Context Histories, Activities, and Abstractions: Ubiquitous Computing Support for Individual and Collaborative Work

Stephen Voida & Elizabeth D. Mynatt
GVU Center, College of Computing
Georgia Institute of Technology
85 5th Street NW, Atlanta, Georgia 30332-0760 USA
{svoida, mynatt}@cc.gatech.edu

INTRODUCTION

As ubiquitous computing technologies find their way into widespread use and become an "invisible" and pervasive part of users' everyday practices, the relationship that users have with these technologies will begin to change. While users may have been content to adapt their practices to match the information management strategies and "application-document" models imposed on them by computers (and their designers) in the past, they will be less willing to do so as computers find their way into more aspects of everyday life and mediate more of our human-human interactions. Gay and Hembrooke have noted a corresponding shift in the language used by HCI practitioners—where *user-centered design* used to be the touchstone of the field, the ideas of *activity-* and *context-centered design* are becoming increasingly prevalent [4].

The workplace is a particularly interesting setting for studying this transition. Although the desktop computer is a long-established fixture in the office, mobile phones and networked devices like the RIM Blackberry have, for many workers, become as common and just as indispensable. The proliferation of Web-based corporate applications, virtual private networks, and VOIP telephony has extended the boundaries of the traditional workplace so that work now occurs in many non-traditional locations—and "on the go." Furthermore, these new technologies, in many cases, have not *replaced* existing technologies so much as they have served to *augment* them; the role of each technology is constantly changing, but the overall complexity of the workplace is, in general, on the rise.

As the *amount and diversity of incoming information* confronting knowledge workers steadily increases, the *devices used to carry out work* multiply, and the *locations in which work takes place* become more varied, more traditional computer-based practices for organizing and managing work begin to break down. Email is the most common example of this trend—it is widely acknowledged that email has become an incredibly overloaded medium, serving not only as a means for communication, but for coordination, scheduling, task-awareness, organizational memory, document sharing and version control (to name just a few) [2].

An increase in the amount of contextual information collected in the workplace and available to knowledge workers can be (and in fact *is*) part of this problem: it is just that much more information to be managed. However, it can also be an asset for helping users to maintain an overall awareness of their work environment, their ongoing work tasks, and the state of their collaborations with others, as well as a memory aid in task resumption. Our research has focused on the iterative development of computing systems that support these goals, based on models of activity created by compiling many sources of virtual and physical context. Such systems provide a structured environment that serves to organize work artifacts and context in a manner more consistent with knowledge workers' actual work practices.

Our research program lies at the intersection of two major bodies of research: activity-based computing and context-aware computing. Several field studies on the role of tasks and activity in the workplace have recently been published (e.g., [1, 5]) and some initial activity-based computing prototypes have been developed (e.g., [8]). Other research has focused on how context can be utilized as a part of existing work practices, most commonly as a tool for awareness and interruption management (e.g., [3]). Our initial explorations have been focused on investigating the role of activity modeling, peripheral displays, and integrated context-aware frameworks in supporting individual work. We are interested in expanding the scope of our inquiry to explore how adding collaboration support changes the requirements for activity- and context-aware systems.

Activity and Context in the Kimura System

Our prototype system, Kimura, was developed to help us understand how activity models, peripheral displays, and context-awareness could be used to support task-awareness and multitasking in knowledge work [7]. The Kimura prototype combines a desktop computer running a custom virtual desktop manager with an electronic whiteboard and context-aware infrastructure. As in previous systems like Rooms [6], users create virtual desktops on the computer to separate and organize their various work activities. Kimura builds a model of activity based on the "virtual context" of users' interactions with the desktop computer and virtual

Figure 1. The Kimura system, including a desktop component, two interactive peripheral displays with electronic whiteboard capabilities, and a third non-interactive peripheral display. The images projected on the electronic whiteboards are *montages*, representations of activity that integrate history and context information.

window manager. It then integrates other virtual and physical context sensed by the context-aware infrastructure into the model. We call the resulting clusters of computational artifacts and contextual cues *working contexts*, and display a representation of each, called a *montage*, on the electronic whiteboard. Users can view the whiteboard as a passive peripheral display and monitor the state of all ongoing work activities. They can also interact with the whiteboard directly to annotate, organize, and switch among working contexts.

