Judging People's Availability for Interaction from Video Snapshots

Brad Johnson
Department of Educational Psychology
University of Calgary
Calgary, Alberta CANADA T2N 1N4
Tel: +1-403-220-5910
bfjohnso@acs.ucalgary.ca

Saul Greenberg
Department of Computer Science
University of Calgary
Calgary, Alberta CANADA T2N 1N4
+1-403-220-6087
saul@cpsc.ucalgary.ca

Abstract

Several groupware systems support casual real time interaction over distance by providing periodically updated snapshots of other people's offices. People then monitor these snapshots to determine how available others are for communication. In this research, we try to isolate what information people use from these snapshots to help them infer another's availability. Research participants examined video snapshots of people posed in typical office situations, and judged how available those people were for interaction. Our first result suggests that people have difficulty extracting information from these images unless image resolution was at least 128x128 pixels. Our second result indicates that people interpret stereotypic situations as indicating varying degrees of availability. In general, people are judged as less available when they are seen to be absent from their office, or in conversation with others. People are judged more available when they are in transition (e.g., entering or leaving a room), and when they do not appear to be working. People at work seem to portray a more ambiguous situation. However, all situations had a minority of people who interpreted the image quite differently. These results have implications on the design of both video and non-video based awareness and availability systems

Keywords: *Real time groupware, casual interaction, awareness, contact facilitation.*

1. Introduction

Casual real time interaction is an essential ingredient of group cohesiveness. Yet the bottleneck to rich spontaneous interaction is distance (Kraut, Egido and Galegher [11]), and distant-separated team members will be at a disadvantage unless a prosthesis that overcomes distance barriers is available. Consequently, many groupware researchers are designing interfaces that facilitate how distance-separated team members can establish real-time contact with one another across computers. A crucial component of these systems is that they provide information that helps people stay aware of who is around in their team, whether those people are available for conversation, and whether it is socially acceptable to initiate a conversation with them (Cockburn and Greenberg [5]). (Several types of these awareness and availability systems are reviewed in Section 2).

While demonstrably useful for supporting casual interaction, these techniques are also fraught with problems. These include privacy violations (Bellotti [2]; Lee, Schlueter and Girgensohn [12]; Hudson and Smith [9]), camera shyness (Lee, Schlueter and Girgensohn [12]), excessive requirements on technology (such as video cameras, active badges, and other 'esoteric' peripherals), bandwidth costs, scaling to large communities, and so on. We believe that some of these issues come about because we do not know precisely what information people require if they are to determine availability of others. Consequently, we either put too much information on the channel (resulting in problems with privacy, bandwidth, and/or scaling), or too little (resulting in inappropriate contacts and/or lost opportunities).

In this research, we try to isolate what information people use to decide whether another person is both available and interruptible. We want to articulate the types of information used, and to determine the weight that

people place on these types of information. We believe that this information is crucial for:

- understanding the prospects and limitations of today's innovative but admittedly ad hoc awareness systems;
- fine-tuning decisions for exactly what information should be captured and transmitted to others, as well as how that information should be displayed;
- creating new, principled designs for non-video awareness systems, where the systems are based upon the awareness needs that people find useful rather than a particular technology.

This paper is the first step in this research program. After providing a brief survey of related work on awareness and availability systems (Section 2), we focus on systems that display intermittent video snapshots (Section 3). These are now popular, are available via the World Wide Web, and have proven successful at portraying awareness (Dourish and Bly [7]; Lee, Schlueter and Girgensohn [12]). In the remaining sections, we present two studies based upon how people interpret awareness and availability from intermittent video snapshots. These studies ask and tentatively answer two questions: what image resolution is necessary to convey the cues people required to make availability decisions; and how do people judge availability given a set of typical poses?

2. Related Work: Awareness and Availability Systems

A variety of technological approaches and experimental systems have been developed that promote awareness and availability.

