In Proceedings of the Western Computer Graphics Symposium 2000, (Panorama Mountain Village, BC, Canada), March 26-29, 2000.

Awareness Projected: Moving Awareness to a Public Space

Michael Rounding, Saul Greenberg and Sheelagh Carpendale
Department of Computer Science
University of Calgary
Calgary, Alberta, CANADA T2N 1N4
+1 403 220 6015
[rounding,saul,sheelagh]@cpsc.ucalgary.ca

ABSTRACT

This paper concerns the design of public displays that project casual awareness information, where passerby's quickly acquire a sense of who in a small group is around and available. After describing what we mean by casual awareness, we present several prototype designs. In particular, we take awareness information garnered from a small group inhabiting a physical environment and abstract it onto a scene. We use the Three-Dimensional Pliable Surfaces visualization technique to emphasize people's particular activities.

Keywords: casual interaction, awareness, visualization, distortion viewing, public displays

1. INTRODUCTION

Casual interactions are the spontaneous and one-person initiated meetings and encounters that occur over the course of the day [10]. It happens so easily and naturally that it is something that most of us take for granted.

For casual interaction to work, people must be aware of what is going on in their environment: who is around, what others are doing, and seeing how available people are for conversation. People often gather this *casual awareness* information unconsciously or at the periphery of attention. For example, we hear the sound of an office door opening, which may indicate that someone down the hall has just arrived. Approaching footsteps indicate that someone is coming, and a quick glance (or our peripheral vision) will note whom it is. When we walk down a hallway we acquire a subtle but detailed record of "who is in" or "who is available."

Traditional casual awareness can only work when people inhabit a common space. Kraut, Egido and Galegher [10] correlated the number of interactions between people as a function of their proximity to each other: their results dramatically illustrate that the incidence of casual interaction drops exponentially with distance. People sharing offices have high amounts of casual interaction; people whose offices are separated by tens of feet have lesser interactions; and people in different floors of the same building have far less encounters.

While conventional workplace design often places potential collaborators in a common space (thus increasing casual interaction), this is not necessarily the case in new workplaces. Large distances may separate collaborators: across buildings, cities, or continents. Rapidly changing work demands may also require quickly changing relationships: a new co-worker required for a current project may be located on a different floor. Even if it makes sense for them to share a space, the cost and constraints may make this impractical.

Consequently, there has been much work in supporting causal interaction between pairs of distant collaborators, usually through media spaces or other information displayed on a personal computer screen e.g., [1]. In contrast, our research goal is to support casual interaction between small groups of people by projecting casual awareness information into a public space via a computer–controlled medium. We expect to present types of casual awareness information similar to what now exists in a physical work environment, where people can gather and maintain this information at the periphery of their attention.

We begin in Section 2 by exploring related concepts of awareness, including situation awareness and workspace awareness. In Section 3, we introduce a real scenario where casual interaction in our work-a-day world was breaking down: we will use this scenario to motivate our investigations. Section 4 then takes a detour into an information visualization technique called three-dimensional pliable surfaces. In subsequent sections, we show how this technique has potential to be used as a method to unobtrusively display casual awareness of a group.

2. AWARENESS

Awareness can be most simply defined as 'knowing what is going on in an environment' [6], which 'involves states of knowledge as well as dynamic processes of perception and action' [9].

More specifically, Gutwin and Greenberg [9] identify four basic characteristics of awareness.

1. Awareness is knowledge about the state of some environment, a setting bounded in time and space.
2. Awareness knowledge must be kept up-to-date: because environments change over time, a person must continually gather and update what they already know.
3. Awareness is gathered from the environment both through sensory perception and by actively acting upon the information acquired.
4. Awareness is rarely the primary goal of a person; rather, it is almost always part of some other activity.

There is a history of research in awareness, although particular researchers have tended to focus on specific situations and thus different awareness demands. Because much of it applies to casual awareness as well, it is worth reviewing them. Two particular sub-areas are visited here: situation awareness and workspace awareness.

Situation awareness can be defined as the up-to-the minute cognizance required to operate or maintain a system [6, 9]. The maintenance of this information can be broken down into three steps:

- being able to discern information in the environment that is relevant to the specific task being performed
- taking the perceived information and integrating it correctly into the previously existing knowledge of the system
- being able to predict how this information could change in the near future

Situation awareness has been explored heavily by human factors specialists who look at how highly trained machine operators act on awareness cues in an information rich and often stressful environment: jet fighter pilots, surgeons, and so on.