Kimura's integration of virtual and physical context is unique [9]. The system creates a high-level framework of working contexts based on the virtual context—the user's manipulation of the virtual desktops and other interactions with the desktop computer—within which other virtual- and physical-context information is classified and interpreted. 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.

This combination of interpreted context information provides detailed representations of each of the user's activities and is used to generate the montage visualizations displayed on the electronic whiteboard. The montage designs take advantage of several visualization techniques to express the working contexts' semantics. To show a summary of a working-context at a glance, montages contain thumbnail images of the user's desktop computer applications as well as icons representing relevant external context for each activity. These representations are also adapted to reflect the history of each activity, including the *relevance* of individual aspects (for example, time spent interacting with a given artifact or the inferred importance of a contextual cue) as well as their relative *recency* (providing a sense of the temporal evolution of the activity).

For a typical knowledge worker, Kimura might monitor a number of concurrent work activities, displaying a montage for each on the electronic whiteboard. Currently, these montages convey to the user what applications and documents have been accessed over the course of each work activity, which documents have been most important, any annotations the user has provided, and other context information about each activity such as whether colleagues affiliated with an activity are available for face-to-face collaboration (if they have been sensed in an office common area) or whether a print job relating to an activity has been completed and is awaiting retrieval.

CHALLENGES OF MODELING ACTIVITY AND CONTEXT HISTORIES FOR INDIVIDUAL WORK

Our experiences with the Kimura system confirmed our intuitions (and others') that activity can be a potentially powerful organizing principle for dealing with the increasing complexity of knowledge work. We feel that there are strong benefits for providing these representations of activity and context to both desktop and ubiquitous computing applications so they might assist the user in switching among ongoing tasks, creating new ongoing tasks that resemble previous ones, and maintaining an awareness of the tasks in which they are currently engaged. However, our initial models for representing activities and their associated context have proven to be somewhat inadequate for authentically modeling real-world work practices.

Models of activity should enable the expression of different classes of activities such as routine tasks and recurring tasks and different types of activities such as information analysis tasks and content production tasks. They should also be able to encode a broad range of affiliated context such as the location in which an activity was accomplished, the time (or frequency) at which it occurred, the individuals with whom the activity was carried out and what specific contributions each made. Systems implementing these sophisticated models will further benefit from maintaining details of activity and context over time, so that trends can be monitored and patterns detected, leading to representations of emergent behavior and enabling systems to suggest procedures or artifacts that have been useful in similar situations.

We envision a system like Kimura that enables users to demarcate their work activities and to organize their computational artifacts, relevant communications, colleague contact information, and personal reminders as an implicit part of their existing work practices (or with as little additional overhead as possible). This system should also allow users to search for past material using rich contextual cues as indices into past activities or recommend relevant information based on contextually-similar situations to ones the system has seen before.

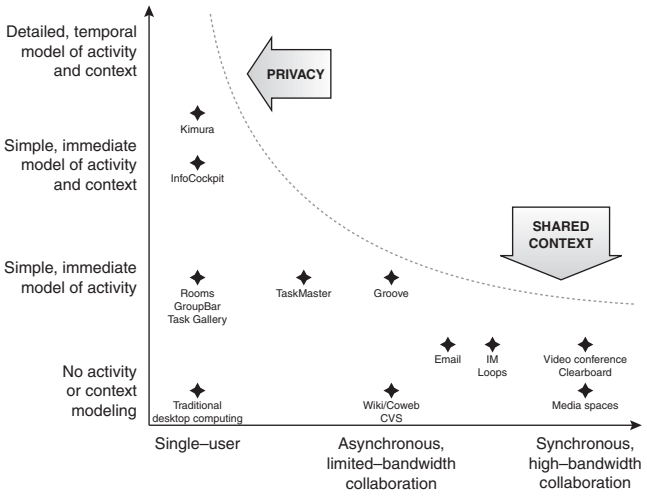

Figure 2. A design space for collaborative, activity- and context-aware applications. The large arrows indicate potential influences on the distribution of current applications within the design space.

***Challenge #1*: Integrating virtual and physical context to create a coherent model and history of user activity**

What are the critical characteristics of modeling activity and context over time? A successful model should reflect the findings of research on workplace activity and enable useful individual task management in ways not available with today's systems. However, maintaining a balance between flexibility and complexity will be important in order for applications to be able to utilize the modeled data and for users to be able to manage their representations.

