UNDERSTANDING AND SUPPORTING TRANSITIONS WITH LARGE DISPLAY APPLICATIONS

by

ANTHONY TANG

B.Sc., Simon Fraser University, 2002

M.Sc., University of Calgary, 2005

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES

(Electrical and Computer Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

January 2010

Abstract

Interactive large displays offer exciting new opportunities for collaboration and work. Yet, their size will fundamentally change how users expect to use and engage with computer applications: a likely reality is that such displays will be used by multiple users for multiple simultaneous tasks. These expectations demand a new approach for application design beyond the conventional desktop application model, where applications are single-user, and intended to support a subset of user tasks.

In this research, we develop such a framework based on the premise that large display applications should support *transitions*—users' desires to shift between multiple tasks and activities. We build this framework from models of how traditional large surfaces such as whiteboards are used to facilitate multiple tasks—often simultaneously.

Based on studies of users' whiteboard use, we construct a classification scheme of users' activities with whiteboards, and the role of whiteboards in supporting the transitions between these activities. From a study of meeting room activity, we then develop a classification for collocated activity around traditional surfaces. We further develop models of how users' needs change during their use of large display applications, exploring two contexts: a digital tabletop application for focused collaboration, and a public large display. These studies reveal how users engage and disengage with one another during collaborative work, and the dynamic needs of bystanders.

Next, we design and evaluate a prototype that supports transitions between tasks in a scheduling activity using viewing changes. The results demonstrate that users transition between related tasks during such activities, and that viewing changes can support these transitions. Finally, we describe a design space for supporting transitions in large display applications.

Taken together, the findings of this research illustrate the fundamental need to develop a new framework for designing large display applications. This work provides a step in this direction by providing rationale and empirical evidence for supporting transitions in this framework. In so doing, it suggests that we realign designers' efforts from the predominant desktop-centric model of application development, and instead to a model that engenders smooth transitions between multiple, related activities.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I admit that this section has received more thoughtful consideration more frequently and more consistently than any other section in the entire dissertation. I thought about it when things were going well, and I thought about it when things were going poorly—a testament to what the PhD challenge was for me: an intellective one, certainly, but also an emotional one. I now see those times, the good ones and the bad ones, differently—they were all things that helped to prepare me to become the person that I am and will become. And—as challenging as many of those moments were—I am grateful for them, and for the people that saw me through them.

Cheryl, my best friend, my partner—thank you for the long nights of listening to me whine and complain, for sharing in the highs and lows, and always trying to support me in the ways you best knew how.

Sid, my supervisor, my biggest critic, and the one that best knew how to make me into a better researcher—thank you for pushing me the way you did. It was undoubtedly challenging, but helped make me into the researcher I am today: one who is more thoughtful, critical, and careful.

Carman, my friend, my colleague, my mentor—thank you for your friendship, your thoughts, and the countless hours of reading drafts, commenting on them, and for talking me through the hardest parts.

My committee, Kelly and Peter, were immeasurably helpful at each critical time. Thank you in particular to Kelly for helping me to frame my thoughts, and to help me find my confidence in my own analytic ability.

To my local "research family," Joel, Nelson, Mike, Leah, Karyn, Kirstie, Karen, Nicole, Meghan, Garth, Rock, Matt and Ian, thank you for your companionship and support: you helped me ride through the worst of times, and pushed me up onto the saddle when I was too bashful during the best of times.

To my remote research family, Petra, Nelson, Mark, Gregor, I have appreciated your support and friendship—it will be fun continuing to make a splash in the research world with you.

Saul, Sheelagh, and Joanna, my life-supervisors, you have been invaluable as life mentors, helping me to understand my needs, my goals, my strengths and my weaknesses.

I thank and acknowledge NSERC, NECTAR, and UBC for financial support.

To my volleyball family, Jill, Becky, and Heather, thanks for keeping life light, and for helping to remind me that there can be important things outside of work ... and to not take life too seriously. Thank for helping to keep me sane when the research threatened to drown me.

Finally, to my buddy and brother Jonathan: you are the best friend a guy could ask for. I can't count the times I needed to chill out, and you were there. Whenever I think about the possibility of having done this PhD elsewhere, I can't help to think where our relationship would be. To have been able to spend the last five years growing up with you as I went through this challenge is a blessing that I cannot put into words.

Dedication

For my mother Eva, and my father Tom.

