	<i>p</i> ≈ .26
No Conversation	1230	77	6.26%			
	Conversation	95	9	9.47%		
Idle	In Building	559	26	4.65%	18.32	<i>p</i> < .0001
	Location Unknown	681	5	0.73%	17.34	<i>p</i> < .0001
	Calendar Free	902	12	1.33%		
	Calendar Busy	338	19	5.62%		

Figure 5. Influence of a colleague's context details on initiation of instant message sessions.