CHALLENGES POSED BY COLLABORATION

However beneficial enhanced models of activity and context might be for supporting individual users, potential tensions exist *because* most activity-aware systems are targeted at individual use and many "real-world" knowledge work activities are inherently collaborative.

In order to understand how this tension has played out in existing systems, we constructed a design space illustrating the sophistication of activity-awareness and collaborative complexity of several commercial and research Ubicomp and workplace applications (Figure 2). Most of these systems cluster toward the individual-use, activity-aware portion of the diagram (the left-hand side) *or* toward the collaborative, non-/marginally-activity-aware portion of the diagram (along the bottom). We speculate that two forces may be acting on the position of systems in this design space: *privacy* and *inherently shared context*. The cluster of systems along the vertical axis may be constrained by concerns about privacy. These systems encode significant details about individual activity and context but are not equipped to represent these models appropriately for collaborative situations. In contrast, the cluster of collaborative applications along the horizontal axis may inherently convey some degree of shared context and activity-awareness as a by-product of the collaboration process. As a result, it may not be necessary for these applications to explicitly encode models of activity or context in order for the interaction to be successful in the context of working in a group.

***Challenge #2*: Addressing privacy concerns when collaborating with sophisticated models of user activity and context history**

As more and more detail about a user's actions and the context in which he or she carried out their work are captured and stored, the risk of having this potentially personal information inadvertently shared with others over the course of collaboration grows. Finding a balance between activity- and context-awareness and collaboration support requires difficult design trade-offs.

***Challenge #3*: Accommodating differences in granularity of activity specifications**

There will almost certainly be cases in which two users need to coordinate activities and context histories established independently. The way in which these models are specified will determine the complexity of "merging" the two models, particularly for cases in which the users conceive of and manage their activities at different levels of granularity. Resolving these differences elegantly is critical to these systems' success. The development of user interfaces and visual representations to ease merging models will likely be a critical area for research.

The Role of Abstractions

We are interested in developing tools that support all aspects of knowledge work, including individual work and collaboration. However, in order to do so, we need to find ways to overcome the potential privacy issues involved in sharing personal activity and context information, and, if possible, integrate the representations of shared context in the collaboration process itself, as do many existing tools.

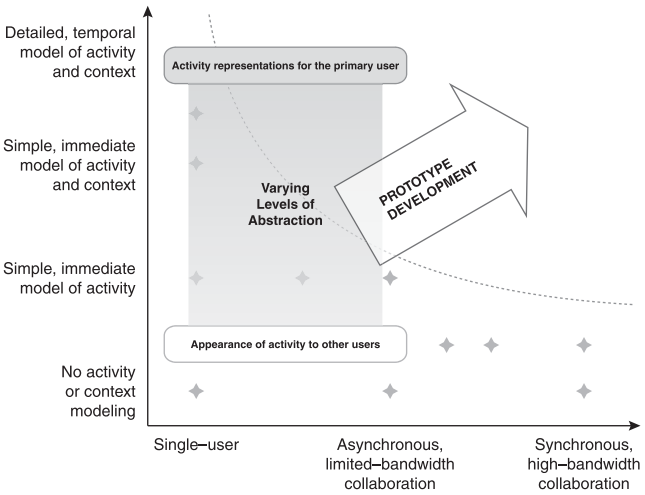

Figure 3. Our proposed work focuses on the development of systems that encode rich activity and context histories, but also provide user interfaces and representations for controlling abstractions, so that the systems can be appropriated for individual use *and* collaboration.

We hypothesize that providing varying levels of abstraction in our activity and context histories can allow users to specify the level of detail most appropriate for a given situation: while collaborating with a particular group of colleagues, working in a specific location, or working on a particular device (Figure 3). We believe that this approach gives users the most flexibility, allowing them to take full advantage of activity- and context-awareness when working individually and providing them access to activity and context information when needed during collaboration.

Challenge #4: **Identifying critical characteristics of activity and context histories for which collaboration hinges on having the right abstraction(s)**

Abstractions will likely be more critical for some aspects of activity and context histories than others. Due to the potential complexity of these histories and the myriad ways abstraction could be used to limit the disclosure of personal information, identifying the information users are most interested in protecting—and to what degree that information needs to be aggregated, anonymized, or excluded from histories shared with others—will be critical in informing the design of appropriate abstractions.