The majority of systems leverage desktop video cameras. *Media spaces*, for example, let team members observe each other's offices and public areas through continuous audio/video channels (Bly, Harrison and Irvin [3]). Through these, people can see who is around and what others are doing at remote sites (Abel [1]). Alternatively, *video glimpses* give a person a short video-only view into another's office, simulating what happens when a person walks down a hallway and glances into an open office door (Tang, Isaacs and Rua [15]). Eschewing full video, *video snapshots* provide one with periodically-updated still images of other people's offices (Dourish and Bly [7]; Lee, Schlueter and Girgensohn [12]). People typically select and display a matrix of these snapshots on their screen, so that they can glance into the remote offices of people of their choice and gather a sense of who is around their community. To preserve privacy (Bellotti [2]), these snapshots can be altered to show only the essence of the activity while masking the details. Hudson and Smith [9]), for example, use video differencing techniques to pixelate a person appearing within a scene so that only a coarse outline of that person is visible. Similarly, Lee, Schlueter and Girgensohn [12] provided a user-controllable option to add blur to outgoing images, which would mask the details of what the person in the office was doing.

There are also many computer-based techniques that indicate awareness. A variety of systems monitor a person's keyboard and mouse activity, and use that information to tell others that the person is present. This information has been displayed as a list of all people who were active within the last few minutes (Chang [4]), as sounds that reflect activity e.g., the sound of someone typing (Lövestrand [13]; Cohen [6]; Greenberg [8]), and as icons whose shape and color correspond with activity levels (Greenberg [8], Wax [18]). Another approach involves video differencing, where the computer compares successive video images for changes e.g., as would happen when a person moves in a room (Lee, Schlueter and Girgensohn [12]). Rather than sending the video image, the computer only transmits to others whether a person was present or absent in the room.

Awareness systems are also moving "out of the box". *Active badges* transmit location information as people move around a site (Want, Hopper, Falcao and Gibbons [17]). Ostensibly developed for automatic call-forwarding, a variety of software systems used this location information to tell others about people's whereabouts (Lövestrand [13]). For example, a map can be raised to show where people are located. Another technique gathers information from *input sensors*—motion detectors, instrumented furniture and telephones, heat sensors (Ishii et al [10]; Mantei et al [14]). Depending on how these are configured, sensors can capture general activity within a room (such as motion) as well as specific activity e.g., that a phone is off the hook or that the office door is open. This information can then be transmitted and displayed to others by a variety of means. In our own lab for example, we are developing sensor-equipped physical surrogates, each indicating the activity of the person depicted by the surrogate. An example is a figurine that rotates to face you when the remote person is active, and away from you otherwise.

Finally, *awareness servers* act as a central database that collects awareness information from a variety of sources, and distributes that information to the community (e.g., Walker [16]). These servers could simplify how the many technologies mentioned above could work together, and could provide various controls to mediate access and privacy to the information contained by them.

3. A First Step in Studying Contact Facilitation

Dourish and Bly [7] suggested that video snapshots were reasonably successful for contact facilitation. A set of small video snapshots taken from cameras in people's offices and updated every minute or so suffices to give people a sense of who is around and what they are doing. Its disadvantage is that screen real estate is consumed, privacy violations are easy (Bellotti [2]), and that it relies on video cameras. In order to understand if alternate strategies can provide the same information in a more concise and secure form, we decided to determine what information people found useful in these snapshots.

Intuitively, we expect that people's judgement of other's availability will rely on several attributes. Some of these are part of the video image e.g., whether a person is actually present. Others are part of a culture e.g., whether it is polite to interrupt at this particular moment; part of the history e.g., inter-personal relationships, past responses; and part of the current needs e.g., level of urgency. We consider the first of these in this paper: to discover the essential awareness information contained within a video image. Our specific questions are below. Subsequent sections describe the studies motivated by these questions and the results that answer them.

1. *Determine the critical threshold of image resolution necessary to convey the cues people require to make availability decisions*. The visual cues people use to make availability decisions may be difficult to extract if image resolution is reduced beyond some threshold (see Figure 2). Consequently, decisions or determinations regarding availability will deteriorate to guesswork.

