

Integrating Virtual and Physical Context to Support Knowledge Workers

Stephen Volda, Elizabeth D. Mynatt, Blair MacIntyre, and Gregory M. Corso,
Georgia Institute of Technology

ABSTRACT

The Kimura system augments and integrates 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.

INTRODUCTION

Our research seeks to design an office that better supports knowledge workers—business professionals who interpret and transform information [1]. 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 [2, 3, 4, 5]. The diversity of these work practices and the complexity of implementing flexible computing tools make it difficult to meet all the workers' needs.

In typical office environments, knowledge workers use a variety of tools and information sources [Figure 1], including desktop and laptop computers, whiteboards, and desks. Information sources are frequently distributed throughout the office environment, both within and beyond the individual office walls. Although 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 this information, synthesize the disparate pieces of information, and act appropriately on that information.

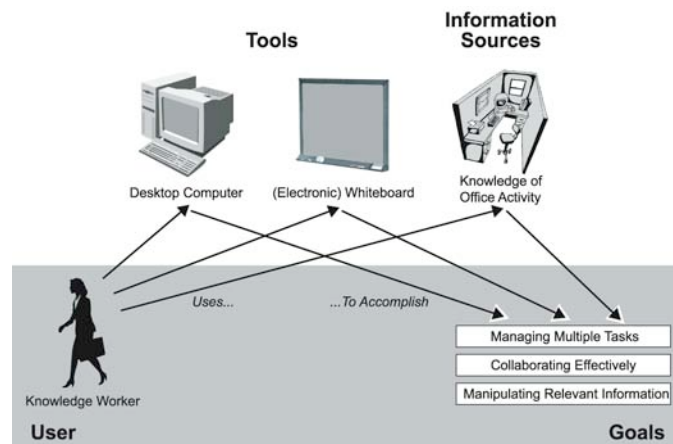


Figure 1: Knowledge workers must rely on a variety of independent tools and information sources to successfully accomplish their work goals.

We are working to augment and combine these independent tools into an integrated, 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 that we refer to as montages on a wall-size peripheral display—the electronic whiteboard [Figure 2].

The whiteboard lets users monitor each ongoing work activity, transition smoothly between activities, access a wide variety of contextual information designed to facilitate collaboration, and maintain awareness about relevant activity changes. Additionally, the interactivity provided by the electronic whiteboard allows the user to informally annotate and spatially organize the montages.

The montage design relieves the user of burdens associated with maintaining 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. The montages are designed to present this information without intruding on the user’s focal activity and in a manner that supports the needs of knowledge workers.

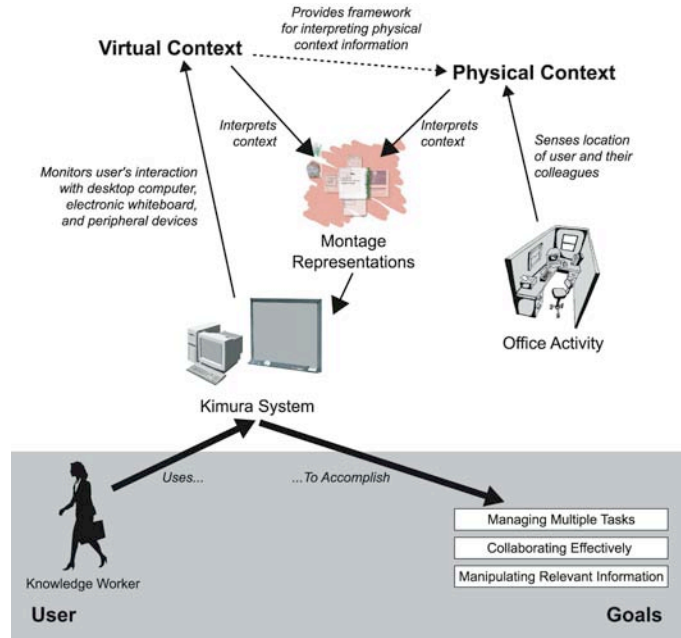


Figure 2: Our vision of a pervasive computing system leverages both physical and virtual context to reduce the number of tools and information sources that knowledge workers must use and synthesize to accomplish their goals.

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 [6]. 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 [8], 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 [8, 9, 10], but none, to the best of our knowledge, approach the scale of an integrated office environment.

Elsewhere, we presented the design and architecture of the Kimura system [11]. Here, we focus on the ways in which Kimura integrates physical and virtual context information to create a pervasive computing environment for knowledge workers. Kimura makes two key contributions. First, we use virtual context from the user’s desktop actions to help classify, interpret, and visualize other forms of virtual and physical context. Second, we integrate virtual- and physical-context information into visualization of the user’s disparate activities to help the user interpret and act on this available information.