Thank you for sharing in my dreams.

Statement of Collaboration

The thesis work in this dissertation was conducted under the primary advisement of Dr. Sidney Fels. Dr. Kellogg Booth provided invaluable help in the written reports of this work in general. Dr. Sheelagh Carpendale played a supervisory role in the tabletop study reported in Chapter 4.

I am the primary contributor to all aspects of this research. For each publication listed below, I detail the contributions of each co-author.

- An earlier version of parts of Chapter 2 was presented as:

  Tang, A. and Fels, S. 2008. Four lessons from traditional MDEs. ACM CSCW 2008 Workshop on Beyond the Laboratory: Supporting Authentic Collaboration with Multiple Displays.

- Parts of Chapter 2 also appeared in:

  Tang, A. 2006. Surface use in meeting room collaboration. In *Conference Companion of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY, 43-44.

- With regard to the whiteboard study presented in Chapter 3, Dr. Joel Lanir was a fellow graduate student at the time, and helped in discussions and early framing of the work. The study was discussed with Dr. Saul Greenberg. Both co-authors provided feedback on drafts of this work. The study was previously published as:

  Tang, A., Lanir, J., Greenberg, S., and Fels, S. 2009. Supporting transitions in work: informing large display application design by understanding whiteboard use. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work* (Sanibel Island, Florida, USA, May 10 - 13, 2009). GROUP '09. ACM, New York, NY, 149-158.

- The tabletop study presented in Chapter 4 was previously published. Dr. Melanie Tory and Dr. Barry Po were post-doctoral fellows at the time of this work, and aided in the planning, execution and analysis of the data. Petra Isenberg (née Neumann) aided in the execution of the study.

  Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. 2006. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 1181-1190.

- The study of MAGICBoard presented in Chapter 4 was previously published. Dr. Matthias Finke and Dr. Michael Blackstock helped design and manage the development of the prototype. Rock Leung and Meghan Deutscher provided helped with the visual design of the system, while Dr. Rodger Lea provided intellectual grounding for the work.

  Tang, A., Finke, M., Blackstock, M., Leung, R., Deutscher, M., and Lea, R. 2008. Designing for bystanders: reflections on building a public digital forum. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 879-882.

Chapter 1

Introduction

Visual displays are the dominant method that computers use to communicate with us. Yet, our everyday experience with them remains mainly limited to the desktop-sized screens that are attached to our PCs. As the price of manufacturing displays decreases, large displays (whiteboard sized or larger) will increasingly become an everyday reality. Already, large displays are beginning to permeate our homes for entertainment, and public spaces for advertising. This dissertation addresses the unique challenges of designing large display applications for our work spaces. The overarching goal of this research is to uncover these complexities as they relate to users' interactions with traditional large surfaces and large display applications in order to serve as a foundation for a framework that guides designers of large display applications. Ultimately, the intention is for this framework to help designers build applications that integrate with familiar workflows while providing the advantages of new technology.

Most current large display systems have been designed using a conventional, desktop-centric application model. Here, computational support for user tasks are embedded in applications, with each application addressing a reasonable subset of user needs. When a user's activity comprises tasks beyond the scope of any one of these subsets, the user needs to use multiple applications (for instance, sending a written document to a colleague requires the use of a word processor and an email client). In this case, the user then needs to explicitly switch between different application windows to find the appropriate functionality for each task. The system model that addresses how these applications work in concert is data-sharing, either through intermediary files (e.g. saving the document, and loading it again) or the virtual clipboard (e.g. copying and pasting the text).

If we believe that large displays are to supplant the role of current traditional large surfaces (e.g. whiteboards), how appropriate are our assumptions about how applications should be designed? The context in which these large surfaces are used is different from conventional desktop applications in at least two fundamental ways:

(1) Large traditional surfaces are often used for multiple concurrent tasks. That is, a whiteboard provides multiple functional roles to its users simultaneously. Figure 1.1,