   Answering this question has direct implication to the design of today's video snapshot awareness systems, as it would determine the minimum snapshot resolution that should be captured and displayed if the system is to be effective. Given advances in technology, this may seem a minor point. However, it is useful to know the *minimum* resolution required: sending very high resolution images may impact privacy if others are able to view details about a person's workspace. For example, we already mention how Hudson and Smith [9] deliberately pixelate awareness images in an attempt to preserve privacy while still providing awareness information.

2. *Determine how people judge availability from a set of typical poses.* The typical ways people work in their offices can be loosely categorized e.g., whether people are present or absent in their offices, whether they are entering or leaving the office, whether or not they appear to be hard at work, whether they are talking to others. We captured these situations as image poses (see Figure 3), and we hypothesize that people's judgement of availability will vary with these poses.

   Answering this question has implications for understanding exactly what awareness information we should be capturing, transmitting, and displaying to others. For example, Whittaker [19] presents pitfalls that users experienced with video-based awareness systems: the answer to these questions could help designers predict failures and successes occur.

4. Methods

4.1 Participants

Research participants were sixty-five university people, all with some computer experience.

4.2 Materials

Ten poses each of two male and two female actors in their actual personal offices were photographed with a digital camera for a total of 40 photographs. Figure 1 illustrates the ten categories of poses of one actor. Other actors had similar poses, and Figure 2 shows an example of how each actor appeared in a single pose. The poses were of stereotypic personal office situations, as noted below and in Figure 1.

General

1. Empty room

Transitions

2. Entering a room
3. Leaving a room

Work activity

4. Working at a desk
5. Working on a computer

In conversation

6. Telephone conversation
7. Face to face conversation

'Idle'

8. Staring into space
9. Eating
10. Standing at a bookcase

The original photo resolution was 640 x 480 pixels and 16 million colors. Photos were then digitally reduced to lesser resolutions: 16x16, 32x32, 64x64, 128x128, 256x256, and 512x512 pixels. Figure 3 shows one photo at various resolutions. Of course, the print quality of the


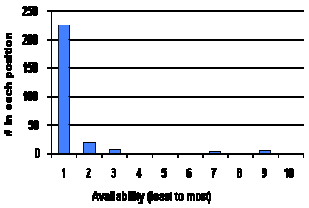

1. Empty room


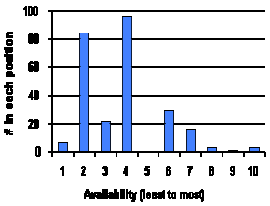

6. Telephone conversation


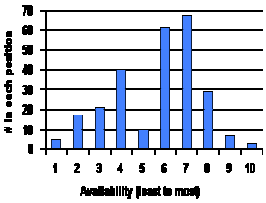

2. Entering the room


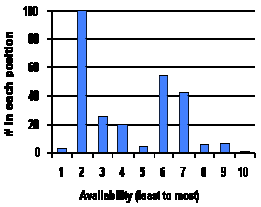

7. Face to face conversation


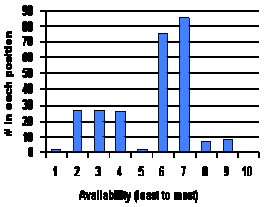

3. Leaving the room

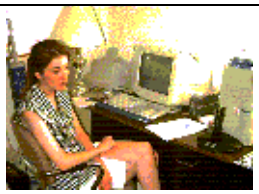
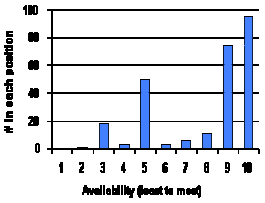

8. Staring into space


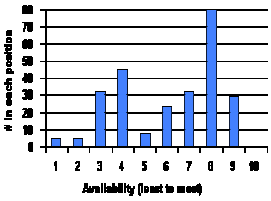

4. Working at the desk


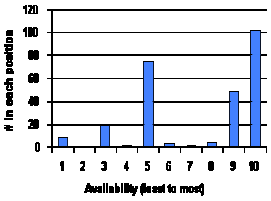

9. Eating


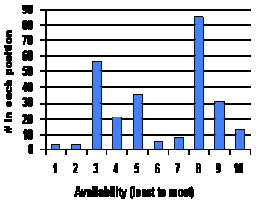

5. Working on computer


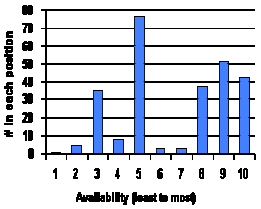

10. Standing at bookcase

Figure 1. The ten poses, photographed over one actor (printed here at low resolution). Accompanying histograms show pooled data for that pose by all actors for the sort performed in task 2. The X axis is the sorted position, from least available (1) to most available (10). The Y axis plots the actual number of images found in each position.

