Chapter appears in O'Hara, K., Perry, M., Churchill, E. & Russell, D. (Eds.) (2003). *Public and situated displays*. Dordrecht, The Netherlands: Kluwer Academic Publishers. This is the author's version of the work; not for redistribution.

Chapter 3

LARGE DISPLAYS FOR KNOWLEDGE WORK

Elizabeth D. Mynatt, Elaine M. Huang, Stephen Voida, and Blair MacIntyre,
College of Computing & GVU Center, Georgia Institute of Technology

Abstract: Knowledge workers, a growing component of our modern workforce, have specialized work practices not well served by existing office technology solutions. Our research in designing office systems that better support this class of workers has included an in-depth study of traditional whiteboard use and the development of three large display systems, each addressing different aspects of and work practices involved in knowledge work. We have identified three cross-cutting themes in making these kinds of large display applications effective for supporting a variety of information management activities: context-awareness, privacy and content relevance, and informal and casual interaction techniques. In this chapter, we reflect on the ways in which each of these themes guided the design of our system prototypes and how they may inform future efforts in integrating large displays into the office environment.

Key words: computer-supported cooperative work, large displays, knowledge work, context-awareness, privacy, informal interfaces, information visualization, peripheral displays, ubiquitous computing, office computing, whiteboards, groupware

1. INTRODUCTION

Our research seeks to design office systems that better support knowledge workers, the growing number of business professionals whose job it is to interpret and transform information. Knowledge workers are individuals who "put to work what [they have] learned in systematic education, that is, concepts, ideas and theories, rather than…manual skill or muscle" (Drucker, 1973). Successful knowledge workers manage multiple tasks, collaborate effectively among several colleagues and clients, and manipulate information that is most relevant to their current task by leveraging the spatial organization of their work area (Kidd, 1994; Malone, 1983; Mynatt, 1999). The diversity of these work practices and the complexity of implementing flexible computing tools make it difficult to meet all of these workers' needs.

Further complicating the design of systems for knowledge workers is the variety and distribution of tools and information in a modern knowledge work environment. In a typical office, workers use desktop and laptop computers, PDAs, and wired and mobile telephones. They also make use of non-electronic tools, such as filing systems, whiteboards, and desk surfaces. Sources of important work-related information are frequently distributed throughout the office environment, both within and beyond the individual office walls. Finally, colleagues often act both as additional information sources and as collaborators. It is up to the individual knowledge worker to make sense of all of these informational resources, coordinate and organize them so that higher-level abstractions emerge and the relationships among pieces of information become clear, and then act on this comprehensive understanding appropriately.

In our research, we have focused on leveraging large interactive surfaces, such as electronic whiteboards, as an effective platform for supporting spatial organization of information, task management, and collaboration. We have expanded the scope of this research over time from supporting informal manipulation of information on a whiteboard (Mynatt, Igrashi, Edwards & LaMarca, 1999, 2000), to supporting multi-tasking and background awareness with an integrated whiteboard and desktop system (MacIntyre, et at., 2001; Voida, Mynatt, MacIntyre & Corso, 2002), and to enhancing awareness and asynchronous collaboration of small groups of co-workers using a shared board (Huang & Mynatt, 2003).

Some of the recurring themes that have emerged in all of these projects are the importance of context-awareness in supporting knowledge work, the necessity of addressing user concerns about content relevance and information privacy on large interactive surfaces, and the usefulness of informal interaction techniques that match the flexibility of knowledge workers' practices.

2. THE WRITING ON THE WALL: A STUDY OF WHITEBOARD USE

Our initial research exploration in this area was completed at (then Xerox) PARC. This project focused on understanding the use of common office whiteboards in individual offices. Using observational data and interviews, we collected information about personal whiteboards from a diversity of participants, ranging from senior managers to visiting researchers and support personnel. Daily photographs of office whiteboards provided useful clues in understanding how whiteboard content changes on a day-to-day basis and proved incredibly useful for grounding subsequent discussions with study participants. Although whiteboard use has been studied in other settings (Abowd, 1999; Moran, et al., 1996) such as meeting rooms, classrooms, and production environments, our observations (Mynatt, 1999) indicate that whiteboard use in the office differs from its use in more public spaces:

- In contrast to meeting room or production environments, office whiteboards tend to be used for a heterogeneous set of tasks in parallel.
- Whiteboard content seems to group in natural clusters or segments. These segments may correspond to different tasks, different people, or writing at different points in time.
- Personal whiteboard content in particular seems heavily context-dependent. Although an office visitor would have great difficulty interpreting a whiteboard, its owner invariably recalls a substantial amount of context when working with each cluster of content.

These characteristics underlie our model of informal whiteboard use for knowledge workers and are relevant to shared whiteboards when the content can be persistent and owned by the group.

3. SUPPORTING INFORMAL WHITEBOARD INTERACTIONS WITH FLATLAND

Our next step was to create an augmented whiteboard, called "Flatland," to better support typical whiteboard use in an individual office (Igrashi, Edwards, LaMarca & Mynatt, 2000; Mynatt, et al., 1999, 2000). In our design, we extended the existing whiteboard look-and-feel with an interface that matches its role in informal office work. Using a touch-sensitive surface and projected display, our design goals were:

- To support a low threshold for initial use while making increasingly complex capabilities available. At the simplest level, Flatland should act like a normal whiteboard, so that users can walk up to it and write on it.
- To provide a look and feel appropriate for informal whiteboard tasks and distinct from production-oriented tools typically found on a desktop computer.
- To support informal whiteboard tasks such as to-do lists and sketching.
- To support clusters of content on the whiteboard. These clusters, or segments, may be created for different purposes, at different times, and by different people.
- To support the use of informal and context-dependent information. For example, content could be stored and retrieved based on its salient

context (e.g. time of creation, people present) instead of requiring a file name.

