Studying Awareness in Contact Facilitation

Saul Greenberg and Brad Johnson
Department of Computer Science
The University of Calgary
Calgary, Alberta CANADA T2N 1N4
email: saul@cpsc.ucalgary.ca, bfjohnso@acs.ucalgary.ca

Cite as:

Greenberg, S. and Johnson, B. (1997) Studying Awareness in Contact Facilitation. Position paper for the *ACM CHI'97 Workshop on Awareness in Collaborative Systems*, organized by Susan E. McDaniel and Tom Brinck, Atlanta, Georgia, March 22-27.

Note: This paper represents work in progress, and as such some of the claims are still subject to analysis and interpretation.

Abstract

In this position paper, we use a controlled experiment to try and isolate what information people use to decide on whether another person is both available and interruptible.

A General Approach to Studying Awareness

Designers of groupware systems now agree that people's awareness of their surroundings and of other people must be supported, if the collaboration is to be natural. Yet "awareness" is a slippery term. While it is something all people have and use, it is not really clear how to operationalize it in software. Dictionaries, for example, define awareness as "having perception, knowledge or realization", and provide synonyms such as " informed" and "conscious". This give some clues about how information should be presented in a groupware setting, for it means that people must *perceive* certain types of awareness information in their environment, *know* how to interpret its meaning in their particular context, and *realize* how to apply that information. All these steps are required before people can use that information to make an informed decision or action. For groupware, the implication is that we must know what fundamental awareness information must be captured, transmitted and ultimately displayed at a person's site; whether that person can perceive the raw information on their screen; whether its meaning is interpreted correctly; and when and how that information is applied to a person's task.

Following from this, we have created a series of steps to help us study and understand how awareness can be applied to groupware.