The Kimura system

Kimura separates the user's "desktop" into two regions: the focal region, on the desktop monitor; and peripheral displays, projected on the office walls. Each work activity is 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 display [Figure 3].



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

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 have iconic indications of ongoing activity—including 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 [12], 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.

Each working context is displayed as a *montage* of images garnered from system activity logs [Figure 4]. These montages are analogous to the "room overviews" provided by other multicontext 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 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 background awareness information that is gleaned from our context-aware infrastructure.



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 context information (the availability of a colleague).

The electronic whiteboard—the primary display surface for the montage visualizations—supports common whiteboard practices [4]. Whiteboards feature an intuitive user interface and are well-suited to supporting informal information management activities. Our system implementation incorporates existing electronic whiteboard interaction techniques with montages and notification cues [13, 14]. This allows the user to annotate montages with informal reminders and 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.

Scenario

Wendy, our knowledge worker, walks into her office Monday morning following a week’s vacation. She scans the piles of paper on her desk and the contents of her whiteboard, recalling the work that has been waiting for her.

After quickly surveying the various whiteboard montages that represent ongoing activities, she annotates the budget plan with “Work on Wed., Due Friday” and throws it to the whiteboard’s far side.

The calendar image in the Acme design project montage reminds her of a design briefing later that day.

She studies the montage for a moment, trying to remember how far into the design briefing activity she was before she left on vacation. She sees opaque images of the documents she worked with most recently: her calendar, an illustration, a presentation file, and a Web search page. The montage also includes several ghostly images of past documents—two important email messages from her group’s client and the original project proposal. She taps on the 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 search 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. As she runs out of the office, she sees the icon for the completed print job. Grateful that someone, or something, is on top of things, she heads to the printer on the way to the meeting.

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
Email 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 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

Table 1: The kinds of virtual- and physical-context information collected by the Kimura system, the effect each has on the montage's appearance, and the specific user goals each supports.

CONTEXT-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 [Table 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.

The 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.

Virtual Context

Our system uses several lightweight monitoring components and proxies to acquire virtual-context information about the Kimura user. The focus is on capturing users' actions on their desktop computers and their interactions with the windows associated with each work activity. We 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 (that is, opening a window, changing the window focus, pressing a key, clicking the mouse) to our desktop monitoring process, which packages the event in a Kimura-readable activity log entry and sends it to a distributed activity log. Additionally, the desktop monitor creates a screenshot of each window every time the window system's input focus changes. The context interpreter integrates these screenshots into each montage so that the visual representations of the user's task can include actual images of the user's work. The images, similar to thumbnails, provide more relevant visual reminders than generic icons or labels.

This virtual-context collection mechanism does not connect directly with the running applications to find out 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 should not be bound to a small set of Kimura-aware applications—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, we can 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, to produce reasonably detailed visualizations of the user's overall activity in a given working context.

Kimura also acquires virtual context through an email monitoring system, tracking 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. It monitors all email messages that the user sends—the system associates each mail recipient with the active working context. The process also adds the recipient to a list of individuals with whom the Kimura user might be trying to connect, and instructs the location-monitoring component to actively monitor the availability of that individual by watching for their presence in public areas of the office.

In addition, Kimura observes 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 in a working context. As the status of each print job changes (for example, a print job is sent to the spooler, prints after being buried in a long queue, or stalls because the printer is out of paper), the context interpreter adds a notification cue to the appropriate montage.

Finally, we monitor all the user's interactions with Kimura to discover which background activities are of most interest given the user's current working context. We also use this information to build a model of

how activities are interrelated. The electronic whiteboard uses this information to draw the user's attention to relevant montages and to influence the initial placement of new montages on the peripheral display.

Physical Context

Kimura helps the user reconstruct the environmental circumstances surrounding a working context and provides cues about the user's colleagues' location and availability using physical context. In our current prototype, we simulate a pervasive, location-aware infrastructure with a series of Dallas Semiconductor i-Button docks distributed throughout the office environment. We designed our sensor network to detect the arrival and departure of known individuals in our augmented office environment, in public areas of the office, and near peripheral devices (that is, 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 he or she is not directly interacting with the desktop or electronic whiteboard. It also lets the system determine the general location of colleagues and infer when they might be available for collaboration or when they have joined the user in the augmented office for an informal meeting.

We also record the user's physical presence and activity at his or her office workstation using a keystroke monitor and a mouse movement monitor, but we hope to extend this 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 how focused the user is on a given task. We plan to use this information to adjust the rate of visual change on the 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.