Gutwin and Greenberg [9] re-defined the work in situation awareness into a more specialized instantiation of awareness called *workspace awareness*, where several individuals require awareness to efficiently operate and interact with each other when using a shared visual workspace (such as tabletops, whiteboards and groupware drawing surfaces and editors). Workspace awareness is formally defined as 'the up-to-the-moment understanding of another person's interaction with the shared workspace' [9].

Gutwin and Greenberg [9] go on to describe how people maintain workspace awareness by using a modified version of Neisser's [11] *perception-action cycle* of how information in a physical setting is gathered and interpreted. Quoting from Gutwin and Greenberg: 'people gather perceptual information from the environment, integrate it with what they already know, and use it to look for more information in the workspace'. People gather information through three means: *consequential communication* (the presence of hands and bodies in the workspace); *feedthrough* (movement and changes to workspace artifacts), and from *intentional communication* between participants (both verbal and gestural).

These notions of situation and workspace awareness suggest a more formal way of viewing casual awareness: while we expect the perception/action cycle to remain the same, the sources of information will likely differ somewhat. This is still work we have to do, but we are particularly interested in how new arrivals in an environment acquire knowledge of what is going on, that is, how they get an entry point to the perception / action cycle.

To explain this further, we will describe in the next section a breakdown in an actual casual awareness situation. We will describe how the original breakdown occurred because of the lack of an entry point into awareness, and then how adding an entry point partly but perhaps inadequately repaired it. This will be used as a context to suggest an interesting and useful research direction.

3. SCENARIO: A BREAKDOWN IN CASUAL AWARENESS

We present a problem that we observed in our own physical workplace, a problem that we believe occurred because there was no 'entry point' to casual awareness information. We use this problem to describe a possible solution, and to motivate our research.

3.1 The Scenario and Its Problem

GroupLab members, mostly graduate students, occupy a shared room full of dedicated workstations (the GroupLab laboratory). The lab is set up in such a way that a person looking in from the doorway may see some but not necessarily all of the people who are present. The layout of the workstations could block their view, and there are two additional rooms attached to the lab that cannot be seen from the doorway.

Throughout the course of a fall semester, all of the members of the lab were involved in many different activities. These activities saw them coming in and out of the lab at unpredictable intervals. Some of the activities occupying the students' time included teaching assistant positions in undergraduate level courses. With these positions came frequent visits by undergraduate students seeking consultation outside of class time.

Undergraduate students would arrive in the lab looking for their teaching assistant. They would typically glance in and find them not within their immediate line of sight. For this reason, they sought the assistance of other members of the lab in locating the desired person. The members of the lab who had been interrupted would then inform the student that their teaching assistant was either present, or that they were not currently in the lab. In the latter case, they would often make an educated guess as to when or if the person sought after would return.

The problem here should be fairly obvious. Information that was immediately apparent to those already occupying the lab space was not apparent to new arrivals in the environment. The new arrivals had to poll occupants of the system for information to gain a fast entry point into the perception cycle. This polling interrupted the workflow of the lab occupants temporarily: they were forced away from their working line of thinking in order to summarize their current state of casual awareness of the desired GroupLab occupant.

As often happens, an ad-hoc solution was found. Near the end of the semester, the lab adopted a somewhat public notification of who was around. As pictured in Figure 1, below, it was basically a picture of each member of the lab, accompanied by textual descriptions of several places that they could possibly be or things that they could be doing. Using a post-it note stuck to the one of the descriptions, each person could communicate their state (i.e., in or out) and what they were doing (i.e. where they were) to anyone who was looking for them at any particular point in time. Of course, there is nothing new about this approach: we often see individual offices where its occupants place similar notices on their door.

