Lilsys: Sensing Unavailability

James "Bo" Begole
Sun Microsystems Laboratories
16 Network Circle
Menlo Park, CA 94025
bo.begole@sun.com

Nicholas E. Matsakis
MIT CSAIL, 32-G585
32 Vassar Street
Cambridge, MA 02139
matsakis@mit.edu

John C. Tang
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120
john.tang@us.ibm.com

ABSTRACT

As communications systems increasingly gather and propagate information about people's reachability or "presence", users need better tools to minimize undesired interruptions while allowing desired ones. We review the salient elements of presence and availability that people use when initiating face-to-face communication. We discuss problems with current strategies for managing one's availability in telecommunication media. We describe a prototype system called **Lilsys** which passively collects availability cues gathered from users' actions and environment using ambient sensors and provides machine inferencing of unavailability. We discuss observations and design implications from deploying **Lilsys**.

Categories and Subject Descriptors

H.5.3 [**Group and Organization Interfaces**]: Collaborative Computing—*synchronous interaction, CSCW*

General Terms

Design, Experimentation, Human Factors

Keywords

Availability, contact negotiation, context-aware computing, Instant Messaging (IM), **Lilsys**, presence, sensors, telecommunication

1. INTRODUCTION

Presence has emerged as one of the most compelling features of Instant Messaging (IM) systems. Most IM systems show whether a prospective recipient can be reached *before* attempting to call. While this capability is helpful to the caller, it potentially makes the recipient more susceptible to interruptions as others are more aware of times when she is reachable.

Presence information has begun to appear in other telecommunication media such as email and soon in mobile phones [7]. As presence information becomes increasingly pervasive, we need mechanisms to help understand not only whether an intended recipient is reachable but also how receptive she is to being called. This paper examines notions of presence and availability and describes problems related to determining availability in telecommunications. We describe the inadequacies of current strategies of proactively managing availability. We discuss a prototype system called **Lilsys** which uses passive detection of unavailability using ambient sensors, while preserving privacy by abstracting the person's context details into an unavailability inference.

CSCW'04, November 6–10, 2004, Chicago, Illinois, USA.

2. PRESENCE ≠ AVAILABILITY

In order for a caller to smoothly initiate a communication, the recipient must both be *present* to receive the call as well as *receptive* to communication at that time. These two components combine to form a person's general *availability*. We examine these two notions in more detail in this section along with the limitations of current strategies users employ to manage their availability.

2.1 Presence

Presence is generally used in today's communication systems to mean whether a person can be reached via a synchronous communication network. In practice, presence is typically equated with *device presence*, determined by whether a person's IM client is running or a mobile phone is in range of the cellular network. However, device presence does not always equate to *person presence* as devices may be left online while the owner is not physically nearby.

To help bridge the gap between the *device presence* that a system collects and the actual *person presence* that users want to know, IM systems typically indicate how long it has been since the owner last used the device; current or very recent use implies the owner is nearby. However, lack of use does not necessarily imply the owner is not nearby; durations of inactivity may occur when the person is present but talking to someone, reading, etc. Additionally, while device activity does indicate presence, it may also occur when the person is *most busy* and least receptive to interruption.

2.2 Receptivity to Interruption

Receptivity can be thought of as one's willingness to be interrupted. Interruptions are often considered to be a source of distraction that diminishes productivity. Indeed, Perlow documents the costs of interruptions among a group of software engineers and found a measurable benefit in reducing interruptions [6]. On the other hand, interruptions are also an inherent and useful part of accomplishing group goals, as Hudson *et al.* found in studying managers in a research organization [4]. The appropriateness of an interruption depends on many factors in the context of both the interrupted and interrupter: current task, the time of day, impending deadlines, relationship, topic, etc.

In comparison to presence, determining receptivity is a much more complex task. Clark [2] describes key elements of the subtle protocols humans use in face-to-face contact initiation: encounter, establish mutual eye contact, position bodies, and perhaps use gestures and other nonverbal cues to start a conversation. Many of these moves are offered tentatively so that they can be gracefully overlooked if no interaction is desired. Users of electronic communication generally do not have the necessary cues to smoothly negotiate contact the way it is done in face-to-face settings.

