IEEE Personal Communications and IEEE Pervasive Computing
Special Issue on Context-Aware Pervasive Computing

Integrating Virtual and Physical Context to Support Knowledge Workers

Stephen Voida, Elizabeth D. Mynatt, Blair MacIntyre, and Gregory M. Corso

Version #2, Submitted 28 June 2002

INTRODUCTION

Our research seeks to explore how an "office of the future" might be designed to better support knowledge workers – business professionals whose job it is to interpret and transform information [1]. Previous studies have shown that the work practices of knowledge workers involve managing multiple tasks, collaborating effectively among several colleagues and/or clients, and manipulating the information most relevant to their current task by leveraging spatial organization of their work area [2, 3, 4, 5]. The diversity of these practices and the complexity of implementing flexible computing tools have made it difficult to meet all the needs of these workers.

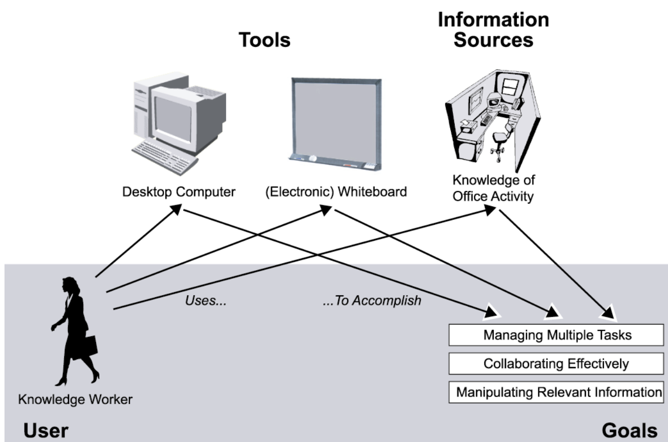

Figure 1: Knowledge workers must rely on a variety of independent tools and information sources to successfully accomplish the goals associated with their work activities.

In typical office environments, knowledge workers accomplish their work activities using a variety of tools and information sources [Figure 1]. Common tools include desktop or laptop computers, whiteboards, and desks. Information sources are frequently distributed throughout the office environment, both within and beyond the walls of the individual office. Although each of these tools and information sources contribute to a comprehensive understanding of the various activities in progress, it is entirely up to the worker to make sense of the information contained in each of the tools and distributed throughout the office, to synthesize the disparate pieces of information, and to act on that information in an appropriate manner.

We are working to augment and integrate these independent tools into a pervasive computing system that monitors a user's interactions with the computer, an electronic whiteboard, and a variety of networked peripheral devices and data sources. Our system also draws from several physical sensors distributed throughout the office. This combination of virtual and physical context drives the creation of activity representations on a wall-size peripheral display (the electronic whiteboard) [Figure 2]. This allows the user to monitor each ongoing work activity, transition smoothly between them, access a wide variety of contextual information designed to facilitate collaboration, and maintain awareness about relevant changes in the physical and virtual aspects of each work activity. Additionally, the interactivity provided by displaying the activity representations on an electronic whiteboard allows the user to informally annotate and spatially organize these representations.