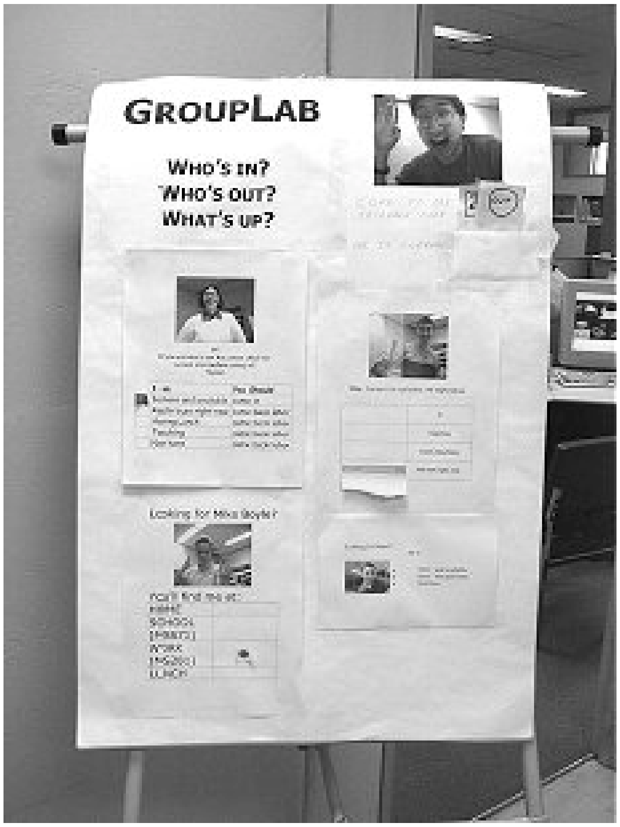

Figure 1: our laboratory's simple solution to the breakdown in casual awareness

This system reduced the interrupting traffic in the lab somewhat but, unfortunately, was still not very noticeable to visitors. It also required lab members to constantly maintain the information panel. Forgetting to do so meant that visitors had the wrong information (e.g., the person is shown as absent when they were actually there). This eventually led to mistrust i.e., even though the information panel would say one thing, visitors would still look inside to confirm it.

3.2 A Possible Solution

This example introduces several design challenges if we are to automate a system for casual awareness. First, we need to track and capture automatically casual awareness information about people: if the system requires explicit acts (such as pressing an 'I am here' button) it will likely fail. Second, we need to present the information about the state of all occupants in the lab in an effective manner: if visitors cannot immediately interpret the display to know what is going on, they will bypass it and go into the lab instead. Third, we need to display information captured over time, where visitors could immediately discern not only the state of the system, but also (if necessary) how it got to that state. For example, a GroupLab occupant may know that one of its members is around (because they have been in and out frequently in the last hour) even though they may not be present at that moment.

Our approach is to build a public artifact—where information held by the computer can be seen, heard, and acted upon in an easy and natural way by the people inhabiting the area around the computer [7]—for casual awareness. We want the public artifact to provide a fast entry point to the perception / action cycle. This could be similar to the "who's around" board mentioned above, although we would have to design it not only to gather and present awareness information, but also to fit well into the physical environment. We also have to design it to be highly visible but still unobtrusive, where those occupying and visiting the space can be gather information at the periphery of their attention.

This description is derivative of the notion of *calm technology* [12]. *Calm technology* does not demand our attention as other things do. The information calm technology may present or represent is available, but unless it is consciously observed, it remains largely a peripheral to the rest of the environment it inhabits [7,12]. The device recedes into the background until some event attracts attention back to it. Located in a public space – such as the lab discussed in the previous subsection – such a device can become a public artifact that is, in effect, ignored until either needed or until something happens that attracts attention to it [7].

While this idea sounds simple enough, there are a number of issues to be resolved. For example, placement of public artifacts is extremely difficult: where to put something so that everyone can see or use it can vary greatly depending on environmental factors. We have to recognize that a device's physical location can encourage or discourage public interaction. Notes placed on a whiteboard are considered more public than notes attached to computer monitors. Similarly, a workstation monitor is a far more private artifact than is a large interactive computer whiteboard (in essence a projected computer monitor). Making information public and visible encourages interaction [7].

Making awareness information public and available can make it easier for new arrivals in an environment to immediately determine its state. The information presented in public view definitely needs to be useful. As our scenario illustrated, not just any and all information can be explored at once. In essence, three questions need to be answered:

1. What information does an entire group need to see?
2. When and for what is the information important?
3. How should that information be conveyed?

Answering these three questions is at the core of designing a casual awareness system. Of course, the idea of calm technology suggests a useful mechanism for tracking and presenting casual awareness information, while the idea of public artifacts forces us to think how such devices fit comfortably within a public space.