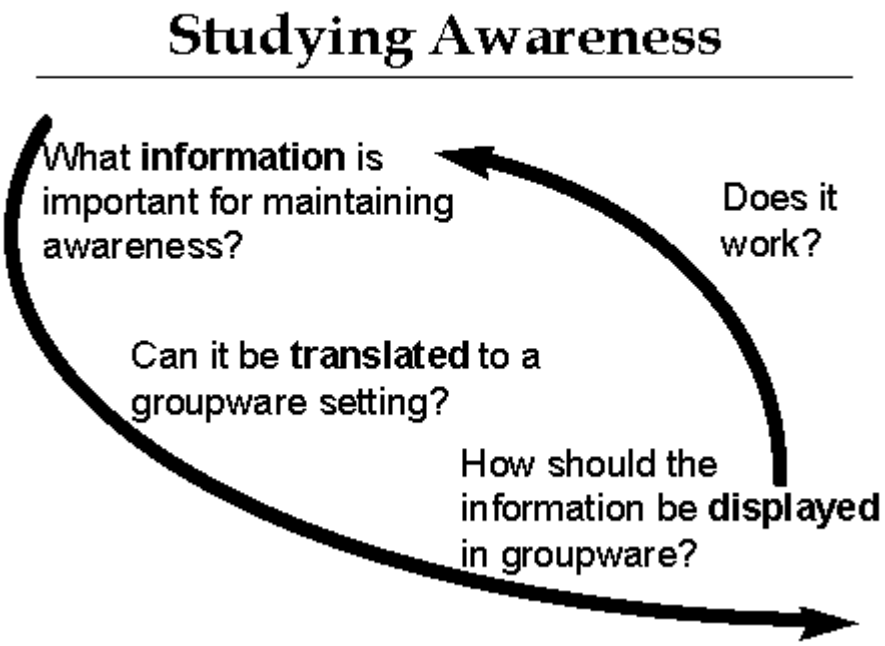

1. In the first step, we try to articulate the critical *information* people need to know to maintain awareness. This includes understanding how that information is obtained, and when/how it is used in practice. Ideally, this information could be used to form frameworks and theories of what people need to maintain awareness.
2. Assuming we know what information is required, the second step asks how that information must be *translated* to fit the constraints and pragmatics of groupware and its uses. For example, we have to see if we can capture the critical awareness information electronically, abstract it so it can be stored and/or manipulated, and transmit it down our channel.
3. Next, we need good interface mechanisms for *displaying* the critical awareness information. Here we ask if the information can actually be presented within our interface, and whether that information is of a form that can be perceived, interpreted, and used in light-weight and timely fashion. We also have to consider issues when displaying information such as the balance between providing awareness *vs.* distracting people from their tasks, and whether we are also violating people's privacy by giving out information too freely.
4. Finally, we have to know if our interface *works* in practice. The problem here is that there is a paucity of evaluation methods that deal directly with awareness. Because awareness is a cognitive activity, it is somewhat difficult to monitor. Good metrics do not yet exist. Because people are quite resilient at adapting their collaborations to fit whatever awareness information is provided, it is difficult to contrast "control" and "experimental" situations that provide subtle differences in what information is presented and how it is displayed.

Case Study: Contact Facilitation

Many groupware researchers are concentrating on designing interfaces that facilitate how people can establish contact with one other. These *contact facilitation* typically contain two components (Cockburn and Greenberg, 1993):

- the provision of information that helps people stay aware of who is around in their community, whether those people are available for conversation, and whether it is socially

acceptable to initiate a conversation with them.
- the ability to actually initiate a communication and/or collaborative session (i.e., a shared application combined with a voice channel)

Several interface techniques have been developed for contact facilitation. Video walls, for example, rely on continuous video and audio for showing people who is around at other sites (Abel, 1990). Video snapshots, such as Portholes, provide periodically updated video snapshots of other people's offices (Dourish and Bly 1992). Video glimpses, as in Montage, give short video-only views into one's offices. Minimalist awareness systems, such as PeepHoles (Greenberg, 1996), indicate how long its been since a person has been active on their computers. Unfortunately, these techniques are also fraught with problems. Recurring issues include privacy violations, excessive requirements on technology (such as video/audio), 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 for contact facilitation. 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). Because we do not know what information is critical, we do not have a methodology to solve these problems (which is why existing designs tend to rely on insight and creativity, rather than any methodological process).

In this research, we are trying to isolate what information people use to decide on whether another person is both available and interruptible. We want to articulate the types of information used, as well as the weight that people place on those types of information.

A First Step in Studying Contact Facilitation

Portholes (Dourish and Bly 1992) was a reasonably successful contact facilitation system. A set of small video snapshots, updated only every few seconds or minutes, was enough 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, 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 people's decisions to contact one another to 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); others are part of history (e.g., inter-personal relationships, past responses); and others on the current needs (e.g., level of urgency). We are now considering the first of these: the information contained within the video image.

We have designed a controlled experiment, where we assume that people use similar criteria to make similar decisions regarding another's availability. If this is true, then people should agree whether or not a particular scene denotes availability. That is, people should sort availability cues into similar categories of availability ranging from available to not available.

The experiment contains two independent variables: image resolution and pose.

Image resolution. As a picture is reduced in size, details of critical awareness information may be lost. We wanted to know what the critical threshold was for image resolution, which would show

up as decreasing agreement between people on how available others were. We rendered high resolutions images at various pixel resolutions (and held the pixels/sq. cm constant): 16x16, 32x32, 64x64, 128x128, 256x256, 512x512. For example, the figures below show the same image at 16x16, 64x64, 128x128, and 256 resolution.