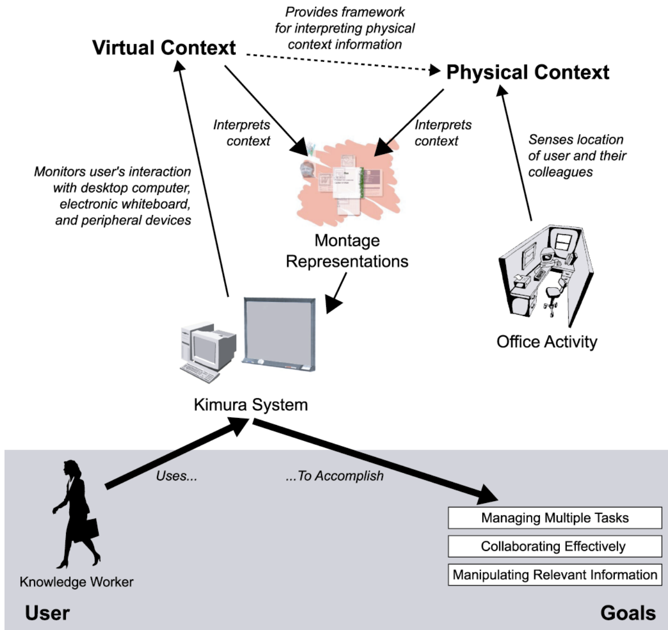

Figure 2: Our vision of a pervasive computing system leverages both physical and virtual context to reduce the number of tools and information sources that knowledge workers must utilize and synthesize to accomplish their goals.

Several researchers have recommended integrating physical and virtual context to provide a better understanding of user activity in future pervasive computing environments [6, 7]. 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 (e.g. [6]). Likewise, members of the intelligent user interface community have often used virtual context to tailor user interfaces and information presentation to match the user’s activity or abilities (e.g. [8]), but rarely is physical context-awareness integrated into these systems. Few research projects, if any, have focused on demonstrating an integration of physical and virtual context. Some projects might be considered to have accomplished this to a limited extent [8, 10], however none, to the best of our knowledge, approach the scale of an integrated office environment.

Previously, we presented the design and architecture of the Kimura system [11]. In this paper, we give an overview of the Kimura system with attention focused on the ways in which Kimura integrates physical and virtual context information to create a pervasive computing environment for knowledge workers. We first describe the basic workings of the Kimura system and provide a scenario of its use, calling attention to our use of "working contexts" to aid in collecting and synthesizing context information about the user's work activities and to the visualization of context information on an interactive whiteboard surface. We then discuss in greater detail the types of virtual and physical context used by Kimura and how this information supports knowledge work practices. We note two key contributions. First, we use virtual context from the user's desktop actions to aid in classifying, interpreting, and visualizing other forms of virtual and physical context. Second, we integrate virtual and physical context information into visualizations of the user's disparate activities to aid the user in interpreting and acting upon this available information.

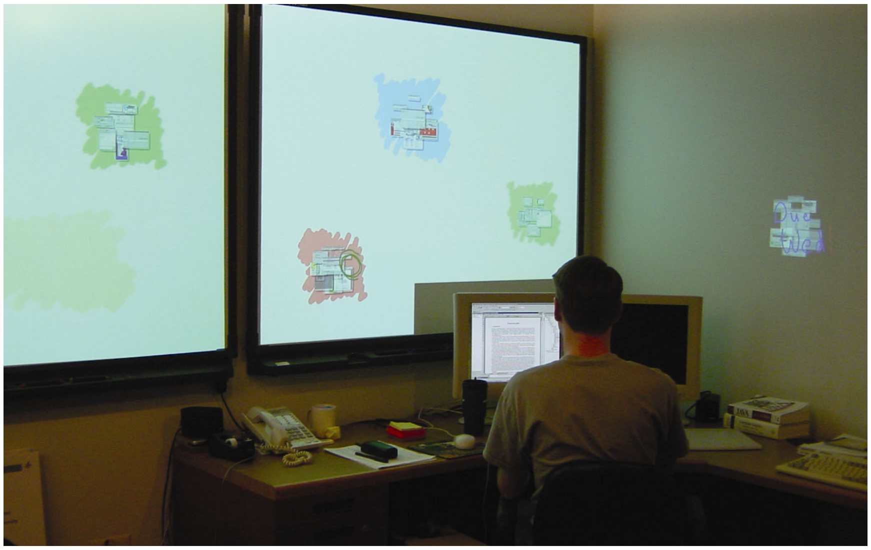

Figure 3: The Kimura system in an office environment, including the focal and peripheral displays.

The Kimura System

Kimura separates the user's "desktop" into two regions: the focal display on the desktop monitor, enhanced with a virtual desktop manager, and peripheral displays that we project on the office walls. Each of the user's work activities is associated with a unique virtual desktop on the focal display. While the user is interacting with a single work activity, representations of his or her background activities are illustrated as visual montages and projected on the peripheral display [Figure 3].