Although this describes the properties of the device that contains casual awareness information, we still do not know how to present that information as a display. The next section suggests one approach. It details briefly a visualization technique called three-dimensional pliable surfaces, which we will then apply to present casual awareness information.

3. 3-DIMENTIONAL PLIABLE SURFACES

One possibility for visualization is the use of distortion techniques that magnify some information at the expense of other information. The purpose of this section is to just give a simple overview of the visualization technique. We begin by describing fisheye views and how it tries to solve screen real estate problems. Three-dimensional pliable surfaces are then presented as an alternate solution to the fisheye approach. As we will see in Section 4, we will use this idea to make particular awareness cues more salient on a public display.

3.1 Fisheye views

The small screens and poor resolution of existing computer display often means that they are inadequate for displaying large amounts of detailed information in a single scene. This is the "screen real estate problem;" and the challenge is how to display more information onscreen in a useful manner [2].

When viewing a set of information, centering or zooming in on a particular point in the data eliminates the context that existed between the focus of the zoom (the point of interest) and the whole space. As an example, taking a map of a city and zooming in on a particular subdivision eliminates the sense of that subdivision's location with respect to the rest of the city. This loss of context is the problem that distortion-oriented display techniques attempt to solve. One particular approach to maintaining focus plus context is *fisheye views.*

With fisheye views, all the data for a given data set is presented onscreen with a moveable focal point that magnifies the area underneath it in place using a fisheye lens effect (see Figure 2, below). With a large amount of detailed data present, displaying all of it at once in sufficient detail becomes difficult. Scaling to fit is inadequate. For example, text can quickly become too small to read. With a fisheye lens applied to a particular focal point, the area underneath the lens is magnified and thus visible in detail, while still showing its position and relation to other items in the global context.

Many fisheye implementations work best when the data being viewed is discrete, as they apply a fisheye effect only to the object under the focal point. When the data is one continuous entity (such as a map), a more sophisticated technique is required to achieve the same effect. One such technique is three-dimensional pliable surfaces, discussed next.

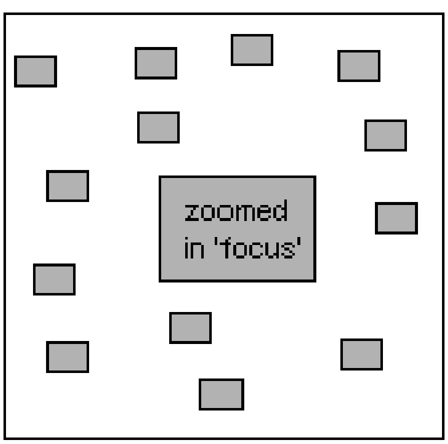

Figure 2: fisheye viewing

3.2 Pliable Surfaces

The three-dimensional pliable surfaces technique, developed by Carpendale et. al [2] applies a distortion effect to a two-dimensional surface by changing its shape from a plane to a three-dimensional surface. This visually distorts particular areas, which can be used to effectively draw attention to a particular region of some data representation while maintaining the context between the focus and its surroundings.

When applied effectively, the technique results in an increase in the amount of information that can be displayed usefully on a screen without getting lost. Accentuated distortions can easily provide a mechanism for drawing the eye to particular region. For example, we see a map in Figure 3A, and how the area in its centered has been distorted in the middle in 3B. The 3-d manipulation that stretches the surface effectively expands some regions and shrinks others. What is important is that the area under the lens is more visible. The grid lines in Figure 3B and the added shading in 3C are cues that help a person interpret what areas are distorted and what areas are undistorted.

To create the distortions, a two dimensional image is manipulated in three dimensions, pulling regions towards and pushing regions away from the viewer so as to accentuate certain elements within the image. This allows certain elements to be viewed smoothly in detail without actually losing their context within the whole of the information space.