– To support the flexible management of a dynamic whiteboard space, such as freeing up white space for new input while maintaining the visibility of current content.

The conceptual building blocks in Flatland are called segments. Segments are clusters of whiteboard content that are automatically created based on grouping of input strokes, and that can be enhanced with special, task-specific behaviors, e.g. to-do lists, a calculator utility, and stroke clean-up or enhancement. The notion of segments mirrors the existing work practice of clustering like information together in blocks on the whiteboard surface, and allows for straightforward manipulation of the board's content in an intuitive manner (see Figure 3-1).

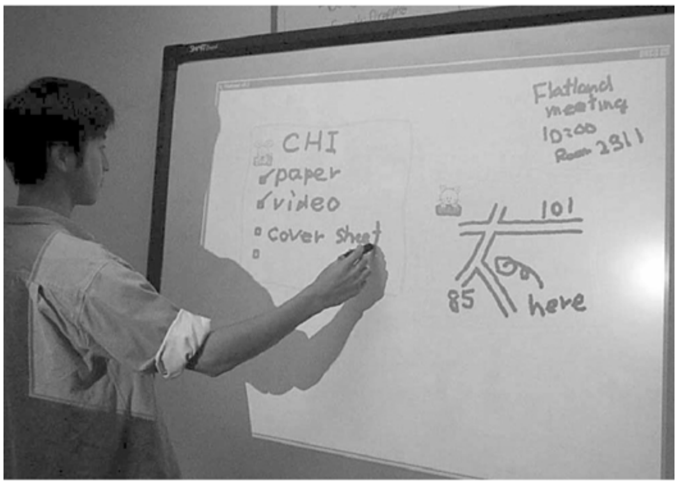

Figure 3-1. Typical use of Flatland.

Although evaluation of Flatland indicated that the system design successfully supported knowledge workers in leveraging an enhanced whiteboard, many of the subjects expressed an interest in integrating Flatland with other existing technologies such as desktop computers and PDAs. Moreover, just as traditional whiteboards are fluidly used for individual and small group tasks, Flatland also enabled informal collaboration in limited ways. However, the potential to leverage content associated with multiple people as a springboard for asynchronous collaboration and awareness remained.

4. LEVERAGING AN INTEGRATED WHITEBOARD AND DESKTOP WITH KIMURA

In our more recent work, we sought to augment and combine many of the common tools used by knowledge workers – desktop computers, whiteboards, and networked peripheral components – into an integrated, pervasive computing system called Kimura (MacIntyre, et al., 2001; Voida, et al., 2002). Moreover, we explored ways to make the content of the whiteboard context-aware to support individual and collaborative work.

Figure 3-2. The Kimura system in an office environment, including the monitor and peripheral displays.

Kimura separates the user's computer desktop into two regions: a focal region, displayed on the desktop monitor; and one or more peripheral displays, projected nearby on the office walls. Each of the user's work activities are associated with a unique virtual desktop that is displayed on the monitor while the user is engaged in the activity. Background activities are projected as visual montages on the peripheral displays, implemented here as an electronic whiteboard (Figure 3-2). The peripheral display allows users to monitor each ongoing work activity, transition smoothly between activities, access a wide variety of contextual information, and maintain awareness about relevant activity changes. The interactivity provided by the electronic whiteboard surface allows the user to informally annotate and spatially organize the visual montages.

From Kimura's point of view, a work activity—such as managing a project, participating in a conference, or teaching a class—is modeled as a cluster of documents and a collection of cues representing ongoing interactions with people and objects related to that activity. We refer to this cluster as the activity's *working context*. Each working context may have numerous documents—including text files, Web pages, and other application files. A working context may also include indications of ongoing activity, such as email messages without replies and outstanding print jobs. Kimura automatically tracks the contents of each working context and tags documents based on their relative importance. As in previous systems, such as Rooms (Henderson & Card, 1986), users define the boundaries of working contexts manually—in our case, by creating virtual desktops. We chose this strategy because these operations are easy for the user to perform and can be easily monitored to detect working-context changes, and because this strategy avoids relying on the system to infer these transitions.

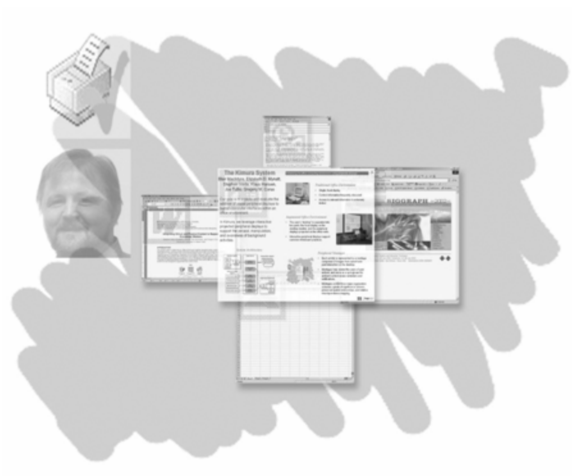

Figure 3-3. A montage of a working context, including a number of application windows and two external context notification cues, representing both virtual- (completion of a print job) and physical-context information (the availability of a colleague).

Each working context is displayed as a *montage* of images garnered from system activity logs (Figure 3-3). These montages are analogous to the "room overviews" provided by other multi-context window managers. But where these systems show the exact layout of the windows in each room, our goal is to provide visualizations of past activity in context. These visualizations are designed to help remind the user of past actions; the arrangement and transparency of the component images automatically create an "icon" for the working context. Additionally, montages can serve as anchors for the background awareness information that is gleaned from our context-aware infrastructure.

