Ubiquitous Awareness Spaces

Michael Boyle
Department of Computer Science
University of Calgary
Calgary, AB T2N1N4 CANADA
+1 403 220 3532
boylem@cpsc.ucalgary.ca

ABSTRACT
In this paper, I describe the present course of research we are pursuing related to the design, development, deployment and evaluation of ubiquitous and reactive environments for supporting casual interactions among intimate collaborators. I begin by briefly describing the work done in this area by us and by others and give some motivation for the design decisions made. I then step back from the problem to re-evaluate the designs, and question the suitability of ubiquitous computing for supporting tele-awareness and facilitation of informal interaction. I examine problems relating to privacy in the ubiquitous media space, and expose them as failures resulting from the drive for a seamless user interface. I conclude by generalizing these problems to other applications of ubiquitous computing, and revisit previously held tenets about how reactive environments should be designed and question their sanctity in light of the problems identified.

Keywords
Ubiquitous media spaces, awareness, informal interaction, video mediated communication

INTRODUCTION

Goals

Our present research at GroupLab, the human-computer interaction laboratory at the University of Calgary, examines the design and implementation of groupware systems that support smooth and graceful transitions between awareness and interaction among distributed groups. The technique we are presently focusing on uses reactive media spaces [3] for supporting casual contact among intimate collaborators. By intimate collaborators, we mean groups (usually small) of people with already established and rather close working relationships. Specifically, group members share a common need to maintain informal awareness of one another to coordinate their activities such that they may take advantage of serendipitous opportunities for informal interaction.

Previous work [3, 4, 5, 9] in media spaces has pointed out a fundamental tradeoff between the benefits of tele-awareness and the threats to privacy and solitude that come with revealing information about oneself in greater and greater fidelity. My own work frames these privacy issues in the context of casual interactions among intimate collaborators. We suspect that privacy issues are relaxed somewhat under the constraint of intimacy, and thus if adequate solutions for the constrained problem cannot be found, then there stands little hope for finding adequate awareness/privacy compromises in more public ubiquitous media spaces.

Boyle, M. (2001) **Ubiquitous Awareness Spaces**. *Yellow Series Report 2001-682-05*, Department of Computer Science, University of Calgary, Alberta, Canada.

http://www.cpsc.ucalgary.ca/grouplab/papers/index.html

Methods

Our first work in this area looked at building digital yet physical surrogates for signalling awareness cues and facilitating interaction. We built the Active Hydra [8], an integrated analog audio/video-enabled camera-display-speaker-microphone unit equipped with physical proximity sensors to control its operation (figure 1).

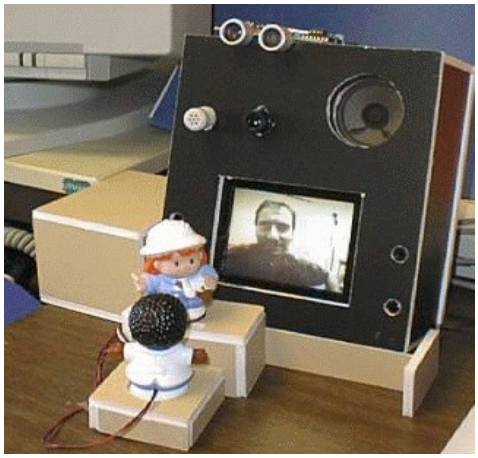

Figure 1. Active Hydra Unit

The Active Hydra was a reactive, physical surrogate for office sharing-like interaction. It was presumed that one nears the camera-display when one is interested in interacting with a remote participant. Thus, as one nears the Active Hydra, it would vary its operation from off, to snapshot-only video, to full audio-video. Physical proximity served as a subtle and implicit yet socially natural cue from which the device inferred its operation.

From this, we transitioned into the multimedia digital domain, to better support the low-cost ubiquity obtainable with inexpensive desktop digital video cameras and microphones, and to explore more sophisticated designs that preserve privacy. In particular, I have been working on

using distortion filters for balancing awareness and privacy (figure 2). In a study [2] run on the blur and pixelize filters, we found that participants were able to find a range of levels to which the blur filter may be applied such that awareness and privacy were appropriately balanced.

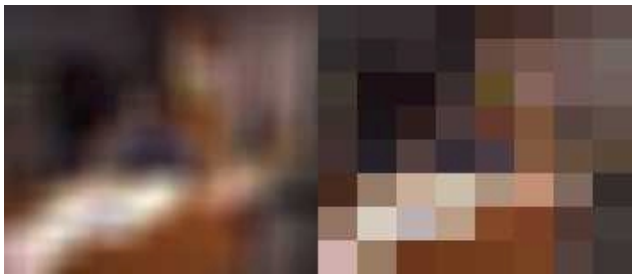

Figure 2. Blur and pixelize privacy-preserving filters.