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 on-going interactions with people and objects related to that activity. We refer to this cluster as the activity's *working context*. Each working context can have numerous documents, including text files, web pages, and other application files, that have been used in the course of the activity, plus indications of ongoing activity such as e-mail 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 [12], users demarcate the boundaries of working contexts manually – in our case, by creating and moving between virtual desktops. We chose this strategy because these operations are easy to perform from the user's perspective, can be easily monitored to detect working context changes, and because attempting to infer transitions between work activities will be error-prone if left to the system.

Each background activity (working context) is visualized as a *montage* of images garnered from system activity logs [Figure 4]. 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 current windows in each room, our goal is to provide visualizations of past activity in context. These visualizations help remind the user of past actions; the arrangement and transparency of the component images automatically creates an "icon" for the working context. Additionally, montages can serve as anchors for background awareness information that is gleaned from our context-aware infrastructure.

We elected to use the electronic whiteboard as the primary display surface for the montage visualizations in order to support common whiteboard practices [4]. Our implementation of the system incorporates existing electronic whiteboard interaction techniques [13, 14] with montages and notification cues.

Montages can be repositioned by the user to indicate the respective priority of background activities and can be annotated with informal reminders. Whiteboards feature an intuitive user interface, and are well suited to supporting these kinds of informal information management activities. Additionally, the whiteboard's large display area is an ideal, unobtrusive location to display contextually relevant information about the user's work activities and context information sensed from around the office.

Figure 4: A montage of a working context, including a number of application windows and two external context notification cues, representing both virtual (a completed print job) and physical (availability of a colleague) context information.

Our montage designs take advantage of a number of visualization techniques to express the semantics of the working contexts. In order to provide a quasi-summary of past activity, literal representations of the user's activity in a working context (document thumbnail images) are manipulated to convey the history of the working context at a glance. For example, the relative size of the document thumbnail images are determined by the amount of time that the user has spent interacting with each desktop application; thumbnail images representing applications that consumed more of the user's time are rendered larger than other thumbnail images. Another example is our use of transparency. The recency of a user's actions is encoded in the transparency of the thumbnail images, so that the most recently used documents are the most opaque. We are also experimenting with a variety of thumbnail layouts within the montage, with the hypothesis that changing the layout over time as a working context "ages" will foster better activity resumption.

Our development of the Kimura system prototype is only one aspect of a multidisciplinary research project. We are also running a series of psychology experiments designed to reveal the thresholds at which humans can perceive changes in their peripheral vision. The results of these experiments will help inform the development of future montage animations, so that the user immediately notices critical changes and notifications, while less important changes do not unnecessarily distract users from their focal activity. In addition, we are performing user tests on the effectiveness of our montage designs in supporting suspended task resumption, and are completing our own task analysis of knowledge work in the office environment to supplement and clarify the findings in the published literature.

Scenario

As our knowledge worker Wendy walks into her office Monday morning following a week's vacation, she scans the piles of paper on her desk and the content on her whiteboard as she recalls the work that has been waiting for her.

After quickly surveying the various whiteboard montages that represent on-going activities, she annotates the budget plan with "Work on Wed., Due Friday" and throws it to the other side of the whiteboard. The calendar image in the Acme design project montage reminds her of the design briefing later that day. She studies the montage for a moment, using the thumbnail images of the documents that she had open most recently as well as the more ghostly images of past documents to remember how far she had gotten in *the design briefing activity before she left for vacation. She taps on that montage to load it onto her desktop.*

The design briefing documents reappear on her desktop computer, just as she left them. After a quick perusal, she resumes her Web searches for details on an interesting technology and fine-tunes one of her sketches. After sending the new sketch to the printer, she decides to spend some time catching up on the theme ideas for the upcoming open house. Using the desktop controls to switch activities (and virtual desktops), the montage for the Acme design activity reappears on her whiteboard annotated with a printer icon, indicating that the print job is in progress.