The montage design is intended to relieve the user of some of the cognitive load associated with remembering a large amount of information—information about each work activity and its related contextual information—and with synthesizing that information on the fly from a potentially overwhelming number of sources. Furthermore, they are designed to present this information without intruding on the user's focal activity and in a manner that specifically supports the needs of knowledge workers. Future work in this area can and should include further iterations on these kinds of work activity visualizations based on usability evaluation.

As in Flatland, Kimura's electronic whiteboard—the primary display surface for the montage visualizations—supports common whiteboard practices (Mynatt, 1999). This surface is also well-suited to supporting informal information management activities. Our system implementation incorporates existing electronic whiteboard interaction techniques with montages and notification cues (Hong & Landay, 2000; Mynatt, et al., 1999). This allows the user to annotate montages with informal reminders and to reposition montages to indicate the respective priority of background activities. Additionally, the whiteboard's large display area is an ideal, unobtrusive location to show contextually relevant information about the user's work activities and the context information sensed from around the office.

We have completed a preliminary series of user studies designed to evaluate the usefulness of various montage designs in supporting task resumption. The initial findings from these studies suggest that there may be some correlation between the layout of the document thumbnails within a montage and the ability of users to resume a postponed work activity represented by the montage. We have also run a number of psychophysics experiments to determine if a relationship exists between the speed of changes in the peripheral display and humans' ability to detect those changes when involved with a demanding focal task; analysis of the results of these experiments are forthcoming.

5. SUPPORTING COLLABORATION WITH SEMI-PUBLIC DISPLAYS

In addition to the use of large interactive displays to support the information needs of individual knowledge workers, we have also been exploring the use of these surfaces to support awareness and collaboration among group members co-located within a shared space. Like the Kimura system, the Semi-Public Displays project (Huang & Mynatt, 2003) aims to help users integrate and manage information from a wide range of sources; the primary goal of the system is to help group members maintain an informal work/social awareness of each other's activities by taking advantage of the various affordances of persistent, group viewable displays. The system aims to provide some of the properties or group viewing that public displays afford, but in the context of a shared space available only to individuals within a tightly-knit group. Therefore, while the information on the display is available to members of the co-located group, it is generally unavailable to passerby or people who are not members of the workgroup. As these displays are neither truly public in their availability, but are still group viewable within a set of people, we have termed them "Semi-Public."

Small group settings are often already equipped with the means to display information, such as projection displays, electronic whiteboards, or large monitors, but these tend to be used only in certain circumstances, such as during a meeting or a presentation. Our aim is to take advantage of these resources by using them as persistent sources of group information and shared workspace. However, determining what content and interactions are most effective for these large displays still remains a challenge.

The members of a co-located lab are likely to be highly aware of each other's activities; even so, we have identified several benefits that shared interactive displays can offer in a small group setting. By making certain types of relevant information persistent in the environment, these displays can provide information about group members, and foster coordination and collaboration. We have found them to be a potentially effective medium for making information from other channels persistent in the environment, thus making information easily available and reducing the need for group members to remember or retrieve it from overloaded channels, such as email.

Based on observations of practices for maintaining awareness and collaborating, we identified potential content for the system, limitations of current methods for sharing this information, and then sought to adapt this content to be effective and appropriate on a large shared display. In exploring the current methods the lab members used to maintain awareness of each other, we found they used information distributed across several methods and tools. In designing the system, we also aimed to provide users with a display on which they could have easy access to this type of information all at once.

The use of Semi-Public Displays supports a heightened awareness of group interests and activities by making information persistently visible in the environment. Additionally, the group viewable properties of the displays promote the use of the board not just as an information source, but as a catalyst for face-to-face collaboration and interaction as well.

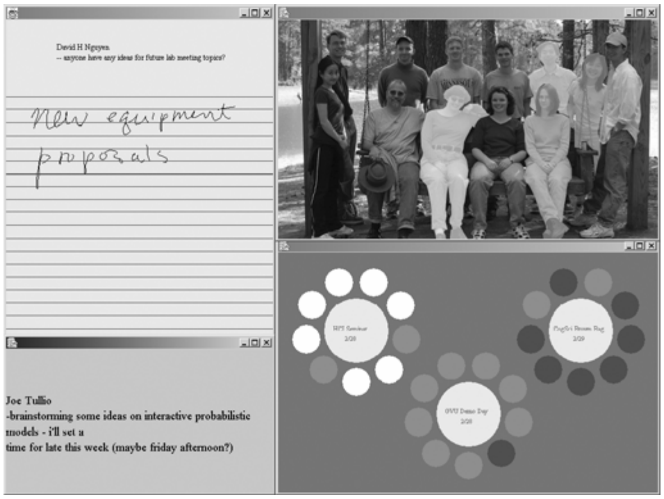

Figure 3-4. The Semi-Public Display prototype. Clockwise from top left: (a) Collaboration Space, (b) Active Portrait, (c) Attendance Panel, and (d) Reminders.

Our first design of a Semi-Public Display (Huang & Mynatt, 2003) aims to compensate for overloaded individual communication channels and to provide desired awareness information via four applications (Figure 3-4):

– Reminders: a slideshow-style application to display help requests and reminders that have been extracted from group members' weekly status reports, in order to provide constant access to the information in the environment.
– Collaboration Space: an interactive application that cycles through lab members' requests for assistance or brainstorming and allows lab members to create and edit ink annotations.
– Active Portrait: an abstract representation of group activity, consisting of a group photograph of the lab members in which each group member's image fades slowly over time when he or she leaves the lab.
– Attendance Panel: an application that provides lightweight group awareness through flower-like representations of upcoming events and the group members' intended attendance of those events.