Figure 2. The four actors in the eating pose.

Figure 3. Various image resolutions of one actor working at their desk. The 32x32 and 512x512 images are not show. This printed version likely differs in quality from the ones used in the experiment.

Figure 4. The interface used to present the images. In this particular case, the subject is going through the 30th of 44 images in the 512x512 resolution condition. The subject can select an availability level from the 7-point Likert-like scale below the image. The progress bar at the bottom shows the number of images seen so far and how much is left to do.

grey-scale image in this paper will differ somewhat from the color image that appeared on the screen.

Figure 3 also shows that we displayed images by exact pixel resolution, where the image size shrank to represent the actual pixels used. This would contrast images scaled to fit a predetermined size. In those, the image size would remains constant, but several pixels on the screen would have to be used to represent a single pixel in the image.

Custom software presented screen sequences to participants, as illustrated by the screen snapshot in Figure 4. Each screen contained a particular image, whose resolution depended on the research condition. Participants used a 7 point Likert-like scale to indicate the perceived availability of the person in the image. The scale ranged from 1 (least available) to 7 (most available).

4.3 Procedure

Participants were randomly placed into one of six conditions corresponding to the six photo resolutions. Participants then completed the two tasks below.

Task 1. Using the software mentioned above, participants looked at and rated individual images according to how available they believed the person in the picture was. The short scenario included in Appendix 1 was given to subjects to provide them with the context for the task.

Task 2. Afterwards, participants were given the four sets of high-resolution printed pictures corresponding to the four actors, and asked to sort each set into a sequential order from most available to least available.

5. Experimental Design and Results

5.1 Critical threshold of image resolution

Design. We measured the *basis of agreement* in task 1 between any two people regarding whether or not the person in a particular image is available. Using a basis of agreement as a statistic assumes that people will generally base their decisions on similar cues or cue sets, and that different people would sort similar cue sets into similar positions or rankings. To determine the critical threshold of image resolution necessary for people to interpret awareness cues, we expect that agreement approaches chance below a particular threshold, with more agreement as image quality improves above it. The key is not necessarily high levels of agreement, because people may interpret similar images differently. Rather, we expect a statistically significant difference between agreement levels across this threshold.The design was one-way between-subjects ANOVA, where the independent variable was image resolution and the dependent variable was agreement, as scored by the Kappa coefficient (K). The Kappa statistic, designed to measure agreement between a number of raters on a set of criteria, corrects for chance agreement. Kappa values may take ranges from 0 (no agreement) to 1 (perfect agreement). Thus any amount greater than 0 is a measure of agreement, however weak that agreement may be. Participant scores were randomly paired in each condition and a Kappa coefficient generated. There were 5 pairs in each cell.

Results. Mean Kappa scores are plotted in Figure 5 for each pose. Scores indicated only modest agreement. A significant difference was found between groups $F(5,24)=6.389$, $p<.001$. Post-hoc test comparisons, using the Freedman test for consistency and/or reliability, showed that the difference occurred between the 64x64 and 128x128 conditions.

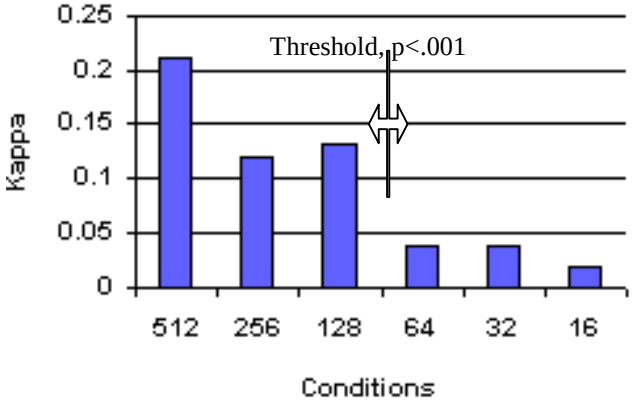

Figure 5: Mean Kappa scores. There is a significant difference between the 64x64 and 128x128 condition.