2.3 Availability Management Workarounds

Today, there are few tools to help a caller determine whether a prospective recipient is receptive to contact. This places the burden of managing availability on recipients who must proactively attempt to control access using only a few techniques. Even on networks that do not provide any availability services, one option is always possible: turn off the device. This simple approach entirely prevents interruptions from that device. However, users often forget to turn off the client, especially when they are busy, and to turn it back on, possibly missing important calls.

Another strategy is for the recipient to *screen* incoming calls. With caller-identification, a person can decide whether to accept an incoming call based on the caller's identity. Because telephone systems do not currently carry presence information, screening calls allows the recipient to use the cover of "plausible deniability" that he was perhaps not present or out of cell range. Screening an incoming message in an IM system is less possible because the caller can see if the receiver is present. Although a receiver can potentially not respond to an IM if she is too busy, many people find it socially difficult not to respond for reasons similar to those that make it difficult to ignore someone waiting at your door. In the end, screening calls does not in fact *prevent* unwanted interruptions because the receiver is still required to attend to each incoming call long enough to decide whether to accept it.

Another technique provided by most IM systems is to allow users to set their status to "away" along with explanatory text. In some systems, setting one's status to "away" blocks incoming communication, though in others it does not. Although many users take advantage of proactively setting their status to "away", a large portion of users do not use it or do so inconsistently.

Users may also manage availability by maintaining multiple addresses, phone numbers or IM accounts and selectively giving them out. Although the above proactive availability management techniques can help reduce unwanted interruptions, they may also reduce *desired* interruptions. Other problems with proactive management are that a user may not anticipate all times when she will be unavailable, may not take the action to block incoming calls (especially at times when she is most busy), or may not remember to unblock calls once she becomes available. For these reasons, proactive management offers only a limited solution.

3. LILSYS: SENSING UNAVAILABILITY

An alternative to proactive management by the recipient is to give callers the context information they need to assess availability, as in face-to-face communication where both parties engage in negotiating availability. Prior work [9] has found this strategy useful, although some users express concern that it may expose more details about the recipient's context than they are comfortable sharing. One proposal to address such concerns is to hide the context details and present an inferred abstraction of the prospective recipient's availability with the aim of providing the caller with the salient information (the recipient's availability) without exposing more details about the recipient's context than necessary.

To explore the practical issues of inferring availability in an office setting, our research group prototyped a system called **Lilsys**, which adds an array of physical sensors to the keyboard and mouse activity monitors already used in **Awarenex**, a research awareness and communication system [9]. To simplify deployment and provide a single place for the hardware user interface elements, the sensors and acquisition computer were combined into a single unit, shown in Figure 1. Mounted on each box were sound and motion sensors, with phone and door switches attached via wires. In addition to the sensors, we added a timer switch which allows users to override the system inference and set their state to the maximum unavailability level. In addition, **Lilsys** provides a switch labeled "online/offline" which allows users to turn off the reporting of sensor data if they desire, making the system blind to their context.

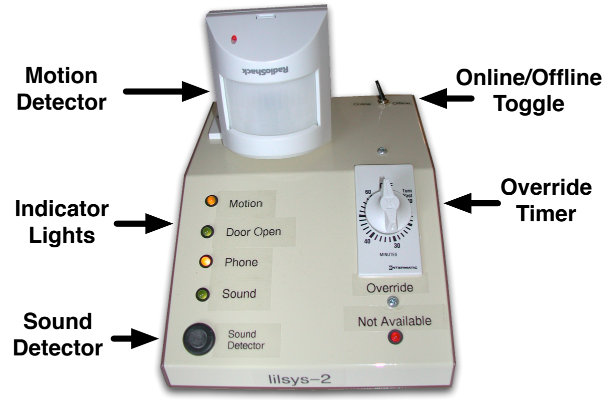

Figure 1: Lilsys sensor and data acquisition module.

We based our selection of sensors on the results reported by Hudson, *et al.* [5], who simulated using sensors to detect unavailability and found that indicators of social engagement were the most salient predictors of non-interruptibility. An inferencing model using only the single simulated sensor to detect when someone was speaking had an accuracy of 76%, arguing for the inclusion of speech detection. Fogarty, *et al* [3] used a laptop microphone as a speech sensor along with device activity and calendar in a system called MyVine. **Lilsys** inherits device and calendar information from its integration with **Awarenex** and adds speech, door and phone sensors to the inference. Although door state was not found salient by the subjects in the Hudson, *et al.* study, other studies suggest that door state can convey a considerable clue about availability [1]. Finally, **Lilsys** includes a motion sensor which is combined with speech and device activity to detect a person's presence.