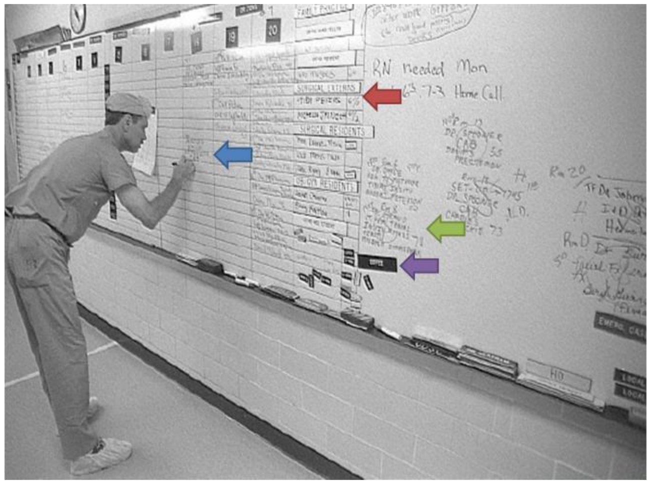

Figure 1.1 A whiteboard at a nursing station fulfills multiple roles simultaneously.

captured from a nursing station, illustrates this concept. Each arrow points to a different area of the whiteboard that is providing a different function to the nurse.

(2) Multiple users may be employing the same large display simultaneously, each with different, changing (and in some cases, competing) information and interaction needs. In a typical presentation meeting, comprising at least one presenter and one audience member, each individual has different information and interaction needs of the large surface: the presenter needs to be able to modify and manipulate the information, whereas the audience member just needs to be able to see the content being presented (and perhaps retrieve it at a later time).

Thus, a traditional large surface, between facilitating multiple functional roles and multiple users, supports multiple simultaneous tasks—some of which may be immediate and requiring direct interaction (e.g. brainstorming), and others which involve ongoing and intermittent use (e.g. ambient task list). The need to support this range of activities is at the heart of this research.

The general approach taken to understand these activities has relied on studying users' needs and work practices around these traditional large surfaces. The consequences of this particular approach are detailed later, but the general rationale is that users' work practices with traditional surfaces reveal insights into the user's model—his or her goals, information needs, and expectations—of the surface, and how it fits into the user's work. Developing our understanding of users is important so that interface design for large display applications can align with users' models—respecting their goals, information needs, and to some extent, the work practices that users employ to achieve them.

The findings of this research suggest that in the context of large display systems, a different framework of application design may be more appropriate than the conventional desktop model. This dissertation lays the groundwork for this alternative framework, in particular focusing on the role of *transitions*—users' ability and desire to shift between different tasks or activities—as a cornerstone in users' interactions with large displays. The research develops models of the roles

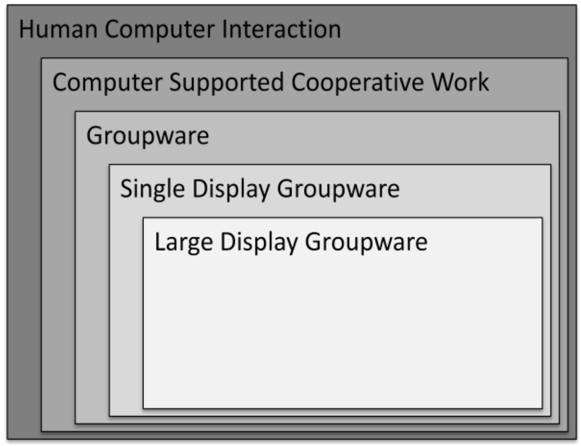

Figure 1.2 A visualization of how our research is situated within the broader research space.

large surfaces play in supporting work, and of users' changing needs with large display applications, where transitions bridge different application states.

Research Context

We situate this work within the broad context of human-computer interaction (HCI), where researchers examine the relationship between humans (users, non-users alike) and computer systems (Figure 1.2). Within HCI, the work fits within the sub-discipline of computer-supported cooperative work (CSCW), which is ultimately focused on the design of tools to support cooperative work. CSCW is informed both by the design and evaluation of tools, as well as the study of contextualized work practice—an approach we emulate in our work. These tools are generally called groupware, meaning that they are designed to support groups rather than just individuals. A useful groupware classification scheme is the time-place taxonomy (Johansen, 1988), which partitions the design space along two axes: time (same vs. different) and place (same vs. different). This classification emphasizes the *temporal* nature of collaboration that the tool supports (same-time—referring to synchronous activity, different-time—referring to asynchronous activity), and the *location* of users (same-place—referring to users who are collocated, different-place—referring to users who are distributed).