6. RECURRING THEMES IN LARGE DISPLAY DESIGN FOR KNOWLEDGE WORK

Although our research has focused on understanding and supporting many varied information activities, we have identified some common cross-cutting themes in making large display applications effective for supporting all of these activities. Among these themes is context-awareness, which allows interfaces to provide richer and more useful information. Additionally, we have also sought to tailor information to be appropriate for its audience through the maintenance of information privacy and the emphasis on content relevance. Finally, the design of large displays need to match their mixed roles of primary and peripheral display and the need for informal and casual interaction.

6.1 Context-awareness

Several researchers have recommended integrating physical and virtual context to provide a better understanding of user activity in future pervasive computing environments (Dey, Abowd & Salber, 2001; Edwards, 1995). However, an overwhelming amount of the actual work in context-aware computing has focused exclusively on collecting and applying physical information—specifically location—in a variety of applications (Dey, et al.). Likewise, intelligent user interface community members have often used virtual context to tailor user interfaces and information presentation to match the user's activity or abilities (Maes, 1994), but rarely is physical context-awareness integrated into these systems. Some projects might be considered to have focused on integrating physical and virtual context to a limited extent (Maes, 1994; Tang, et al., 2001; Horvitz, Jacobs & Hovel, 1999) but none, to the best of our knowledge, extend beyond the bounds of a single application.

In Flatland, we explored the utility of context-awareness as the primary means for storing and retrieving whiteboard content. The notion of "un-filable" content typical of knowledge workers is strongly evident in typical whiteboard content. Requiring users to name, organize, and store this content runs against these flexible work practices. In Flatland, all material is automatically stored with a series of context tags, including virtual context information such as when each segment was created, when it was last edited, and what behaviors and colors were used. Additionally, if possible, physical context information is also stored when changes are made, such as what people were in the room at the time. Users can then search for past segments either by submitting context descriptions (e.g. a map created during the last month) or by searching backwards in time with a timeline scrollbar.

Table 3-1. The kinds of virtual- and physical-context information the Kimura system collects, the effect each has on the montages' appearance, and the specific user goals each supports.

Context description	Context type	Effect in task representation (montage)	Work practice supported
User interaction with windows on a desktop computer	Virtual	Presence, size, and opacity of window images	Multitasking, task awareness
Email messages	Virtual	Colleague availability notification cue	Collaboration
Documents printed by the user	Virtual	Peripheral notification cue	Task awareness
User interaction with Kimura on the electronic whiteboard	Virtual	Size and (initial) location of montage on electronic whiteboard	Multitasking, task awareness
Location (availability) of the user's colleagues	Physical	Colleague availability notification cue	Collaboration
Presence of multiple individuals in the user's office	Physical	Presence, size, and opacity of montages on electronic whiteboard	Collaboration
User's presence at or near peripheral devices	Physical	Peripheral notification cue	Task awareness
User's presence and physical activity at the computer (mouse and keyboard use, conductive paint, eye gaze tracking)	Physical	Rate of peripheral display change, intensity of alert notifications	Task awareness

The Kimura system is designed as a series of distributed components that fall into three classes: context acquisition, context interpretation, and user interaction. As the system acquires context information, it posts it to an activity database. Then, the context interpreter transforms the raw activity data into one or more working contexts and augments these working contexts with relevant cues about how other events in the office environment relate to them (see Table 3-1). Finally, user interaction components running on the whiteboard and the desktop computer display different visualizations of the working contexts. The user can manipulate (and, in the case of the whiteboard, annotate) those representations. The user interaction components on both the desktop and electronic whiteboard displays also act as an interface to a virtual window manager system.

Kimura's context acquisition components capture a wide variety of information garnered from sensors placed throughout the office and from virtual-context sources, such as keystroke and mouse monitoring utilities and mail and Internet use proxies. This integration of context information is an important contribution because the research in context-awareness focuses so heavily—and often, exclusively—on applying the use of physical-location data. A concerted effort is made to capture the user's activity while they run software applications, use documents and networked electronic information, and interact with peripheral devices distributed throughout the office environment.

The Kimura system uses the context information it collects in two ways. First, it creates a high-level framework of working contexts based on the user's activity, within which other virtual- and physical-context information is classified and interpreted. Second, it produces interactive visualizations of the user's working contexts. The system's context interpreter constantly updates the framework and the montage visualizations based on the stream of virtual and physical context captured by the context acquisition components.

The Semi-Public Displays leverage context in much the same way as the notification cues do in the Kimura system. Keystroke and mouse-movement monitors are installed on computers throughout the co-located group's work area to estimate the activity of each group member. This virtual context information is translated into a visualization supporting social awareness in the Active Portrait component. The use of keystroke and mouse-movement provides an inexact approximation of presence information, as it is conceivable that group members are present but not active at their computers. The Active Portrait application therefore provides viewers with approximate data about the presence of their colleagues, while the potential for ambiguity offers plausible deniability and privacy about the specifics of activity to individuals.

6.2 Privacy and Content Relevance

Large displays for groups can often suffer for lack of content that is relevant to its audience, that can be automatically gathered or easily authored, and that is appropriate given differing privacy concerns.

We have analyzed many of the existing and past projects that address awareness among people and have found that they can be meaningfully subdivided along two dimensions (Figure 3-5). First, awareness applications can be categorized by the type of audience they are intended to support. Applications have been built to facilitate awareness between pairs of people, as well as among groups of people. Applications that support groups typically target different types of groups such as a workgroup of five people, or a company of a hundred. Applications typically are intended to support only one of these audiences—pairs, small groups, or large groups.