5.2 Differences between poses

Design. In the second task, each person sorted the sets of poses (grouped by actor) from most to least available. Every image thus had a position in the sorted pile, from 1 (least available) to 10 (most available). As this was an exploratory study, we chose a qualitative rather than statistical analysis. We plotted the pooled data for each pose as a frequency histogram, and we counted how frequently a pose appeared in particular position. We then examined the histograms for trends.[1]

Results. The histograms for the pooled data by pose are illustrated in Figure 3 next to its pose type (although the figure shows only one actor in that pose). The X axis shows the sorted position, from least available (1) to most

[1] We also produced histograms of data pooled by pose and by actor, to see if the actual actor had any effect on how people judged poses for availability. Differences, while present, were slight enough to allow us to consider the data pooled across all actors.

available (10). The actual number of images found in each position is shown on the Y axis.

6. Discussion

6.1 Critical threshold of image resolution.

When picture resolution dropped past some threshold, people found it difficult to make determinations regarding availability. From our results, it appears as if the threshold lies somewhere above 64x64 pixels and below or at 128x128. As we would expect, people's availability cues depend upon sufficient visual information being presented. These cues may be obscured when image quality is degraded.

The practical implication is that systems using a snapshot approach for displaying awareness information should use an image resolution of at least 128x128 pixels. Similarly, any study (such as this one) that has people examine images for availability should supply images of at least 128x128 resolution.

We should note, however, that the expected image quality in a real situation should be considered as well. Video images are degraded by poor lighting conditions e.g., glare from windows; dark rooms where people use only desk lights. Similarly camera placement may compromise the cues that should be visible in an image. Cameras may be far from the subjects, as with 'door cams', where cameras are mounted outside the office door (Lee, Schlueter and Girgensohn [12]), or too close to them to get a full view of the scene (as with cameras mounted on top of the video display). On the other hand, image processing techniques may enhance the intelligibility of the image. Dourish, for example, mentioned[2] that contrast enhancement made a difference in the images displayed in Portholes [7]. Similarly, Lee, Schlueter and Girgensohn [12] found that image sharpening and contrast enhancement compensated for defects in the frame-grabbed photos. Consequently, the 128x128 figure should be considered a reasonable but not exact estimate of the critical threshold of image resolution used in awareness and availability systems.

6.2 Differences between poses

Several poses indicated definite trends on how people ranked their availability. As we would expect, the vast majority (87%) rated the "empty room" (Figure 2, pose 1) as 1 (least available), with 97% scoring it 3 or less. Images involving telephone conversations (pose 6) were also rated as generally unavailable, where 80% judged them as 4 or less. On the other extreme, people generally rated the "staring into space" image (pose 8) as most available, where 65% scored it as a 9 or above. Surprisingly, 27% scored it as modestly unavailable (between 3-5), indicating a dichotomy of how people rated that pose.

This dichotomy in judgement rankings appears in varying degrees throughout the rest of the poses. For example, the way people interpret face to face conversation (pose 7) is similar to how they see telephone conversations: about half thought it indicated unavailability (a score of 3 or less), while 36% thought it indicated modest availability (a score of 6-7). The transitional act of entering or leaving a room (poses 2 and 3) also tended towards availability (60% scored entering between 6-8, while 61% scored leaving between 6-7), although as seen in Figure 2 about 40% of the people scored these poses as four or less. People's judgements on each of the other poses were split almost equally, again in a bimodal distribution. About half thought the "working" images (poses 4 and 5: at desk and at computer) indicated modest unavailability, with the other half interpreting it as quite available. Many thought that "eating" and "standing at bookcase" (poses 9 and 10) indicated high availability, although again a few thought this reflected unavailability.