Subsequently, I have been designing a video media space that modulates frame rate and size and extent of blurring to provide a tailorable balance between awareness and privacy. Moreover, building upon the experience with the Active Hydra, I am combining these privacy-preserving quality-of-service modulations with feedback obtained via physical proximity sensors and simple image processing techniques to build a reactive video media space that implicitly tailors the awareness/privacy tradeoff according to the sensory input. Lastly, during our evaluation of the blur and pixelize filters we found that no matter how well the filter performs, participants will always want to have the option of turning the media space off. Put another way, participants will always need a disconnected mode of media space operation. This somewhat contradicts the goal of making the media space reactive; after all, how can it react when it is turned off?

Thus, our media space tool provides a blocked mode of operation. When I block the video, I see the back of a hand superimposed over the entire video image of you that I see. Because reciprocity is enforced in our system as a measure to help people respect each other's privacy, you will see the image of a palm of a hand superimposed over the image of me that you see (figure 3).

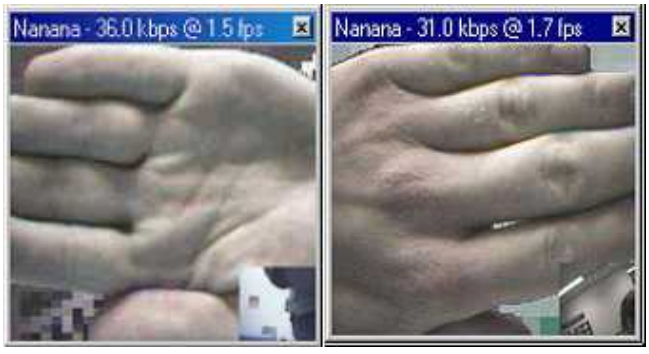

Figure 3. Super-imposed hands to block video.

Rather than having users scrounge around on their busy computer desktops for a button to toggle the block, we instead use a more gestural interface: covering the lens of the camera so the image goes dark for a few frames toggles the block. We feel this is a more socially natural interface, as it is common for people to reach out to cover the camera lens when, in dire circumstances, they do not wish to be recorded.

Meanwhile, through other work in our lab on the Notification Collage [6]—a public information display of ephemeral information—we gained valuable insight into the practical implications of deploying a media space among intimate collaborators. The Notification Collage included a snapshot video client, one without the aforementioned privacy-preserving filters or any sort of reactivity (figure 4).

Figure 4. Notification Collage Snapshot Video

The Notification Collage was deployed internally to our group, but in usage situations that included not only private offices and semi-private shared workspaces, but also private home offices. The present course of research will bridge the two threads of research to design, implement, deploy, and evaluate a reactive video mediated communications client in the context of a semi-private ephemeral media space.

TAKING A STEP BACK TO RE-EVALUATE UBIQUITY

Our own experiences with snapshot-only video and practically implementing a reactive video media space have identified a number of problems regarding the issue of privacy in ubiquitous video media spaces that have prompted us to take a step back and evaluate the fitness of ubiquitous and disappearing user interfaces for supporting informal awareness and casual contact.

First, we have learned that even private home and work offices are not genuinely private spaces, and thus not all media space participants may be part of the core group of intimate collaborators. In our workplaces, we occasionally have visitors stop by; in our homes, we have our spouses and children to contend with. Neither of these two groups have the a-priori relationship nor need to maintain awareness that distinguish intimate collaborators from other forms of work relationships.

Visitors to our semi-private laboratories and private work offices feel no sense of ownership of the space they inhabit, thus they quite often feel obliged to accept the presence of the video media space and the fact that their presence and activities may be broadcast to others whom are essentially strangers. The presence of cameras, microphones, sensors, and other surveillance technologies makes the semi-private space more public. Typically, these visitors are already appropriating themselves in socially acceptable manners, and so while they may feel uncomfortable at the prospect of being observed, the consequences of surreptitious observation by parties unknown to them are not severe.

In contrast, our family members feel very much a sense of ownership over the home environment: it is truly a private space. In the case of the Notification Collage, the home office was set up in a guest bedroom, and thus there was always the risk of spouses, children, or even guests walking in to the home office, when unoccupied, in various states of undress. These incidental third parties, much like visitors to workplaces, gain no benefit from any reciprocity enforced by the surveillance systems because they have no need for the awareness provided by such systems. Indeed, such practical conditions serve only to amplify the perceived intrusiveness of the surveillance technology and the potential for privacy to be severely compromised.