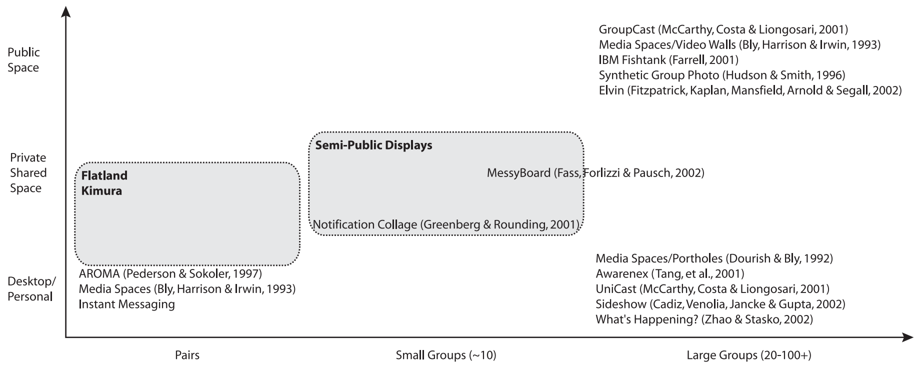

Figure 3-5. An illustration of the current design space of awareness applications categorized by the group size they are designed to support and the type of space in which they are meant to be viewed. Flatland and Kimura explore the space of shared public and personal displays to support very small groups, illustrated by the box on the left. Semi-Public Displays addresses the space of shared public displays to support small, co-located groups, delineated by the box in the center.

Additionally, we have examined the environments in which the applications are intended to be viewed and used. Applications typically reside on the desktop or other individually viewed display such as a PDA, within a shared private space such as an office or lab, or within a public space, such as the kitchen or lounge area of a building.

Looking at the space, we can see that much of the work on awareness applications falls into three areas: applications for pairs on desktop or individual displays, applications for large groups on desktop or individual displays, and applications for large groups in public spaces.

The success of these applications is primarily dependent on two issues: the ability of the application to provide *relevant content*, and the extent to which the application addresses *privacy* concerns.

Applications deployed in public spaces intended for supporting large groups that provide low levels of general awareness have been most effective. Displays that indicate the presence or absence of group members (Farrell, 2001) or video walls that allow for opportunistic conversations (Jancke, Venolia, Grudin, Cadiz & Gupta, 2001) aim to provide low-level or general information awareness. These types of applications have usually been successful in fostering low-level awareness because the lack of information specificity combined with the generality of audience to which the information is targeted mitigate some of the challenges of personal privacy and information relevance.

Applications intended for large groups that attempt to provide higher levels of awareness by providing more detailed information about individuals often face difficulties because of privacy concerns. It is difficult to make personal details available to large groups without impinging upon individuals' privacy. Because of the tradeoff between privacy and depth of awareness, the applications that have been most successful and useful for large groups have been those that provide less information. Another challenge for presenting detailed awareness information to large groups lies in the relevance of the information to its audience.

Because it is unlikely that detailed information will be of general interest and relevance to all members of a large group, rather than a subset of that group, it is difficult to provide appropriate content for these systems. A further challenge is presented when such systems rely on the users to supply content explicitly. Because the information an individual can provide about herself and her activities and interests is unlikely to be of interest to the large group in general, individuals tend not to be motivated to supply content, or else have difficulty identifying appropriate content. As a result, systems for promoting awareness among large groups that rely on user-submitted content tend to be uninformative because of a paucity of content, and because the content is not of interest to much of its audience (McCarthy, Costa & Liongosari, 2001; Tang, et al., 2001).

Desktop displays of applications for large group awareness exacerbate privacy concerns because it is difficult to assess who is looking at what information and how often. In contrast, public displays that exist in a shared social space that bring with them shared social mores and mechanisms for preventing abuse. Moreover, desktop interfaces for pairwise connections have the opposite effect. The privacy garnered by the personal display supports the depth and details appropriate for disclosure by a pair of users. Flatland and Kimura were designed to take advantage of this region of the design space, since their deployment was intended for individual offices – environments which exhibit characteristics of both desktop and private shared space displays, but have limited access, which is typically controlled by the display's owner.

In Kimura, we rely on the techniques underlying ambient and peripheral displays (e.g., the AmbientRoom, Ishii & Ullmer, 1997) to determine what content to display on the electronic whiteboard and how to display that content to mitigate privacy issues. Our montages consist of thumbnail-style images of the user's actual work, which have been manipulated to provide a quasi-summary of past activity. These montage styles are reminiscent of the "piles" (Mander, Salomon & Wong, 1992) that conveyed the age and type of items in desktop folders. However, due to the small size and limited resolution of the montages' component images, the content is largely undecipherable to an office guest who is unfamiliar with the details of the work that the montages represent. Like other ambient displays, the montages provide meaningful information to the initiated, but much less so to a casual observer.

Because of such privacy concerns, for example, it makes little sense to design for some unexplored areas of the audience-environment design space. If the intention of a system is to increase pairwise awareness between people, using a display in a public space is unlikely to be a good choice to support that goal.

However, not all areas with potential for design have been explored in depth. We posit that the space of applications for small groups supported by displays in shared group space has been largely unexplored. Designing applications for this space is an important area of research for several reasons. First, detailed information about individuals is more likely to be of relevance to the group as a whole. It is easier to identify and present content that is going to be both informative and appropriate to its audience. Additionally, whereas the presentation of personal information to a large group may intrude upon an individual's privacy, such information may be more appropriate among a group of co-located co-workers, who are likely to share context and have more personal knowledge of each other. Finally, because the need for communication and collaboration within small groups is often greater than that of large, distributed groups, it may be more important for individuals in a small group to have access to information about their co-workers. The small group audience may benefit more from having access to awareness applications and their content than the larger groups for which most applications have been designed.

The shared public display setting also matches the privacy and collaboration needs of small groups. The use of the display can be integrated into the practices and social customs of the group. The shared space encourages collaboration, especially semi-asynchronous collaboration, without competing for valuable desktop real estate.