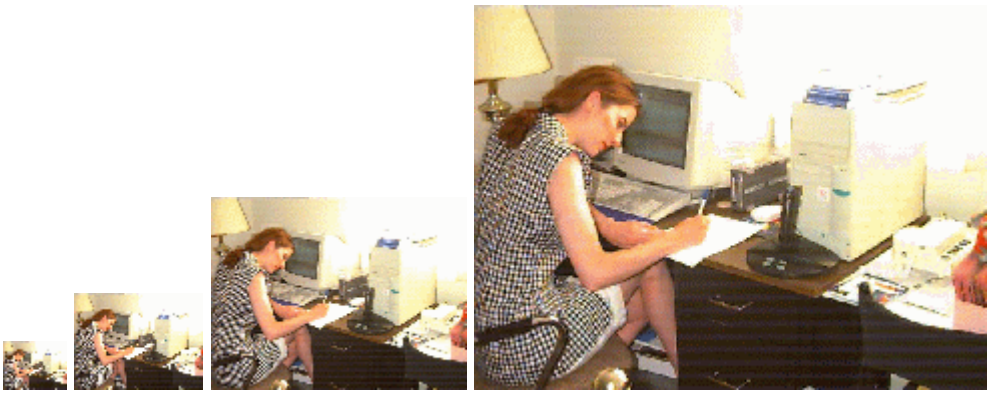

Pose. Snapshots of people working in their office can be loosely categorized according to poses. We wanted to know if people made different but relatively consistent choices about another's availability from these poses. We created 11 stylistic poses, and created four different sets of these showing four different people, providing 44 images in total. One set is provided below in its 128x128 version. The final "pixelated" pose was inspired by Scott Hudson's work at CSCW '96, and we included it for comparative purposes with its unpixilated version.