This has prompted us to re-evaluate the awareness technology in terms of who gets the benefit and who incurs a threat to his privacy. This threat/benefit issue is similar to Grudin's general work/benefit issue in CSCW applications [7], and has been examined in detail by Bellotti in her discussion of Xerox EuroPARC's RAVE and Apple ATG's Oh La La Cafe media spaces [1]. By introducing ubiquity into the design of an informal awareness system, we also introduce the chance—nay, likelihood—that there will be participants who must put up with a threat to their privacy yet will accrue no benefit from the system.

Second, we have observed an air of distrust of the technology from the standpoint of participants—incidental or otherwise—-who are not involved in its design or implementation. This observation has been further substantiated by discussions with participants involved in other, independent media space efforts. In the case of our home office deployment, the spouse of a media space participant has been very vocal in her discontent with the presence of the technology in her home. As noted by Bellotti [1] it seems that because of issues such as screen-savers and poor visibility of the display within the camera's field of view, people cannot reliably tell if the system is recording. The lack of appropriate cues to signal the state and operation of the media space has engendered a sense of distrust. This has prompted the participant to turn the camera around so that it points out a nearby window when he is absent. Along a similar vein, we are told that it was commonplace for participants in a Portholes [4]-like audio-enabled media space to physically unplug the microphone when they wanted to cut off the audio—even though the user interface had a simple and readily accessible graphical button to toggle this operation in software. These accounts reveal the critical problem of providing sufficiently salient and appropriate feedback when making the user interface controlling media space operation seamless.

What is curious to point out about these trust issues is that people generally do not mistrust their colleagues (or even strangers) so much as they do the technology. Though the consequences that arise should their privacy be compromised are small, people generally have an extremely low tolerance of failures on the part of the system to protect their privacy. This matter is further complicated by the fact that the degree to which a circumstance threatens one's privacy varies widely with the individual's personality, the participants involved, and the timing, events, conditions, and the meanings behind them that make up the circumstances of a privacy-threatening situation. Bellotti pointed out that these privacy issues might in part be because actions and the intentions that drive them are disassociated in the digital domain [1]. Intentions, as she points out, are generally not made explicit in ubiquitous computing systems, and making them explicit is undesirable because it trades off lightweightness.

The trustworthiness issue is further exacerbated in reactive environments because it is difficult to find simple-to-measure cues from which to reliably infer expected system behaviour. The cues to be measured must satisfy a number of fitness criteria:

- *Provide a socially natural mapping between participant behaviour and media space control*: in the case of the Active Hydra, the use of physical proximity to the camera/display surrogate mimics the natural social behaviour of moving closer to another when engaging in conversation;
- *Be reliably measurable*: the readings taken from the inexpensive ultrasonic rangers used in the Active Hydra units suffer from severe noise; moreover, environmental factors, such as tables and chairs, often get in the proximity sensor's path, and thus lead to false readings;
- *Lead to consistently drawn inferences*: in the reactive video media space tool that uses a camera lens-covering gesture as a signal to toggle a block, we found that the inexpensive video cameras used sometimes over-compensate for brightness/contrast, and thus even though the camera lens is covered, the video fed into

the image analyzer is insufficiently dim to toggle the block;

- *Corroborate other feedback offered by the system*: in our reactive video media space tool, the camera must always be on, even when blocked or not broadcasting the live video, and as the camera has a small, green LED on it to indicate it is functioning, this more physical feedback conflicts with the superimposed hand image provided as feedback by the software; and,
- *Be salient and purposeful*: with the physical proximity sensors, we find that it is easy to unwittingly fall into the sensor's field of view even when one is uninterested in the remote party and is engaged in other activities.

In taking this step back, we've identified these problems:

- Private places are not always private;
- Threat/benefit dilemma;
- Trust issues;
- Saliency and fitness of feedback;
- Fitness of cues sensed in reactive environment and their meaning in the system;
- Getting the inferences drawn to be consistent and correct; and,
- Making intentions more visible without sacrificing lightweightness;

What's significant to point out here is that these problems are not unique to ubiquitous awareness systems, but may be generalized to all ubiquitous computing systems. For example, we can rephrase some of the problems identified above in terms of those that may afflict ubiquitous accessibility systems:

- Accessible times are not always appropriate for certain kinds of interaction;
- Intrusion/benefit dilemma—who must put up with intrusive technology and who gets the benefit?
- Do people trust that an intelligent ubiquitous accessibility system will route messages appropriately?
- What counts as appropriately salient and suitable feedback for the behaviour of the system when it routes messages so as not to intrude?