Challenge #5**: Providing user interfaces to manage abstractions**

Users will likely need to provide some degree of direct control or fine-tuning over the abstractions used in a given situation. However, this requires imposing additional "meta-work" on top of the work practices users already have in place. What user interfaces are most appropriate for managing abstractions of activity and context histories? Are there instances in which implicit observation of existing work practices can be used to determine the appropriate abstraction to apply?

Challenge #6**: Examining the role of the user's location and the devices they use in selecting an appropriate level of abstraction for a given context**

Our initial explorations have taken advantage of a subset of Ubicomp technologies we felt most appropriate for integration into an individual's existing workspace. Can the virtual or physical context sensed using a broader range of devices (including those specifically designed to support collaboration) be used to reliably infer the level of abstraction most appropriate for a particular situation?

OBJECTIVES FOR THE WORKSHOP

We are looking forward to participating in the ECHISE2005 workshop since it appears be an ideal venue for us to refine and inform our intuitions about the challenges in designing these types of systems based on the research being carried out by others in the field.

We are particularly interested in discussing the methods that are being used to model complex activity and context histories in other systems, the means for abstraction that these other approaches employ, and the user interface conventions others have found successful for representing and providing user control over context histories.

AUTHOR BIOGRAPHIES

Stephen Voida is a Ph.D. student in the Georgia Institute of Technology's College of Computing and a member of the GVU Center and the Everyday Computing lab. His research interests include ubiquitous computing, technology in the workplace, and augmented environments. He received his M.S. in human-computer interaction from the Georgia Institute of Technology in 2001.

Elizabeth D. Mynatt is an associate professor in the Georgia Institute of Technology's College of Computing and the GVU Center. She directs the Everyday Computing research program within the Future Computing Environments group, examining the implications of having computation continuously present in many aspects of everyday life. Her research interests include exploring how to augment everyday places and objects with computational capabilities. She received her Ph.D. in computer science from the Georgia Institute of Technology in 1995.

REFERENCES

1. Czerwinski, M., Horvitz, E., and Wilhite, S. A diary study of task switching and interruptions, in *Proceedings of CHI 2004* (Vienna, Austria, April 2004), ACM Press, 175–182.
2. Ducheneaut, N. and Bellotti, V. Email as habitat: An exploration of embedded personal information management. *ACM Interactions*, *8*, 5 (September–October 2001), 30–38.
3. Fogarty, J., Hudson, S.E. and Lai, J. Examining the robustness of sensor–based statistical models of human interruptability, in *Proceedings of CHI 2004* (Vienna, Austria, April 2004), ACM Press, 207–214.
4. Gay, G. and Hembrooke, H. *Activity-centered design: An ecological approach to designing smart tools and usable systems*. MIT Press, Cambridge, Massachusetts, 2003.
5. Gonzàlez, V.M. and Mark, G. "Constant, constant multi-tasking craziness": Managing multiple working spheres, in *Proceedings of CHI 2004* (Vienna, Austria, April 2004), ACM Press, 113–120.
6. Henderson, J.D.A. and Card, S.K. Rooms: The use of multiple virtual workspaces to reduce space contention in window-based graphical user interfaces. *ACM Transactions on Graphics*, *5*, 3 (July 1986), 211–241.
7. MacIntyre, B., Mynatt, E.D., Voida, S., Hansen, K.M., Tullio, J., & Corso, G.M. Support for multitasking and background awareness using interactive peripheral displays, in *Proceedings of UIST '01* (Orlando, FL, November 2001), ACM Press, 41–50.
8. Smith, G., Baudisch, P., Robertson, G., Czerwinski, M., Meyers, B., Robbins, D. and Andrews, D. GroupBar: The TaskBar evolved, in *Proceedings of OZCHI '03* (Brisbane, Australia, November 2003), University of Queensland, 34–43.
9. Voida, S., Mynatt, E.D., MacIntyre, B. and Corso, G.M. Integrating virtual and physical context to support knowledge workers. *IEEE Pervasive Computing*, *1*, 3 (July–September 2002), 73–79.