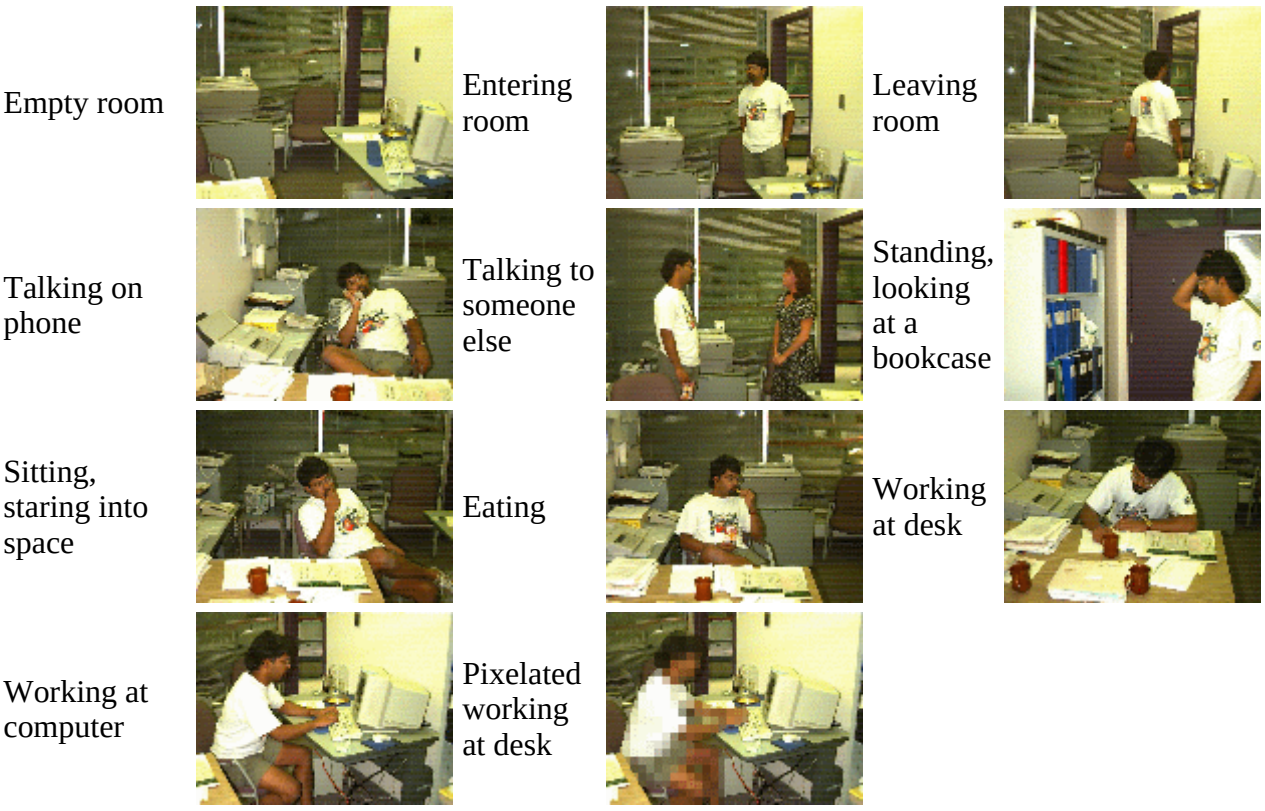

The experimental condition was between subjects for image size, and within subjects for particular poses. Research participants were given the images on a screen one at a time in pseudo-random order, and were asked to rank each according to the perceived availability of the person in the image (using a Likert scale). Afterwards, they were given single sets of images and were asked to rank order all images from most available to least available. This produced a forced choice situation.

Initial Results

We are in the middle of our analysis stage. These results are thus preliminary, and subject to change.

Image resolution.

A significant difference was found between groups F(5,24)=6.389, p< .001. The difference occurred between the 64X64 and 128X128 conditions. The histogram below plots the Kappa statistic (the level of agreement between subjects) against the image resolution. Although the kappa agreement level is fairly low, the difference across the 64 to 128 resolution threshold is a real one.

It was clear from the rating task that when picture resolution dropped below some threshold it became difficult to make determinations regarding availability. This threshold appears to lie between 64X64 pixels to 128X128 pixels. While it is too early to generalize, it makes sense that availability cues are dependent upon sufficient information being presented. These cues may be obscured when picture quality is degraded by reducing resolution. This outcome indicates that image resolution for judging availability should be at least 128X128 pixels.

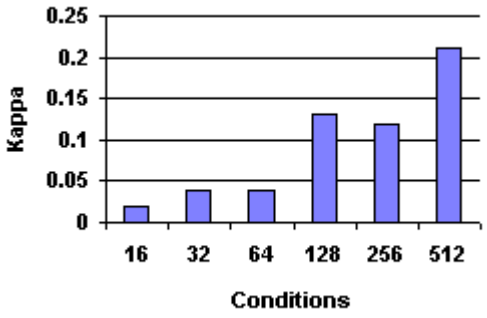

Sets of Poses.

When we chose four people to photograph for the poses, we did not consider that this in itself was an independent variable, with the person in the picture representing certain social cues. The force-choice sorting task proved a rich source of information regarding not only how people rank different activities in terms of availability, but also how they ranked different people doing the same things. For example, actor 4 was an older gentleman (compared to the other actors) and portrayed very definite cues through his body language. Comments from several participants suggested that they found him relatively unapproachable. In contrast, these same participants found the female actors more approachable. Our statistics also indicated that people sorted actor four's pictures in a more consistent fashion than the other actors. Thus a difference (although not statistically measurable) was seen between actors. This may simply mean that some people provide definite cues (in terms of body language and social status) while others provide more ambiguous cues.

Poses

From our forced-choice ranking, we noticed that people interpret different poses differently. In

general, the poses that indicated higher agreement according to their rankings were those that involved interaction with another e.g. talking with someone else either face to face or on the telephone. These poses also were ranked as being relatively unavailable. Poses that were ranked as being more available were those that involved some sort of transition from one working state to another. For example, poses that showed the actor either leaving or entering the room were ranked as relatively available.

Some poses generated a reasonable level of agreement. The first histogram, for example, shows the number of subjects that ranked the image of the "empty room" as indicating the least level of availability (as one would expect). However, many other poses generated dichotomous rankings. For example,we see in the second histogram that the pose 'working at a desk' generated rankings that clustered around position 8 (available) and around position 4 (unavailable). This suggests that people may follow two ways of interpreting the same pose, depending upon one's social understandings. In one view working at a desk suggests 'do not disturb' while in another view this clearly indicates 'available for interruption'. We also saw that the pixelated verison of an image tended to introduce more disagreement about that person's availability.