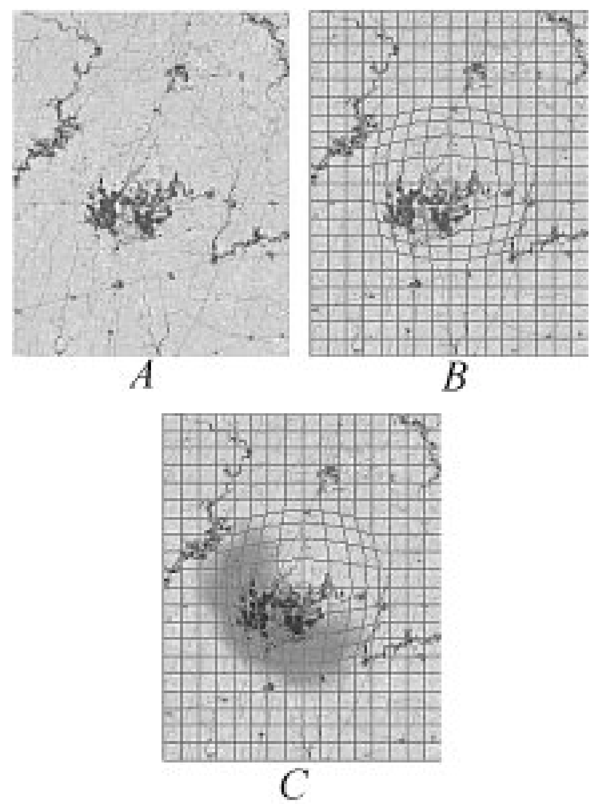

Figure 3: pliable surfaces: A: a simple distorted image. B: simple distorted image with an added grid. C: same image presented in B, with added shading.

A Gaussian curve is applied to a two-dimensional surface to create the distortion. The gently curving profile of the top and base of this curve makes it a good choice. When viewed in perspective, the fall-off of the Gaussian curve is compressed slightly. This is improved by using an auxiliary curve to adjust the

Guassian curve, making the middle of the drop-off much more gradual. Viewed from above, the distortion appears to be much less compressed than with just the Gaussian curve alone [2].

There are several advantages to this distortion approach that make it convenient and easy to use. For one, viewing such a transformed surface from above is akin to having a section of it pulled out of the screen, making the distorted region appear to be closer to the eye of the viewer. It is fairly natural for human beings to manipulate things in three dimensions because they are used to it. Pushing and pulling things makes sense in a real world; pushing and pulling regions in the described distorted world should also be easy to understand [2].

While some distorted images are easy to understand on their own, adding certain features to a distorted image can further imply where and what the distortion is. For example, adding a grid (e.g., Figure 3B) that contours the distortion outlines the bounds of the distortion, making it immediately more comprehensible. Furthering this idea, adding a light model to the model introduces shading and makes the distortion even more apparent to the eye (Figure 3C)[5].

In the next section, we outline the use of this distortion technique in a possible implementation of the casual awareness problem we have previously described.

4. AWARENESS PROJECTED

There are currently a few different ideas for presenting awareness information to others. Two popular approaches include ICQ programs and video media spaces.

Mirabilis' ICQ®, AOL's Instant Messenger®, the Microsoft® Instant Messenger®, and Greenberg's Peepholes [8] show a list of a person's preferred on-line contacts, their current login state, and how long it has been since they have touched their computer. Although this information is captured automatically by the system, people can explicitly change their status through simple control mechanisms. This adds a notion of control and privacy. For example, in Mirabilis ICQ a user can force their network status to "not available" or even to "invisible," thus granting them a certain level of privacy. While these systems are extremely simple and use surprisingly crude mechanisms to capture information, they are very effective and popular. This is because its users immediately go from having no awareness of their colleagues to having some awareness, even though the underlying awareness information is unsophisticated.

Another way to present awareness information is through the use of video media spaces, an always-on audio and video channel between two or more locations [1, 7]. Because colleagues can see each other through this channel, they can fairly accurately track their availability. What is especially appealing about these systems is that the same channel is used for communication. That is, people can act on awareness simply by entering into conversation with the person on the other side. Of course, media spaces such as these have privacy issues associated with them.

While both of these types of systems are successful in conveying awareness information, neither has been developed as a public artifact. That is, they are mostly designed to be used by one person at a time. In media spaces, for example, the pure video information could be too rich to be displayed publicly, as privacy concerns can become very serious indeed [7]. The next two subsections suggest a couple of systems that could be interesting to have displayed publicly. Both use pliable surfaces, as described in Section 3, as a visualization tool to represent awareness.