All of the sensors are "off-the-shelf" and consist of binary switches (phone and door) or sensors that are modeled as binary switches (motion and sound). The sound sensor is a voice-activated switch of the type used to control tape recorders and is not capable of recording sound. An internal Dallas Semiconductor TINI computer accesses the switch states through a 1-Wire interface and communicates with the rest of the system using TCP/IP over ethernet.

We found it necessary to filter the raw motion and sound sensors to more accurately capture phenomena of interest. For example, the motion sensor might be triggered by someone walking past an empty office, falsely indicating presence. The raw data were filtered by integrating the activity over time and changing state when specified threshold values were achieved. We found that these filters reduced spurious activations caused by the short periods of silence and rest that are common in conversation or presence. However, this filtering imposes an undesired latency on detection. Whereas a human observer would immediately be confident that a conversation has begun, the sensors in **Lilsys** cannot sense the social context of detected sound, nor even whether the sound is speech or a drawer closing. It is therefore necessary to wait for sustained activity before concluding that a conversation is underway.

Figure 2 illustrates the information flow from data sources to the client interface. When sensors change state, events are sent from the data acquisition module to an inferencing engine which continually assesses the person's *presence* by combining data from the motion detector, sound, phone and keyboard/mouse activity. It also assesses *unavailability* by combining the sound, phone, and door sensors. **Lilsys** uses a simple Decision Tree for the inferencing model because it is easy to implement and, in a comparison of a number of techniques, Hudson, *et al.* [5] found the Decision Tree to have the highest level of accuracy at 82%. Changes are posted to a presence service which propagates the changes to clients.

The inference is displayed in an **Awarenex** contact list, as seen in Figure 2. **Lilsys** uses a U.S. traffic sign metaphor to suggest a person's inferred unavailability to a prospective caller: neutral (no inference), a diamond yellow "warning" sign indicates *possibly unavailable*, and a triangular, red-bordered "yield" sign indicates the person is *probably unavailable*. Hovering the mouse over an icon displays the full text description of the icon. We deliberately choose vague symbols for two reasons. First, to reflect the fact that the predictions are not 100% accurate. Second, we did not want the system to overly discourage necessary interruptions; The indicator should not too strongly convey unapproachability so that an observer can determine for herself whether and how to heed the indicator based on the circumstances of the interruption.

4. OBSERVATIONS & IMPLICATIONS

Lilsys has been deployed in four users' offices from September 2003 to March 2004. The system's user population also consists of ten **Awarenex** users who have one of the four "**Lilsys**-enabled" people on their contact list. Examination of usage logs along with solicited feedback from both user types has revealed several implications about portraying unavailability in user interfaces.

4.1 Unavailability Icons

With a goal of allaying privacy concerns and providing some amount of "plausible deniability", **Lilsys** presents an abstraction of the person's context rather than the full details. This abstraction raises the question of how best to convey availability status based on an uncertain machine inference.

Users had mixed reactions to the traffic sign metaphor. As expected, the vagueness of the "yield" and "warning" icons left some users unsure of their meaning. Some users ignored them and one reported that they made him "nervous" about communicating with the recipient. On the other hand, half reported that they would wait to communicate with a recipient who had either sign showing.

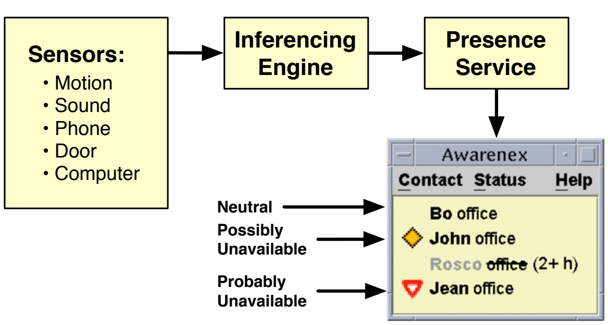

Figure 2: Lilsys system data flow and client interface.