As Wendy contemplates her reply to an interesting theme idea from one of her colleagues, she notices that his face has appeared on her whiteboard. Ah, Joe must be in the coffee room. Deciding that a face-to-face discussion would be more useful than posting another message, she goes to join Joe for coffee and brainstorming.

Later that day she decides to go ahead and start working on those budget numbers. From the corner of her eye she notices the softly changing calendar in the Acme design montage. It is time for the meeting. This jet lag is killing her internal clock. As she runs out of the office, she sees the icon for the completed print job. The print out! Grateful that someone, or something, is on top of things, she heads to the printer on the way to the meeting.

CONTEXT-AWARENESS IN THE AUGMENTED OFFICE

The Kimura system is designed as a series of distributed components that fall into three classes: context acquisition, context interpretation, and user interaction. As context information is collected, it is posted into an activity log database. It is then transformed by the context interpreter into high-level representations of the user's work activities – working contexts – and augmented with salient cues about how other events in the office environment relate to each working context [Table 1]. Finally, user interaction components run on both the electronic whiteboard and on the desktop computer, displaying different visualizations of the working contexts generated by the context interpreter, and allowing the user to manipulate, and in the case of the electronic whiteboard, annotate those representations. The UI components on both the desktop and electronic whiteboard displays also act as an interface to a virtual window manager system. A detailed description of the system's architecture and implementation is available in [11].

The context acquisition components capture a wide variety of information garnered both from sensors placed throughout the office environment and from virtual context sources, such as keystroke and mouse monitoring utilities and mail- and Internet-use proxies. We consider this integration of context information an important contribution of our system, since much of the research in context-awareness focuses so heavily – and often, exclusively – on applying the use of physical location data. In order to support knowledge workers whose primary job function is to interpret and transform information, the ability of our system to reference the virtual activity of the user is critical. Knowing that the user is sitting behind their desk and using their computer provides very little detailed information about the nature of the user's work activity. We make a concerted effort to capture the activity of Kimura's user while they run software applications, use documents and networked electronic information, and interact with peripheral devices distributed throughout the office environment, in addition to more traditional context-awareness information, such as the location and identity of their colleagues.

Collecting virtual context

Our system uses several lightweight monitoring components and proxies to acquire virtual context information about the Kimura user. Currently, our focus is on capturing the user's actions on their desktop computer and their interactions with the windows comprising each work activity. We have developed a desktop monitoring system for Microsoft Windows™ using the hooks feature exposed through the Win32 API. When the Kimura system is running, Windows sends notification of each relevant user action (i.e. opening a window, changing the window focus, pressing a key, clicking the mouse) to our desktop monitoring process, which packages the event into a Kimura-readable activity log entry and sends it to a distributed activity log. Additionally, the desktop monitor creates a screenshot of the content of each window each time the window system's input focus changes. These screenshots are made available to Kimura's context interpreter so that visual representations of the user's task can include actual images of the user's work. The images, akin to thumbnails, provide more salient visual reminders than generic icons or labels.

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

Table 1: A wide variety of virtual and physical context information is collected by the Kimura system and reflected visually in the montage representations to support a number of specific user goals.

This virtual context collection mechanism does not currently attempt to connect directly with the running applications to find out more detailed (and potentially useful) information such as the open document's file name or the contents of a selected region. Although we have considered this approach, we decided that an initial prototype of the Kimura system should not be bound to a small set of Kimura-aware applications, as this would limit the practicality of its deployment and user testing in the field. Even without this close integration between the Kimura system and the applications running on the user's desktop, we are able to collect a significant amount of information about the user's activity. We use metrics such as the amount of time a particular window has been in focus and the number of focus switches between open windows on the desktop to produce reasonably detailed visualizations of the user's overall activity within a given working context.