We focus specifically on the same-place part of this design space. Some researchers have explored the use of multi-display environments to support cooperative work (e.g. Biehl et al., 2008; Johanson et al., 2002a; Stefik et al., 1987); our interest, however, is on how single displays can be used to support cooperative work (e.g. Elrod et al., 1992). While it is possible to consider a range of display types (e.g. visual, aural) in a range of sizes (e.g. PDA-sized to wall-sized), we are specifically interested in designing large, whiteboard-sized surfaces that support work.

Consistent with the desktop application model, most prior systems focus on only a small subset of functionality, such as *knowledge work*, facilitating creation, modification or sharing of information (as in Stefik et al., 1987; Pedersen et al., 1993; Russell et al. 2002; Brignull et al., 2004), or *ambient awareness* (Miller & Stasko, 2001; Churchill et al., 2004). Yet this application-centric approach to supporting subsets of user tasks can be problematic when applied to the broader context of work.

An instructive example that demonstrates the consequences of this approach comes from the reports of the use of MERBoard by Huang et al. (2006b). The MERBoard was an interactive large display application designed for engineers and scientists at the NASA Jet Propulsion Labs (JPL) in California to aid their work on the NASA Mars Exploration Rover mission (Trimble et al., 2003). The design of the MERBoard comprised a suite of applications, including SolTree (a planning tool for controlling the rover units), a typical whiteboard application, the Mars clock application (i.e. the universal mission clock), and a tool for accessing the internal schedules of other researchers at the facility. After the MERBoard had been deployed for several months, Huang et al. interviewed participants on their usage of the system. Some of the findings pertinent to this discussion include:

- ***Poor support for multiple tasks***. Users complained of the need to explicitly export plans generated from the SolTree and whiteboard tools into presentation tools. Even though idea generation (using the SolTree or whiteboard tools) was to be conducted on the same display surface as presentation, MERBoard required users to explicitly export the information, switch applications, and then re-import the data into a different tool. This instance mirrors the example described earlier about having to use a word processor to write a document, and an email client to then send the document to a colleague: in both cases, the general activity comprises two user tasks, which are supported in different applications, and requiring an explicit mechanism for information interchange between the applications.
- ***Competing needs of users dissuade use of some applications***. After several months of use, *ambient* uses for the MERBoard became common (e.g. displaying the Mars clock or using the whiteboard for displaying images from the Mars rovers), at the expense of the use of the SolTree or whiteboard tools for idea generation. When not "in use," for active idea generation, the displays would be turned to these ambient uses; however, because would-be users of the SolTree or whiteboard tools could not be certain if anyone was making use of the Mars clock or photos, they were disinclined to use them. The competing needs of users meant that some felt averse to using the MERBoard.

These findings likely do not reflect the patterns of behavior expected by the designers; however, they are a powerful example of how the designers' model of the large display's role did not account for either the need to support multiple tasks, nor the competing needs of different users. The design neither supported smooth transitions between tasks, nor between the needs of users. In defense of those researchers, the MERBoard was one of the first real-world deployments of large display technology; furthermore, to expedite the deployment process, they relied on off the shelf applications (which had likely been initially written for desktop computers).

Studies of work practice around traditional surfaces suggest that they are used in ways that will often support multiple classes of activities simultaneously (e.g. Mynatt et al., 1999; Xiao et al., 2001; Teasley et al., 2000). Indeed, the whiteboard in Figure 1.1 illustrates that whiteboards can function in many roles simultaneously. Such use has benefits to users: it reduces the cost of context switching between related tasks or activities. Similarly, we would expect that providing smooth support to transition between tasks and activities is likely to provide benefits.

Designing with this conceptual framework is different. To support these transitions means shifting our focus from the functionality of a given *application*, and instead to the changing, organic needs of users as they conduct their work, dynamically switching between tasks as their information needs evolve. Such a perspective realigns the designer to one more in tune with users' practices with large traditional surfaces. Here, a whiteboard is not considered an application—instead, it is a flexible medium that users can appropriate in ways that allow them to conduct and perform work.

Research Goals

The overarching purpose of this research is to formulate and develop a framework that guides the design of large display applications. This framework should capture the richness of the multiple roles that large displays play in supporting users' work, and the interaction of users around such displays, with specific attention to aspects of large displays that distinguish them from the conventional desktop computer. Our thesis work has four goals. The first two goals provide grounding for this framework, developed from observations of users and their use of traditional large surfaces, and large display applications. Together, they motivate our concept of transitions, and their potential role in large display application design. The third goal illustrates how transitions can be supported in a concrete, proof-of-concept tool. Finally, the fourth goal provides designers with a way of thinking about transitions, and how they can be addressed in future research.