These results suggest that although there may be trends on how people judge availability in these images, different people may interpreting the same pose or situation differently depending upon their social understandings. For example, while a majority thought that it was not appropriate to interrupt people in conversation, a few thought that this was acceptable e.g., some viewed the conversation in pose 7 as an informal one and thus interruptible. Similarly, activity transitions, such as when a person enters or leaves a room, were flagged as an indication of availability by most, but not all people. The same applies to "idle" activities, such as standing by the bookcase, staring into space, and eating. However, work activities, whether performed behind a desk or computer, generates a mixed response. Some thought it appropriate to interrupt that person, but others felt they should not disturb that person.

7. Conclusions and implications

This study provided some initial data points in terms of how people tend to view availability.

First, resolution is a factor when images are used to present awareness information. Implications here are straightforward. As a preliminary guideline, the image resolution should be at least 128x128 pixels. However, this assumes that the image was taken in reasonable

[2] Personal communications.

conditions. Poor images may require higher resolution, while image enhancement techniques may allow lower resolutions.

Second, people do interpret stereotypic situations as indicating varying degrees of availability. In general, people are less available when they are seen to be absent from their offices, and in conversation with others. They are more available when they are in transition (e.g., entering or leaving a room), and when they appear to be not working. People at work seem to portray a more ambiguous situation. However, these statements are not absolutes, as all situations had a minority of people who interpreted the image quite differently.

The way people interpret stereotypic situations has implications for all awareness and availability systems, regardless of whether they use video snapshots or not. Any system has to capture information from the environment and display it in a form that allows others to determine one's availability. The results indicate that some types of information (or situations) are more valuable for determining awareness than others. If critical situations and opportunities are not captured, it will be more difficult for others to determine availability. Conversely, capturing and displaying unnecessary information may raise privacy concerns (Bellotti [2]).

As a guideline, other interface devices, such as iconic indicators (Greenberg [8], Wax [18]), must be careful on how they capture and display these stereotypic situations. A straight representation of a particular activity may provide enough information to allow others to estimate awareness For example, icons that simply indicate that a person is present and active on their computer likely reflects the critical information found in pose 5. However, we must remember that this is only one of the many factors that people use to determine availability. Perhaps icons can be extended to indicate other information e.g., that people are in conversation, that they have just entered the room, and so on.

Another danger is that systems may try to infer and provide a measure of availability, rather than displaying what the person is actually doing. For example, availability could (perhaps) be displayed on a screen through an "availability" meter. However, this approach is error-prone for all but the simplest situations. As illustrated by the bimodal distributions presented in Figure 3 and discussed Section 6.2, people interpret particular situations differently. Thus there is no single measure of availability that a system could infer and present that would consistently reflect how people would determine availability of the actual situation.

This study has just scratched the surface, and there is much left to do. We need to appraise other stereotypical situations. We have to consider not just single snapshots, but a sequence of snapshots and how availability is judged over time. Social context must be included as well, as the relationships between the local and remote person will almost certainly dictate how availability is determined. We believe that this future work is necessary if we are to create principled designs for awareness systems.

Appendix 1. The scenario used for the study

Thank you for taking the time to participate in this study. During the next 15 or 20 minutes you will be asked to perform a small number of tasks.

Please read the consent form and if you agree to participate sign and date the form. One copy is for your records.

You are an administrative assistant in a large company. You perform jobs for a number of different people. In the course of your job you must periodically speak with your employers for clarification of various points. You may, for example, need to determine where employer #1 would like her reports sent. This means that you must determine when your employer is available to speak with you. To assist you the company has installed a video conferencing system that allows you to see snapshots of the offices of your employers. Based on your evaluation of these snapshots you will either interrupt them and ask your question or wait for a later opportunity.

The pictures you are about to see represent video snapshots of people at work in their offices. In the context of the above scenario rate how available the people are i.e. very available to not at all available.

Acknowledgements.