Kimura also acquires virtual context through an e-mail monitoring system, which helps to track the user's interaction with colleagues during a particular work activity. A small process running on the user's mail server monitors changes in each of the user's mailboxes. Currently, we monitor all e-mail messages that the user sends – the system associates each mail recipient with the active working context. We also add the recipient to a list of individuals with whom the Kimura user may be trying to connect, and instruct the location-monitoring component to actively monitor the availability of that individual by watching for their presence in public areas of the office.

We are also interested in the user's interactions with distributed peripheral devices over the course of a work activity. We have implemented a printer proxy that records the ID and status of pending print jobs within a working context. The proxy is used to capture the frequency with which the user interacts with a given peripheral over the course of a work activity. It also tracks the status of individual print jobs. As the status of each print job changes (i.e. a print job is sent to the spooler, a print job buried in a long queue finally prints, or a print job stalls because the printer is out of paper), a notification is generated by the context interpreter and presented to the user.

Finally, we monitor all of the user's interactions with the Kimura system itself, in order to discover which work tasks will be of most interest to the user given their current working context, and to build a model of how tasks are related over the course of the user's work. This information is used by the electronic whiteboard to draw the user's attention to relevant montages and to influence the initial placement of new montages on the peripheral display.

Collecting physical context

Kimura also makes use of physical context in order to assist the user in reconstructing the environmental circumstances surrounding a working context and to provide cues about the location and availability of the user's colleagues. In our current prototype system, we simulate a pervasive, location-aware infrastructure with a series of Dallas Semiconductor i-Button™ docks distributed throughout the office environment. We anticipate integrating the Kimura system into a much more robust network of context sensing devices through an infrastructure such as the Context Toolkit [6] as the sensors and infrastructure become more widely available at our location.

Currently, our sensor network is designed to generate events describing the arrival and departure of known individuals in our augmented office environment, in public areas of the office, and near peripheral devices (i.e. next to the printer). Although the granularity of this information is somewhat coarse, it provides enough detail for the system to determine the general whereabouts and activity of the user when they are not directly interacting with their desktop or electronic whiteboard. It also allows the system to determine the general location of colleagues and infer when they may be available for collaboration, or when they have joined the user in the augmented office for an informal meeting.

We also record the physical presence and activity of the user at their office workstation. This is currently accomplished with a keystroke and mouse movement monitor, but we hope to extend it to other areas around the user's workstation (for example, by using conductive paint on various surfaces or taking advantage of a vision-based tracking system) so that we can estimate the user's focus on a given task. We plan to use this information to adjust the rate of visual changes on the electronic whiteboard and to predict the user's availability for interruptions and notifications. We have begun experimenting with an off-the-shelf, vision-based gaze tracking system, but have not yet successfully integrated it into the Kimura system.

Making use of context-awareness

The Kimura system makes use of the context information it collects in two ways. First, the system uses virtual context information to create a high-level representation of the user's activity, within which other virtual and physical context information is classified and interpreted in a meaningful way. Second, the Kimura system uses physical and virtual context information to produce interactive visualizations of the user's activities. These visualizations include a montage of the windows with which the user has interacted on the desktop computer, as well as event notification cues (e.g. a change in the status of a print job or the arrival of a colleague) associated with each working context. The system's context interpreter constantly updates the high-level representation and the montage visualizations based on the stream of virtual and physical context generated by the context acquisition components.

Classification and Interpretation
In the current implementation of the system, we use some of the virtual context collected by the system to guide the classification and interpretation of the rest of the virtual and physical context. The system's context interpreter utilizes the virtual context of the user's interaction with the windows, keyboard, and mouse on the desktop computer, as well as the activity demarcations that the user explicitly provides, to construct a framework of working contexts. As these working contexts are created, each is associated with a unique virtual workspace on the desktop computer. Additionally, as the user moves between working contexts, the context interpreter and various context acquisition components in the system are notified, insuring that incoming physical and virtual context information is interpreted with respect to the proper working context.