4.2 Hardware Interface Elements

To allow users to easily stop the **Lilsys** monitoring at any time, we included an "online/offline" toggle switch. Although logs indicate occasional usage, users report that they have not used this switch other than for testing. One user reported an incident where it had accidentally been set "offline". Nevertheless, one user reports that he finds the presence of this capability reassuring. This suggests that future systems continue to provide an "online/offline" switch, but one that is not as easy to accidentally set.

Lilsys also includes an "Override: Not Available" timer which users can turn to explicitly set their status to "probably not available". Logs indicate occasional use, but users report that they only used it to test or demonstrate the functionality, and did not use it otherwise. This is consistent with the general non-use of proactive availability management techniques. In contrast to the "online/offline" toggle, no user reports finding the mere presence of the override useful or reassuring.

In addition, most users had some negative reaction to one or more characteristics of the off-the-shelf sensors we used, such as awkward appearance, the audible ticking of the timer and the audible clicks that the sound and motion sensors make.

4.3 User Image Control

Lilsys users expressed discomfort at being portrayed as "unavailable". They also reacted to our initial design which used a "do not enter" symbol. They perceived that symbol as too strongly prohibitive, so we changed to using "yield". Taken together, these reactions may indicate a social desire to appear approachable (at least to close colleagues) even under adverse conditions.

People weigh the significance of context cues differently. For example, some people only close the door when they do not want to be disturbed, while others may close it to block outside noises. Another example is differing opinions on the degree to which being on the phone but not talking indicates unavailability: is the person listening intently and therefore less available or has she "tuned out" of the call and is more available? These reactions suggest that the system provide more user control over what conclusions are inferred and how they are represented.

It remains to be seen how best to balance the recipient's desire for control with the caller's desire to understand the recipient's context. Putting control of accessibility entirely in the hands of the recipient is not consistent with face-to-face communication where the caller and the recipient share the context and where impropriety may be felt by either the recipient *or* the caller. For example, the caller might prefer not to call someone in a restaurant, even if the recipient is quite willing to take such a call. If the recipient can customize his appearance, the caller must be able to trust that the appearance is appropriate from the *caller's* point of view as well as the recipient's.

4.4 Asymmetric Capabilities

Only a subset of **Awarenex** users have sensors monitoring them. It was necessary to differentiate those users so that observers could assess the veracity of the presence information. In **Awarenex**, only device activity is used to detect presence, which is not a good indicator of non-presence. Using **Lilsys**, with the addition of motion, sound and telephone activity, non-presence is more certain. Therefore, the interface more definitively indicates when a **Lilsys** user has left a location by striking through the location label. For non-**Lilsys** users, the interface only strikes through the location label when they log out entirely. Because of the different semantics of this interface element, it is important to differentiate the different user types. Any number of means could be used; we simply use a different background color for **Lilsys** users.

4.5 Passive & Active Availability Management

The tradeoffs between proactive status setting and passive collection of context information are complementary, such that an interface can include both types. Whereas proactive status is obtained only when the user remembers to set it, passively collected information is updated as soon as a change is detected. Thus, proactive status information is only sporadically available and possibly stale, while the passively collected data are obtained frequently and are fresh. On the other hand, a person's intent and state of mind cannot entirely be determined from passively collected sensor data but a proactively entered status description can make such things clear.

4.6 Diminished Interruptions?

Finally and most importantly, users report that interruptions still occur even when the strongest unavailability icon is indicated, regardless of topic or urgency. This is consistent with the findings of Fogarty, *et al.* [3] who also found that the number of interruptions was not lower regardless of the indicated interruptibility level. However, our interpretation of this result differs somewhat in that it does not necessarily indicate that users are not using or respecting the availability information. Rather, **Lilsys** users report that others seem to use the availability indicator to "shape" the interruption. For example, the caller might say, "I see you're busy, but I have a quick question," or "can you call me when you're free?"

Such tentative approaches are analogous to how people behave in face-to-face communication as described by Clark [2] and Tang [8]. Someone who has traveled just a short distance down the hall may interrupt even if the recipient appears busy. In that case, the caller may merely indicate a desire for future communication, using phrases that acknowledge the extent of the intrusion. Humans have developed social protocols to smooth out potentially awkward interactions allowing us to interrupt gracefully despite not being entirely sure of each other's state of mind. Mutual awareness of the extent to which a caller has intruded on a recipient allows both parties to be aware of the extent to which the recipient has accommodated the intrusion, which feeds back into the parties' social relationship and shapes future interactions. While **Lilsys** did not perceive a decrease in the number of interruptions, some users report a qualitative improvement in the interruptions that occur.