We are grateful to Mark Wylie, who photographed the actors in the particular poses, and developed the software that presented the poses to subjects. Drs. Jeff Caird and Carl Gutwin provided valuable input into the experimental design. This research was partially funded by Intel Inc. and by NSERC of Canada.

References

[1] Abel M. (1990). Experiences in an exploratory distributed organization. In J. Galegher, R. E. Kraut, and C. Egido (Eds.), Intellectual Teamwork: Social Foundations of Cooperative Work, pp. 489-510. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

[2] Bellotti, V. (1996) What you don't know can hurt you: Privacy in Collaborative Computing. In *Proceedings of HCI '96*, p241-261, Springer.

[3] Bly, S., Harrison, S. and Irvin, S. (1993) Media spaces: Bringing people together in a video, audio, and computing

environment. *Communications of the ACM* 36(1), p28-46, ACM Press.

[4] Chang, E. (1989) "Protocols for group coordination in participant systems." In M. Taylor, F. Neel, and D. Bouwhuis (eds) *The Structure of Multimodal Dialogue.* 229-247, North-Holland, Elsevier.

[5] Cockburn, A. and Greenberg, S. (1993). Making contact: Getting the group communicating with groupware. Proc ACM Conference on Organizational Computing Systems, ACM Press.

[6] Cohen, J. (1993) "Kirk Here:" Using Genre Sounds to Monitor Background Activity. In *Proceedings of InterCHI '93 Conference on Human Factors in Computing Systems*, p63-64, ACM Press.

[7] Dourish P. and Bly S. (1992). Portholes: Supporting awareness in a distributed work group. Proc ACM Conference on Human Factors in Computing Systems, pp. 541-547, ACM Press.

[8] Greenberg, S.(1996) Peepholes: Low Cost Awareness of One's Community. ACM SIGCHI'96 Conference on Human Factors in Computing System, Companion Proceedings, p206-207, ACM Press.

[9] Hudson, S. and Smith, I. (1996) Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In Proc ACM CSCW'96 Conference on Computer Supported Cooperative Work, p248-257, ACM Press.

[10] Ishii, H., Wisneski, C., Brave, S., Dahley, A., Grbet, M., Ullmer, B. and :Yarin, P. (1998) ambientROOM: Integrating ambient media with architectural space. In *CHI'98 Video Program*, ACM Press.

[11] Kraut, R., Egido, C. and Galegher, J. (1988) “Patterns of contact and communication in scientific collaboration.” In Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW ‘88), pp. 1-12, ACM Press.

[12] Lee, A., Schlueter, K. and Girgensohn, A. (1997) NYNEX Portholes: Initial User Reactions and Redesign Implications. In *ACM Proceedings of Group'97*, November 16-19, ACM Press.

[13] Lövestrand, L. (1991) Being Selectively Aware with the Khronika System. In *Proceedings of ECSCW '91 European Conference on Computer Supported Cooperative Work*, p265-277, Kluwer Academic Publishers.

[14] Mantei, M., Baecker, R., Sellen, A., Buxton, W, Milligan, T. and Wellman, B. (1991) Experiences in the Use of a Media Space Remote Synchronous Collaboration. Proceedings of ACM Conference on Human Factors in Computing Systems, p.203-208, ACM Press.

[15] Tang, J.C., Isaacs, E., and Rua, M. (1994). Supporting Distributed Groups with a Montage of Lightweight Interactions. In Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW ‘94), p23-34, ACM Press.

[16] Walker, W. (1997) Rapid Prototyping of Awareness Services using a Shared Information Server. Advanced Technology Group, Apple Computer www.research.apple.com/technology/proj/sis/

[17] Want, R., Hopper, A., Falcao, V. and Gibbons, J. (1992) The Active Badge Location System. *ACM Transactions on Information Systems*, 10(1) p.91-102, ACM Press.

[18] Wax, T. (1996) Red light, green light: Using peripheral awareness of availability to improve the timing of spontaneous communication. Proceedings of CSCW '96 Conference on Computer Supported Cooperative Work (Short papers), ACM Press.

[19] Whittaker, S. (1995) Rethinking video as a technology for interpersonal communications. *International Journal of Human-Computer Studies* 42, p501-529.