We are also beginning to explore ways in which the physical context information gathered by the system can be used to guide the interpretation of other context-awareness information. For example, sensed information about the user's location (i.e. at the computer, at the printer, in a public space, or working with a colleague) can help determine which of the working contexts the user is engaged in. This capability would allow the system to continue updating the user's working contexts with incoming context information, even without the explicit interaction on the desktop computer currently used to determine the user's activity and active working context. A simple example currently implemented in our system is that the context interpreter has the ability to remove completed print job notifications from all working contexts when the user is sensed in the vicinity of the network printer. A more complex example is demonstrated when a colleague is sensed to have entered the augmented office environment – when an individual besides the office's primary user enters the office, the system increases the relative importance metric of working contexts associated with that individual, subtly changing the high-level framework of working contexts. As a result, the montages on the electronic whiteboard that are associated with the visitor become more prominent and easily accessible in anticipation of an informal meeting about one of those working contexts.

Visualization
Only some of the virtual and physical context collected by the system is used to develop the high-level working context representations. The rest of the context information is aggregated and used to generate the montage visualizations, which are then projected onto the electronic whiteboard.

Each working context is visualized as a montage of window thumbnails and notification cues. The basic content of each montage, a cluster of document thumbnails, is derived from the high-level representation of one of the user's working contexts. However, the placement and size of the montage, the layout of the thumbnail cluster, the appearance characteristics of each of the document thumbnails, and the presence and type of notification cues in the montage are all computed based on a wide variety of context information related to the montage's working context. Additionally, the montage also displays historical information about a working context in order to provide a complete picture of the user's activity over the entire course of a working context. These historical artifacts might include thumbnails representing documents that were once significant, but have since been closed.

For example, a montage representing a budget activity that was used very recently might contain several thumbnails of open spreadsheet and word processing documents, a thumbnail of an email message from a financial consultant, a semi-transparent thumbnail of a bank web site that was referenced earlier and closed, and several nearly-transparent thumbnails of background applications always running on the user's machine, such as an instant messaging client and a digital music player. Because the working context was used recently, the thumbnails would be arranged to resemble the layout of the original documents on the computer screen. The montage might also include notification icons representing a pending print job and the arrival of the company CFO in the conference center lobby.

CONCLUSIONS AND FUTURE WORK

Our system, Kimura, leverages both virtual and physical context-awareness to provide a seamless pervasive computing experience for knowledge workers in an office environment. We have integrated the various tools of the knowledge worker – the desktop computer and an electronic whiteboard – and information distributed throughout the rest of the office environment, and simplified the process through which knowledge workers accomplish goals such as multitasking, collaborating, and manipulating relevant information. We have used virtual context from the user's desktop to provide a framework within which a variety of context information can be classified, interpreted, and visualized. We have also used both virtual and physical context information to generate semantically meaningful visualizations of knowledge workers' activity.

With these visualizations, our goal was to create an environment that provides context-aware summaries of the user's work activities on a single, integrated display. The montage design relieves the user of the burdens associated with maintaining a large amount of information – information about their work activities and each activity's related contextual information – and with synthesizing that information on the fly from a potentially overwhelming number of sources. The montages are designed to present this information without intruding upon the user's focal activity and in a manner that supports the needs of knowledge workers.

We are preparing to deploy several Kimura systems over the next several months in order to evaluate the usability and usefulness of our design based on real-world use of the system. We hope to combine our findings with those from the ongoing psychology experiments and montage user tests to inform the designs of future system prototypes.

An important area of future work is to enhance our virtual context acquisition component so that it can extract more detailed information from common desktop applications (e.g. Netscape Navigator™, Microsoft Office™, etc.). We chose to implement an application-independent system so that we could produce a somewhat realistic prototype environment without limiting its use to specially written, Kimura-aware applications. However, our current approach limits the system to monitoring applications currently running in memory on the desktop computer, and provides an abstract view of the user's desktop activity through their interactions with the desktop computer at the windowing system level. We anticipate that by implementing a tighter coupling with key desktop applications, we can improve our interpretations of user activity, allow more interesting interactions with the system, and expand the visualizations to include recommendations for related documents or external reference materials.