FRESH FOCUS ON SOLUTIONS

The basic question I pose here is: given the problems described, just how appropriate are ubiquitous and disappearing UIs for awareness, accessibility, and casual interaction applications? Ubiquity has a number of attractive features, including immediacy; must these come with a lack of control or an inability to find solace? Disappearing UIs are attractive for these applications in that they are lightweight; the keep this, need we endure ambiguity over state, operation, and control that builds a fence of mistrust around users?

Physical or tangible user interfaces may be able to help. I have already pointed out a few cases—such as turning the media space camera to point out the window—where physicality was used not only to control the media space state and operation, but also to signal it in a trustworthy manner. Physicality has a number of good qualities: persistency, immediacy, saliency (yet not overpoweringly so), are largely intuitive, simple yet meaningful, and appear to engender trust.

Next, we need to re-think interface visibility and lightweightness—the reactivity of reactive environments. Perhaps reactivity is desirable for controlling only certain aspects of the system. Consider the following: a video media space tool monitors the conversation in progress. When the conversation is deemed to have ended—perhaps audio analysis shows the participants stopped talking, or proximity sensors show one is very far away, or image analysis shows a decided lack of activity in the scene, suggesting absence—the reactive environment automatically rotates the camera 180° around to point at a blank wall. Upon return of the absent party, or resumption of conversation, they system does not automatically turn the camera back around, and instead the party must turn the camera around manually. To turn the camera around automatically runs a very high risk of violating privacy. Instead, only half the operation of the media space is made reactive: the other half, the more risky part, is left under manual control. The state of the system is always immediately visible, and is changed in an intuitive manner: the stepping motor used to turn the camera around makes noise as the conversation ends, and to turn it back one would simply use one's hand to turn the camera around.

Lastly, in the context of ubiquitous awareness systems, we must re-think the golden rule of reciprocity in light of the threat/benefit dilemma described earlier. Why should reciprocity be enforced when one party cannot be reasonably expected to gain benefit from it? What if we relax the reciprocity requirement and instead allow people to willingly choose to broadcast information about themselves to another without reciprocally receiving the same type of information in the same fidelity about the remote party? So long as the state and operation of the media space were sufficiently salient and understood by the user, and capturing takes place only with explicit consent, the system would merely be offloading the task of enforcing responsibility onto social protocol—arguably, where it belongs.

REFERENCES

1. Bellotti, V. (1998) **Design for privacy in multimedia computing and communications environments.** In P.E. Agre and M. Rotenberg (eds) *Technology and privacy: The new landscape*, MIT Press, Cambridge, MA.
2. Boyle M., Edwards C., and Greenberg S. (2000) **The effects of filtered video on awareness and privacy**, in *Proceedings of CSCW'00* (Philadelphia, PA, Dec 2000).

3. Buxton, W. (1997) **Living in augmented reality: Ubiquitous media and reactive environments**. In K. Finn, A. Sellen and S. Wilbur (eds) *Video Mediated Communication*, Hillsdale, NJ.
4. Dourish, P. and Bly, S. (1992) **Portholes: Supporting awareness in a distributed work group**. In *Proceedings of CHI'92* (Monteray, CA, May 1992).
5. Gaver W., Moarn T., MacLean A., Lövstrand L., Dourish P., Carter K., and Buxton W. (1992) **Realizing a video environment: EuroPARC's RAVE system**. In *Proceedings of CHI'92* (Monteray, CA, May 1992).
6. Greenberg S. and Rounding M. (2000) **The Notification Collage: Posting Information to Public and Personal Displays.** To appear in *Proceedings of CHI'01* (Seattle, WA, April 2001).
7. Grudin J. (1988) **Why CSCW applications fail: Problems in the design and evaluation of organizational interfaces**. In *Proceedings of CSCW'88* (Portland, OR, Sep 1988).
8. Kuzuoka H. and Greenberg S. (1999) **Mediating awareness and communication through digital but physical surrogates**. In *ACM CHI'99 Video Proceedings and Conference Extended Abstracts* (The Hague, NL, May 1999).
9. Tang J., Issacs E. and Rua M. (1994) **Supporting distributed groups with a Montage of lightweight interactions**. In *Proceedings of CSCW'94* (Boston, MA, April 1994).