Interpreting and Using Context

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.

In creating a high-level framework, we use some of the virtual context to guide the classification and interpretation of the rest of the virtual and physical context. The system uses 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, Kimura notifies the various system components, insuring that incoming physical- and virtual-context information is interpreted with respect to the proper working context.

The interpreted context information provides the details about each working context and is used to generate the montage visualizations for display on the electronic whiteboard. The montage designs take advantage of a number of visualization techniques to express the working contexts' semantics. To show a summary of a working-context at a glance, montages are based on thumbnail images of the user's activity in a working context and adapted to reflect the history of that activity. The size of each thumbnail image is determined by the amount of time that the user has spent interacting with the respective desktop application; large thumbnail images represent applications that consumed more of the user's time. We also represent recency of document use through image transparency; the most recently used documents are the most opaque. We are experimenting with a variety of thumbnail layouts, with the hypothesis that changing the layout over time, as a working context "ages," will foster better activity resumption.

FUTURE WORK

We are preparing to deploy a number of Kimura systems over the next several months to evaluate our design's usability and usefulness. We are also interested in enhancing our existing prototype to improve its robustness and flexibility and to introduce more advanced tools for collaboration.

One of our immediate goals is 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 approach to physical-context sensing. Leveraging a more mature context-aware infrastructure also lets us experiment with a more diverse array of sensing devices.

Another area of future work is to enhance our virtual context acquisition component so that it can extract more detailed information from common desktop applications (for example, Netscape Navigator, Microsoft Word, and so on). 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 running in memory on the desktop computer and providing an abstract view of users' interactions with the desktop computer at the windowing system level. By implementing a tighter coupling with key desktop applications, we expect to 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 beginning to explore ways in which the system can use the physical-context information it gathers to guide the interpretation of other context information. For example, sensed information about the user's location (that is, at the computer, at the printer, in a public space, or working with a colleague) can help determine which working contexts the user is engaged in. This capability lets the system update the user's working contexts with incoming context information even without explicit interaction on the desktop computer. A simple example implemented in our system is that the context interpreter can remove completed print job notifications from all working contexts when the user is sensed in the network printer's vicinity. A more complex example is this: when a colleague is sensed to have entered the user's office, the system will increase 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 associated with the visitor will become more prominent and easily accessible in anticipation of an informal meeting about one of those working contexts.

We are also interested in studying the implications of collaborative use of the Kimura system. We envision an office environment in which multiple work colleagues each have their own Kimura system. We would like to understand the interaction issues raised as colleagues collaborate over working contexts. These issues might include topics such as privacy, information access and visibility, and synchronization of shared information. We would also like to explore ways to adapt our existing montage designs to visualize the concurrent activity of different people in the same working context. Finally, we want to understand how the electronic whiteboard can be used more effectively as an informal meeting and collaboration space.

This work 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 aid the development of future montage animations so that users will immediately notice critical changes on the peripheral display and ignore less important changes. We are also performing user tests on the effectiveness of our montage designs in supporting suspended task resumption. Finally, we are completing our own task analysis of knowledge work in the office environment to supplement and clarify the findings in the published literature. We hope to combine all of these findings with the results of our Kimura pilot study to inform the design of future system prototypes.

ACKNOWLEDGEMENTS

NSF grant number 9988712 and research and equipment grants from IBM and Ricoh Innovations supported this work. We also thank our colleagues and students for their feedback and ideas on this article.

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. Office Information Systems*, vol. 1, no. 1, Jan. 1983, pp. 99–112.
4. E.D. Mynatt, "The Writing on the Wall," *Proc. 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: Human Factors in Computing Systems*, 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 J.*, vol. 16, no. 2–4, 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," *Comm. 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. CHI 2001: Human Factors in Computing Systems*, ACM Press, New York, 2001, pp. 221–228.
10. E. Horvitz, A. Jacobs, and D. Hovel, "Attention-Sensitive Alerting," *Proc. Conf. Uncertainty and Artificial Intelligence (UAI '99)*, Morgan Kaufmann, San Francisco, 1999, pp. 305–313.
11. B. MacIntyre et al., "Support for Multitasking and Background Awareness Using Interactive Peripheral Displays," *Proc. ACM Symp. User Interface Software and Technology (UIST '01)*, 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. Graphics*, vol. 5, no. 3, July 1986, pp. 211–241.
13. E.D. Mynatt et al., "Flatland: New Dimensions in Office Whiteboards," *Proc. CHI'99: Human Factors in Computing Systems*, 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. ACM Symp. User Interface Software and Technology (UIST '00)*, ACM Press, New York, 2000, pp. 63–72.