4.1 Small Group Awareness

The level of awareness provided by ICQ-like systems simply present people's state as a single data point, one per person. A given contact, listed by name or by image, is labeled as "connected," "away," "not available," etc. The first problem is that seeing who is around means scanning individual items in the list, which means that maintaining awareness information would become a foreground rather than background activity. As well, the discrete data points declaring status are crude. For example, in the default configuration of Microsoft Messenger, a person is displayed as 'present' if they have touched their computer within a half-hour (or whichever time delay is specified by the person). This label is thus, at best, just an estimate of presence, and there is much room for error. The problem is that we need better ways to represent people's presence.

We can use the pliable surfaces technique to present the information in a scene that can be interpreted at a glance, and to present the notion of presence as a continuous variable [7]. Our suggestion is to use a picture or series of pictures representing a group, where we magnify particular people by an extent that reflects their level of activity. For example, Figure 4A shows the group where no one is on line. In figure 4B we see that Michael's head is larger (front row right), and thus he appears somewhat more present than the others. This works for multiple people as well: in Figure 4C we now see that the person at the middle row far right is also present. The degree of presence can be marked by the amount of magnification: people on-line now would be magnified more than people who have not done anything for a while. They would thus stand out in the scene.

In practice, we would build the actual application so that a group could import an image into it, and specify areas of the image to represent different people. When running, the distortion lens could be applied to the surface of the image based on some measure of activity.

As a public display, we could show this image on a large display situated in a reasonable place. In our lab scenario mentioned previously, this could be a $1m^2$ flat panel display positioned on a wall or easel just outside our laboratory door for others to see. Unlike the current paper version, lab members would not have to update the display explicitly, and visitors or people just walking by could interpret the scene at a glance. While we recognize that some initial training / signage may be necessary to indicate to people what this new technology is doing, we believe that it will be easy to interpret afterwards.

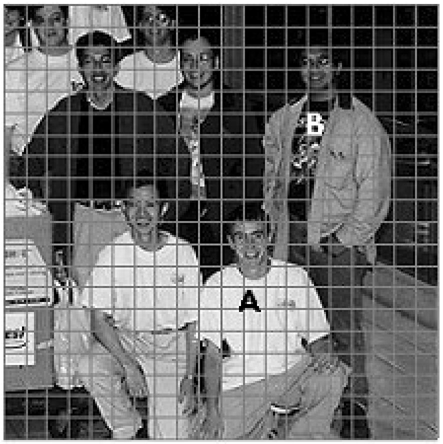

A

B

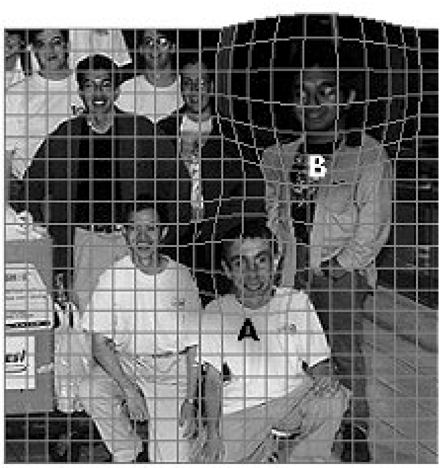

C

Figure 4: Distortion providing awareness. In frame B, person A has initiated some activity while in frame C person B has initiated some larger activity.

4.2 Awareness Maps of Physical Space

Another potentially interesting application of distortion techniques in an awareness application could be to use it to represent activity in a physical space on a two dimensional map. This sort of system could be implemented using motion sensors placed in locations relative to a floor plan.

For example, the representation illustrated in Figure 5 suggests some level of activity near the computer and a larger level of activity near the exit. Someone viewing this data would be able to determine that someone was probably using the computer and that someone was either entering or leaving the room. This later case would be clearly visible if the distortion effect were applied in real time, where the focal point would track the person as he or she moved through the room.

If the same technique were used to represent a lab, the proximity of people to computers or phones could be measured visually, as well as their level of activity (to a certain extent). For example, if a person were on a telephone we would see the telephone magnified as well. This could allow a visitor to make a more accurate decision about how to best go about interrupting or contacting the individual in the represented space. Of course, there are "Big Brother" implications here that have to be addressed.