The Semi-Public Displays application that we use to display an abstract representation of group activity consists of a group photograph of the lab members (Figure 3-4b). In this "active portrait," each person's image is displayed in full color if he or she is present in the lab. A person's image fades slowly over time when he or she leaves the lab. The resulting composite image provides viewers with a quick, at-a-glance picture of colleagues' recent presence in the lab. It allows the viewer to have some context about lab activity, especially when she has just entered the space. Unlike tools such as shared calendars or in-out boards, however, the image does not provide specific information about the exact times when a person left or entered the space, or a person's current whereabouts.

Our design provides low-fidelity presence information to give an overall picture of group presence, unlike instant messenger status, which gives high-fidelity information about presence, usually accurate to within minutes. Like instant messenger status cues, our design also monitors keyboard activity on desktop machines to measure idle time. However, rather than conveying information such as "Beth has been away from her computer for 7 minutes," the Active Portrait allows users to make inferences such as "The lab has been mostly empty for the weekend" or "Most people seem to be around this morning."

The application that we use to provide lightweight group awareness is the Attendance Panel (Figure 3-4c) on the Semi-Public Display, which uses a flower metaphor to provide an aesthetically attractive and easily interpretable abstract visualization to provide information about group members' upcoming plans. The panel displays several "flowers," which consist of a large circle with an event title as a label, surrounded by a ring of smaller circles. Each flower is a representation of an upcoming event, such as a seminar or a talk. The center circle represents the event description; the smaller circles, or "petals," represent users. Each petal has three states: blue for "not planning on attending," bright pink for "planning on attending," and white for "haven't decided yet." When a new event is added to the panel, it creates a flower whose petals are all white on a blue background, signifying that no one has yet updated his or her status. If a user elects not to attend the event, he toggles one of the petals to the "not attending" state, which is a slightly darker blue than the background, therefore blending with it. If a user chooses to attend the event, she toggles a petal to the "attending" state, which is a bright pink, contrasting significantly with the background. Users are free to select any petal that has not already been taken; they are not bound to any particular position on the flower. Their identities, therefore, cannot be discerned by the position of the petal on the flower, thus protecting their privacy.

The colors of the states create a visual image that brings the petals in the attending state to the foreground, while camouflaging the petals in the not attending state. A viewer can easily discern what events are of importance or interest to the group, or of potential relevance to her, by noticing how "complete" the flowers are. This simple interaction and visualization allows users to view planned attendance at near future activities, without compromising group members' privacy.

6.3 Display Design and Interaction Techniques

There has been substantial research in developing adequate interaction techniques for augmented whiteboards in conference room settings, including Tivoli (Moran, et al., 1996) and i-LAND (Streitz, et al., 1999). Some of the interaction techniques presented as a result of these projects – for example, the ability to “throw” content from one side of the board to the other in i-LAND – are clearly applicable for our projects, and many other large surface interactions. In general, our systems, like these conference room whiteboards, require informal interaction techniques to provide a fluid match between work practices and the interaction medium.

However, the unique work practices exhibited by knowledge workers and the specific environments that our projects are designed for lend themselves to additional kinds of informal interactions not well addressed by previous work. The whiteboard surfaces often serve as persistent displays of information in the work environment and thus incur additional requirements regarding information structure, readability, and privacy. Also, because these whiteboards often function as peripheral displays, various visualization techniques are needed so that the interface is appropriately obtrusive.

Our explorations into informal interaction techniques began in Flatland, where we designed and implemented a few behaviors to support typical office whiteboard tasks. Flatland’s design goals of simple, informal interaction extend past the general look-and-feel of the interface into the design of individual behaviors themselves. Since the purpose of the behaviors is to support informal, preproduction tasks, we strongly favored ease of use over providing features for producing a detailed artifact. Common themes in designing individual behaviors were:

– Few explicit commands exist; strokes are interpreted on-the-fly.
– Generated output is rendered in a “hand-drawn” style.
– Minimal (if any) control widgets are added to the segment.
– “Infinite” undo-redo supports easy error recovery.
– Handwriting recognition is generally not used, to limit the need for error correction and recovery. The one exception was the calculator behavior, which favors handwriting in lieu of push buttons.

This final design choice limits some potential uses of the system, but significantly simplifies user interaction.

The difference between behaviors and traditional applications becomes more apparent when combining multiple behaviors over time. For example, starting first with the map behavior, a user can sketch out relevant streets and intersections. After removing the map behavior and applying the 2D drawing behavior, the user can sketch relevant buildings and other landmarks. Finally, with no behaviors present, the user can label the map.

Kimura was designed to compliment Flatland's interface. Each of the montages displayed on the electronic whiteboard is essentially a Flatland segment with a special rendering behavior that creates a visualization of one working context on the desktop computer. Montages can be manipulated in many obvious ways on the wall display: they can be moved, deleted and so on using simple gestures. Montages also respond in many of the same ways as hand-drawn segments do in Flatland. For example, montages adjust their size to fit their contents, particularly user annotations, and they "push" neighboring montages when moved to prevent segments on the wall display from overlapping. Currently, the behaviors connected to montages are *moving* when selected, and *annotating* when de-selected.

We present two primary day-to-day uses for Kimura's large interactive display surface. First, and most importantly, users can use the wall display as a peripheral interface for keeping track of the existence of, and changes in, background activities. Second, users directly manipulate the montage images, in conjunction with other whiteboard tasks, while standing at the wall display.