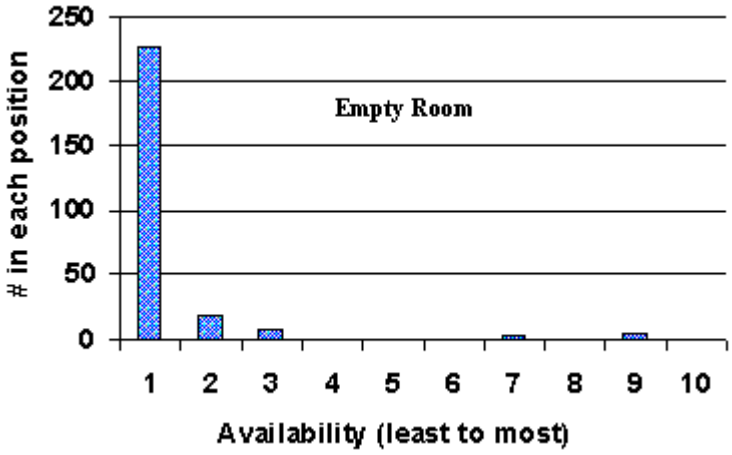

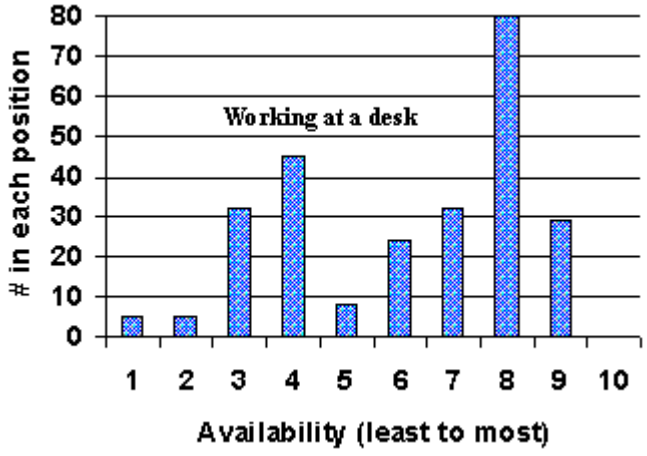

Conclusions

This is still a preliminary study, so it is somewhat dangerous to form definite conclusions. The trends analyzed so far indicates that people's agreement on others' availability in video snapshots is affected by:

- image resolution, with 128x128 snapshots being the threshold
- poses of people within those images, where people generally see transitions as favorable times to interrupt others
- apparent status and body language of the people, where people with definite body language and higher status seems to signal less availability

If these results hold, they could be applied to contact facilitation systems in several ways. First, for systems using video snapshots we now have an estimate of the minimum reasonable size of the displayed pictures. Second, we can consider ways to capture and abstract critical awareness information in non-video systems. These include a person's actual presence (which could be captured by motion detectors, active badges, or estimated by keyboard activity), whether they are engaged in conversation (perhaps captured by instrumenting phones and through active badges), and where they are working (e.g., computer-based work can be estimated by keyboard/mouse activity; seating at a desk by instrumented furniture) and so on. Third, we can consider how social status affects availability. For example, if a person indicates to the system their social relationship others (e.g., peer, friend, boss, etc), we could factor these into system estimates other's availability.

Of course, there are many variables not considered in this experiment, and many things left out in this discussion. However, the study does illustrate that we can begin to articulate what information people use to determine availability of others.

References

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.

Cockburn, A. and Greenberg, S. (1993). Making contact: Getting the group communicating with groupware. *Proc ACM Conference on Organizational Computing Systems*, Nov. 1-4.

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, May 3-7.

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.