We are also interested in connecting the Kimura system with a robust, context-aware infrastructure, such as the Context Toolkit [6], and replacing our network of i-Button docks with a variety of passive context sensors. We will gain a much more ecologically valid sense of how the system supports the work practices of knowledge workers when we eliminate the artificiality and overhead of our physical context sensing approach. Leveraging a more mature context-aware infrastructure also allows us to experiment with a more diverse array of sensing devices.

Finally, we are interested in exploring the computer-supported cooperative work implications of linking two or more instances of the Kimura system together. We envision an office environment in which multiple work colleagues each have their own Kimura system. We are interested in understanding the interaction issues raised as colleagues collaborate over working contexts. These issues may include topics such as privacy, information access and visibility, and synchronization of shared information. We are also interested in exploring ways to adapt our existing montage designs to visualize the concurrent activity of different people within the same working context. Finally, we hope to understand how the electronic whiteboard can be used more effectively as an informal meeting and collaboration space.

ACKNOWLEDGEMENTS

This work has been supported by the NSF under grant 9988712, as well as by research and equipment grants from IBM and Ricoh Innovations. We would also like to thank our colleagues and students for their feedback and ideas on this paper.

REFERENCES

1. P.F. Drucker, *Management Tasks, Responsibilities and Practices*, Harper & Row, New York, 1973.

2. A. Kidd, “The Marks are on the Knowledge Worker,” *Proc. of CHI’94: Human Factors in Computing Systems*, ACM Press, New York, 1994, pp. 186-191.
3. T.W. Malone, “How Do People Organize Their Desks? Implications for the Design of Office Information Systems,” *ACM Trans. On Office Information Systems*, vol. 1, no. 1, Jan. 1983, pp. 99-112.
4. E.D. Mynatt, “The Writing on the Wall,” *Proc. Of INTERACT’99*, IOS Press, Amsterdam, 1999, pp. 196-204.
5. J. Grudin, “Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use,” *Proc. of CHI 2001*, ACM Press, New York, 2001, pp. 458-465.
6. A.K. Dey, G.D. Abowd, and D. Salber, “A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications,” *Human-Computer Interaction (HCI) Journal*, vol. 16, 2001, pp. 97-166.
7. W.K. Edwards, “Coordination Infrastructure in Collaborative Systems,” doctoral dissertation, College of Computing, Georgia Inst. of Technology, Atlanta, 1995.
8. P. Maes, “Agents that Reduce Work and Information Overload,” *Communications of the ACM*, vol. 37, no. 7, July 1994, pp. 31-40.
9. J.C. Tang et al., “ConNexus to Awarenex: Extending Awareness to Mobile Users,” *Proc. of CHI 2001*, ACM Press, New York, 2001, pp. 221-228.
10. E. Horvitz, A. Jacobs, and D. Hovel, “Attention-Sensitive Alerting,” *Proceedings of UAI '99, Conference on Uncertainty and Artificial Intelligence*, Morgan Kaufmann, San Francisco, 1999, pp. 305-313.
11. B. MacIntyre et al., “Support for Multitasking and Background Awareness Using Interactive Peripheral Displays,” *Proc. of ACM UIST’01 Symposium User Interface Software and Technology*, ACM Press, New York, 2001, pp. 41-50.
12. J.D.A. Henderson and S.K. Card, “Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in Window-based Graphical User Interfaces,” *ACM Trans. On Graphics*, vol. 5, no. 3, July 1986, pp. 211-241.
13. E.D. Mynatt et al., “Flatland: New Dimensions in Office Whiteboards,” *Proc. of CHI’99*, ACM Press, New York, 1999, pp. 346-353.
14. J.I Hong and J.A. Landay, “SATIN: A Toolkit for Informal Ink-Based Applications,” *Proc. of ACM UIST 2000 User Interfaces and Software Technology*, ACM Press, New York, 2000, pp. 63-72.