Tapping a selected montage triggers a task switch, an operation that can also be performed from the desktop computer. The contents of the past activity disappear from the desktop and reappear as a montage on the wall display. Simultaneously, the contents of the new task appear on the desktop. The montage for the current task is also displayed near the desktop monitors. This "near-periphery" display allows the user to remain aware of contextual cues, such as a past browsing activity, that are no longer part of the active desktop, as well as the annotations that the user may have made to the montage on the electronic whiteboard. Montages retain their position on the wall display so that a background task will return to its prior location unless the user explicitly rearranges the montages.

Annotating montages is an example of an interaction that is particularly well-suited to the wall display: using the dry pens of various colors provided with the SMART Boards, the user may annotate montages with simple electronic ink.

In keeping with the Flatland design guidelines, and in order to suggest that the information on the wall display is peripheral to the user's active working context, we have tried to create an informal look for the montages. Montages are shown with sketchy backgrounds in soft colors using a separate color for each montage (see Figure 3-3).

The Semi-Public Display project also required the development of informal interaction techniques, although the goals for these techniques differ somewhat from those underlying the Flatland and Kimura projects. Because the Semi-Public Display is designed for a shared, group setting, its interactions must be lightweight enough to encourage casual, walk-up use, and be comprised of observable, easy-to-learn actions so that a novice user is not intimidated by a steep learning curve.

By creating a Collaboration Space component on the Semi-Public Display, our aim was to provide a dynamic, captured space that was clearly designated as editable and viewable by anyone at anytime (Figure 3-4a). Using help requests extracted from the group members' weekly status reports, the Collaboration Space cycles through each request, displays each one for several minutes at a time, and provides a space in which users can create and edit ideas. Users add content using freeform ink and a variety of common gestures provided by the SATIN toolkit (Hong & Landay, 2000). The annotations affiliated with each item are captured and redisplayed every time the item is shown. Unlike conventional whiteboards, whose content group members are often wary of erasing, editing, or adding to outside of brainstorming sessions, our design keeps the request visible and provides an explicit space for brainstorming and scribbling at all times.

Similarly, the Attendance Panel allows users to view other group members' intended participation in an upcoming event, and to add their own preferences with a minimum amount of additional effort. Because the center of each "flower" and its petals are displayed on a touch-sensitive surface, users do not need to interact with the board using additional technology or specialized implements – they simply touch the portion of the board with which they want to interact. Furthermore, the fact that individual users are not permanently associated with any particular "petal" means that users can simply pick any uncolored petal to encode with their planned activity attendance. Based on our informal observations, a user could approach the board and provide their attendance intentions for three to four upcoming activities in a matter of seconds. This low-effort interaction helped to foster use of the Attendance Panel, which increased the informational value of the Semi-Public Display as a group resource.

7. CONCLUSION

In this chapter, we have provided glimpses of the future for large interactive displays and, in particular, their utility for knowledge workers.

The use of large interactive displays follows a similar historical path as the use of previous computing technologies. Due to their expense, the deployment of early large interactive displays was limited to formal situations and required the use of specialized applications and expert system operators. In the past decade, these displays have seen more widespread deployment in specific environments – most commonly, meeting rooms – and, while typically controlled by a technical expert, offer novice users the opportunity to explore the use of these displays with general-purpose software. We are now at the point where these large interactive displays are becoming commonplace and will materialize as a general resource for individuals and informal groups.

This shift to commonplace use requires new software infrastructures and user interface techniques to make large interactive displays a useful office tool. This need is particularly evident when we investigate the work practices of knowledge workers. When large interactive displays become a fixture in these workers' office environments, they must complement the informal, flexible and improvisational practices of knowledge workers, in contrast to forcing the workers to shape their practices around the use of a specialized tool. One distinction is the shift to providing persistent information on large displays, in contrast to their common use as a typically-blank display harnessed for focused interactions. By using these displays as a persistent source of information, the content on the display becomes an additional resource for manipulating information. Additionally, the interaction with these displays needs to exhibit informality to meet the flexibility of knowledge work and mirror the transient nature of information more likely to be modified than filed. Finally, this resource must become part of a distributed network of information sources. It can no longer serve as a stand-alone technology. By leveraging a context-aware infrastructure, the content displayed on the board can provide additional benefit by lowering the burden of retrieving information and supporting awareness.

Similar benefits accrue to groups of knowledge workers. As a persistent source of information, large interactive displays are an excellent resource for coordination, information sharing and awareness. However, these displays face a variety of challenges in providing benefit to a group while avoiding pitfalls of requiring too much effort to use or breaching privacy concerns.

ACKNOWLEDGEMENTS

This work was completed under NSF IIS grant 99-88712, and with additional support from IBM Research, the Sloan Foundation and Ricoh Innovations. We would like to thank the individuals who participated in our numerous studies, and Dr. Gregory Corso and the members of the Everyday Computing lab for their patience and feedback. Thank you to Jaroslav Tyman for his valuable assistance in the development of the prototypes. We would also like to recognize the valuable comments provided by the organizers and participants of the Large and Situated Displays workshop at CSCW 2002. Thanks to Amy Voida and the book editors for reviews of earlier versions of this chapter.

REFERENCES

Abowd, G.D. (1999, October). Classroom 2000: An experiment with the instrumentation of a living educational environment. *IBM Systems Journal, 38*(4), 508–530.

Bly, S.A., Harrison, S.R., & Irwin, S. (1993, January). Media spaces: Bringing people together in a video, audio, and computing environment. *Communications of the ACM, 36*(1), 28–46.

Cadiz, J.J., Venolia, G., Jancke, G., & Gupta, A. (2002). Designing and deploying an information awareness interface. *Proceedings of the 2002 ACM conference on Computer supported collaborative work* (pp. 314–323). New York: ACM Press.

Dey, A.K, Abowd, G.D., & Salber, D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction (HCI) Journal, 16*(2–4), 97–166.