5. FUTURE WORK

We are designing a new version of **Lilsys** based on recent wireless sensor technologies and incorporating the lessons learned from the initial deployment, described previously.

Another aspect we hope to explore is the ways in which availability continues to be negotiated during the course of a conversation, not only at its inception. After contact initiation, participants continue to reassess their desire to continue the discussion. They convey this desire through multiple cues: phrasing, tone of voice, prosody, body position, restlessness, attention to each other, etc. In addition to analyzing the voice signal alone, as explored by Yu *et al.* [10], other cues of conversation engagement may be detected via ambient sensors. Future work will explore analyzing sensor data to determine the extent to which engagement and ongoing availability can be detected and conveyed during remote communication.

6. SUMMARY AND CONCLUSIONS

Instant Messaging, and increasingly other telecommunication media, provide "presence" services that indicate users' reachability for synchronous communication. Presence information potentially leaves people more open to undesired interruptions which can only be avoided today by the recipient taking proactive steps, such as selectively connecting to the communication network, screening calls, or explicitly setting her status to "away". In contrast, both parties in face-to-face communication negotiate contact at the time of an encounter using subtle cues and expending minimal effort or pre-planning.

As an alternative, we have presented **Lilsys**, a prototype system which passively detects certain cases of lower availability through machine interpretation of ambient sensors. We described an interface for integrating inferred unavailability into an IM contact list. After a small-scale deployment, usage logs and user interviews found a number of issues regarding the portrayal of unavailability assessment in the client, hardware interface elements, user image control, asymmetric capabilities, and the integration of passive and active availability management in a telecommunication client. Most intriguingly, users report that interruptions seem to occur with no less frequency despite the unavailability assessments, but in many cases the interruption is tempered by acknowledgment that the recipient appears less available. Users report appreciating this qualitative improvement in the manner of interruptions.

With today's telecommunication technologies, the entire burden of managing accessibility and availability rests on the recipient. Our continuing research is exploring techniques that more closely model face-to-face situations where the act of initiating contact is a negotiation involving both the recipient and the caller.

7. REFERENCES

[1] W. Buxton. Living in augmented reality: Ubiquitous media and reactive environments. In Finn, Sellen, and Wilber, editors, *Video Mediated Communication*. Erlbaum, 1997.

[2] H. H. Clark. Chapter 18: Language use and language users. In G. Lindzey and A. Elliot, editors, *The Handbook of Social Psychology*, pages 179–229. Harper and Row, NY, 1985.

[3] J. Fogarty, J. Lai, and J. Christensen. Presence versus availability: The design and evaluation of a context-aware communication client. *International Journal of Human-Computer Studies (IJHCS)*, to appear.

[4] J. Hudson, J. Christensen, W. Kellogg, and T. Erickson. "I'd be overwhelmed, but it's just one more thing to do": Availability and interruption in research management. In *Proceedings of CHI '02*. ACM, 2002.

[5] S. Hudson, J. Fogarty, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee, and J. Yang. Predicting human interruptibility with sensors: A Wizard of Oz feasibility study. In *Proceedings of CHI '03*. ACM, 2003.

[6] L. A. Perlow. The time famine: toward a sociology of work time. *Administrative Science Quarterly*, pages 57–81, 1999.

[7] C. Saunders. Presence and availability: Coming to a mobile phone near you. *IT Management, Jupitermedia Corporation*, Apr 22, 2003. published at http://itmanagement.earthweb.com/entdev/article.php/2194831.

[8] J. C. Tang. Approaching and leave-taking: Negotiating contact in computer-mediated communication. *ACM Transactions on Computer-Human Interactions (ToCHI)*, under review.

[9] J. C. Tang, N. Yankelovich, J. "Bo" Begole, M. Van Kleek, F. Li, and J. Bhalodia. ConNexus to Awarenex: Extending awareness to mobile users. In *Proceedings of CHI '01*, 2001.

[10] C. Yu, P.M. Aoki, and A. Woodruff. Detecting user engagement in everyday conversations. In *Proc. 8th Int'l Conf. on Spoken Language Processing (ICSLP '04)*, Oct 2004. to appear.