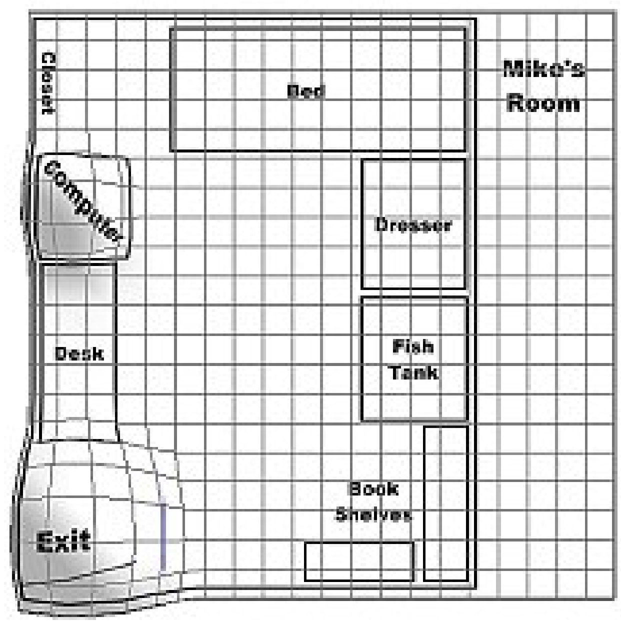

Figure 5: Activity in a physical space mapped and represented using a pliable surface.

5. CONCLUSION

This paper lays out a research direction. We summarized the casual awareness problem, and related how existing theoretic models of awareness can give us the intellectual foundations to understanding what is required of awareness systems. We then presented a particular problem in maintaining casual awareness, that is, how visitors and passersby's can discover the state of people inhabiting a common area. Using this as our motivation, we introduced and then proposed a casual awareness display that uses the notion of distortion-oriented techniques and public displays.

The work is still in the early stages. The pliable surfaces system now exists [2], but is not in the form that it can be readily applied to this problem. We need to extend its API, and to craft an interface that lets users specify individuals in an image. We need to control the degree of magnification in real time, where sensors or other approaches can drive the degree of magnification. We need to explore the device characteristics itself, and understand architectural issues of how devices can be perceived as public artifacts. Of course, we need to evaluate all this. It is unlikely that this design is the best or even an effective approach. We need to understand where it succeeds, and where it fails. We expect new and better designs to come out of this as well as an understanding of the design issues associated with casual awareness displays.

6. REFERENCES

[1] Bly, S., Harrison, S. & Irwin, S. (1993) Media spaces: Bringing people together in a video, audio and computing environment. *Communications of the ACM*, 36(1):28-47.

[2] Carpendale, M. S. T., Cowperthwaite, D. J. and Fracchia F. D. (1995) Three-Dimensional Pliable Surfaces: For Effective Presentation of Visual Information. UIST'95, Proceedings of the ACM Symposium on User Interface Software and Technology, pp 217 - 226, ACM Press, 1995.

[3] Carpendale, M.S.T. (March 1999) A Framework for Elastic Presentation Space – Simon Fraser University

[4] Carpendale, M. S. T., Cowperthwaite, D. J. and Fracchia F. D. (1997) Extending Distortion Viewing Techniques from 2D to 3D Data. IEEE Computer Graphics and Applications, Special Issue on Information Visualization, IEEE Computer Society Press, Vol. 17(4) pages 42 - 51, July 1997.

[5] M. S. T. (1997) Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. Making Distortions Comprehensible, IEEE Symposium on Visual Languages, pages 36 - 45, IEEE Computer Society Press, 1997.

[6] Endsley, M., (1995) Toward a Theory of Situation Awareness in Dynamic Systems, *Human Factors*, 37(1), 32-64.

[7] Greenberg, S (1999). Designing Computers as Public Artifacts. - *International Journal of Design Computing:* Special Issue on Design Computing on the Net (DCNet'99), November 30 – December 3, University of Sydney.

[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.

[9] Gutwin, C., and Greenberg, S. (1999) A Framework of Awareness for Small Groups in Shared-Workspace Groupware. - Technical Report 99-1, Department of Computer Science, University of Saskatchewan, Canada

[10] Kraut, R., Egido, C. & Galegher, J. (1988) Patterns of Contact and Communication in Scientific Collaboration – In Proceedings of the ACM Conference on Computer Supported Cooperative Work, 1988. pp 1-12.

[11] Neisser, U. (1976) *Cognition and Reality*, W.H. Freeman, San Fransisco.

[12] Weiser, M. and Seely Brown, J. (1996) Designing Calm Technology. *PowerGrid Journal*, Version 1.01, http://www8.electriciti.com/1.01/