Dourish, P., & Bly, S. (1992). Portholes: Supporting awareness in a distributed work group. *Proceedings of the SIGCHI conference on Human factors in computing systems* (541–547). New York: ACM Press.

Drucker, P.F. (1973). *Management tasks, responsibilities, and practices.* New York: Harper & Row.

Edwards, W.K. (1995). *Coordination infrastructure in collaborative systems.* Unpublished doctoral dissertation, Georgia Institute of Technology, Atlanta.

Farrell, S. (2001). Social and informational proxies in a fishtank. *CHI '01 extended abstracts on Human factors in computing systems* (pp. 365–366). New York: ACM Press.

Fass, A.M., Forlizzi, J.. & Pausch, R. (2002). MessyDesk and MessyBoard: Two designs inspired by the goal of improving human memory. *Proceedings of the conference on Designing interactive systems: processes, practices, methods, and techniques* (pp. 303–311). New York: ACM Press.

Fitzpatrick, G., Kaplan, S. Mansfield, T., Arnold D., & Segall B. (2002). Supporting public availability and accessibility with Elvin: Experiences and reflections. *Computer Supported Cooperative Work (CSCW): The Journal of Collaborative Computing, 11*(3-4), 447–474.

Greenberg, S., & Rounding, M. (2001). The notification collage: Posting information to public and personal displays. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 514–521). New York: ACM Press.

Henderson, J.D.A., & Card, S.K. (1986, July). Rooms: The use of multiple virtual workspaces to reduce space contention in window-based graphical user interfaces. *ACM Transactions on Graphics, 5*(3), 211–241.

Hong, J.I., & Landay, J.A. (2000). SATIN: A toolkit for informal ink-based applications. *Proceedings of the 13th annual ACM symposium on User interface software and technology* (pp. 63–72). New York: ACM Press.

Horvitz, E., Jacobs, A., & Hovel, D. (1999). Attention-sensitive alerting. *Proceedings of UAI '99: Conference on Uncertainty and Artificial Intelligence* (pp. 305–313). San Francisco, Morgan Kaufmann.

Huang, E.M., & Mynatt, E.D. (2003). Semi-public displays for small, co-located groups. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 49–56). New York: ACM Press.

Hudson, S.E., & Smith, I. (1996). Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (pp. 248–257). New York: ACM Press.

Igrashi, T., Edwards, W.K., LaMarca, A., & Mynatt, E.D. (2000). An architecture for pen-based interaction on electronic whiteboards. *Proceedings of the Working Conference on Advanced Visual Interfaces* (pp. 68–75). New York: ACM Press.

Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits, and atoms. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 234–241). New York: ACM Press.

Jancke, G., Venolia, G., Grudin, J., Cadiz, J., Gupta, A. (2001). Linking public spaces: Technical and social issues. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 530–537). New York: ACM Press.

Kidd, A. (1994). The marks are on the knowledge worker. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 186–191). New York: ACM Press.

MacIntyre, B., Mynatt, E.D., Voida, S., Hansen, K.M., Tullio, J., & Corso, G.M. (2001). Support for multitasking and background awareness using interactive peripheral displays. *Proceedings of the 14th annual ACM symposium on User interface software and technology* (pp. 41–50). New York: ACM Press.

Maes, P. (1994, July). Agents that reduce work and information overload. *Communications of the ACM, 37*(7), 31–40.

Malone, T.W. (1983, January). How do people organize their desks? Implications for the design of office information systems. *ACM Transactions on Office Information Systems, 1*(1), 99–112.

Mander, R., Salomon, G., & Wong, Y.Y. (1992). A ‘pile’ metaphor for supporting casual organization of information. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 627–634). New York: ACM Press.

McCarthy, J., Costa, T., Liongosari, E. (2001). UniCast, OutCast & GroupCast: Three steps towards ubiquitous, peripheral displays. *Proceedings of UbiComp 2001: Ubiquitous Computing* (pp. 332–345). Berlin: Springer-Verlag.

Moran, T.P., Chiu, P., Harrison, S., Kurtenbach, G., Minneman, S., & van Melle, W. (1996). Evolutionary engagement in an ongoing collaborative work process: A case study. *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (pp. 150–159). New York: ACM Press.

Mynatt, E.D. (1999). The writing on the wall. *Proceedings of INTERACT ’99* (pp. 196–204). Amsterdam, IOS Press.

Mynatt, E.D., Igrashi, T., Edwards, W.K., & LaMarca, A. (1999). Flatland: New dimensions in office whiteboards. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 346–353). New York: ACM Press.

Mynatt, E.D., Igrashi, T., Edwards, W.K., & LaMarca, A. (2000, July–August). Designing an augmented writing surface. *IEEE Computer Graphics and Applications, 20*(4), 55–61.

Pedersen, E.R., & Sokoler, T. (1997). AROMA: Abstract representation of presence supporting mutual awareness. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 51–58). New York: ACM Press.

Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., & Steinmetz, R. (1999). i-LAND: An interactive landscape for creativity and innovation. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 120–127). New York: ACM Press.

Tang, J.C., Yankelovich, N., Begole, J., Van Kleek, M., Li, F., & Bhalodia, J. (2001). ConNexus to Awarenex: Extending awareness to mobile users. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 221–228). New York: ACM Press.

Voida, S., Mynatt, E.D., MacIntyre, B., & Corso, G.M. (2002, July–September). Integrating virtual and physical context to support knowledge workers. *IEEE Pervasive Computing, 1*(3), 73–79.

Zhao, Q.A., & Stasko, J. T. (2002). What's happening?: Promoting community awareness through opportunistic, peripheral interfaces. *Proceedings of the Working Conference on Advanced Visual Interfaces* (pp. 69–74). New York: ACM Press.