Thesis Goal 1: Create a classification scheme for the different roles traditional large surfaces play in supporting users' work. The first goal was to identify and articulate the multiple roles of traditional surfaces. In particular, our intention was to capture and describe the essence of users' activities around traditional large surfaces, and further to understand how affordances of traditional surfaces supported users' transitions between these activities. This classification would motivate the types of activities that should be supported with large display applications, and should reveal how transitions between these activities could be managed.

Thesis Goal 2: Formulate models that describe users' changing information and interaction needs during their use of large display applications. The second goal was to explore scenarios that demonstrate how users' needs change, and where users have differing needs of the large display application. These examples would provide insight into the dynamic role of users with regard to the design and use of large display applications—that users are not symmetric, static agents, and instead that their goals and needs change. The models should therefore motivate deeper consideration into how to support the needs of multiple users on a shared large display.

Thesis Goal 3: Design and develop a prototype proof-of-concept tool that supports transitions. The purpose of the third goal was to demonstrate how to apply the concept of transitions (developed in the first two thesis goals) to large display applications, and to realize this support in a concrete instantiation.

Thesis Goal 4: Create a design space for supporting transitions to inform the future designs of large display applications. This design space should reveal areas of design that have not been explored by prior work, but may be promising. Furthermore, to demonstrate its validity, it should be possible to categorize other systems that support transitions within the design space.

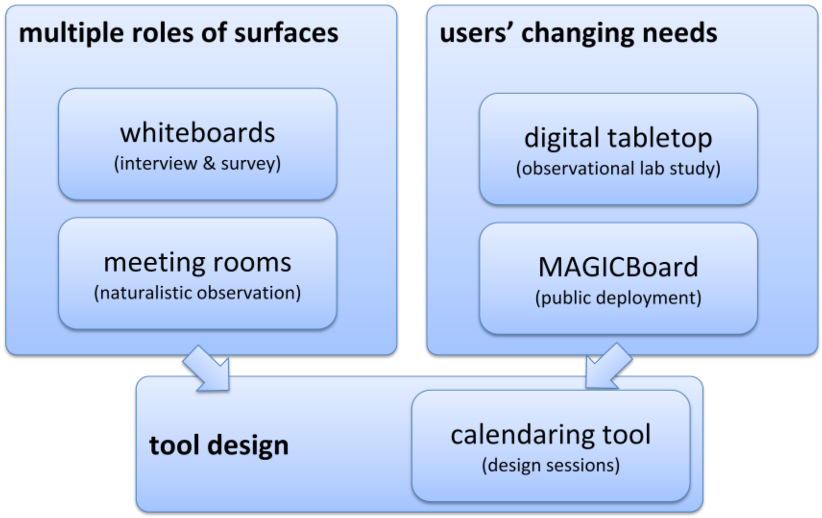

Figure 1.3 A visualization of our research process.

Research Approach

We used a research process that explored both users' existing practice with traditional large surfaces, and users' behaviours with novel large display applications. Addressing our overall research goals required incorporating both these perspectives to allow the findings from each context to complement one another. We provide details of the specific approach of each study in its respective chapter, but our overall process has been pragmatic: we used methods that suited the particular research question and context.

Figure 1.3 illustrates the three major phases of our research process. We describe each of these phases in turn.

The first major phase of work addresses our first research goal, which was to classify the different roles of traditional large surfaces. To accomplish this goal, we conducted two studies: first, an exploratory study of whiteboards, and second, a naturalistic observational study of collaboration in meeting rooms. As reported in Chapter 3, the whiteboard study explored users' whiteboard practices, focusing on how their work practices addressed their goals and information needs. A study of meeting room collaboration examined the roles of traditional surfaces for team collaboration. Based on these studies, we produce two classification schemes: first, a taxonomy of activity modes that circumscribes the ways in which whiteboards are used; second, a descriptive framework of the different roles traditional surfaces support in meeting room work. Together, these studies showed that traditional surfaces such as whiteboards play multiple roles in users' activities, and further illustrate how users appropriated traditional surfaces (whiteboards in particular) to bridge these activities and roles.

The second research goal is addressed by the second major phase of work, which was to model users' changing needs around large display applications. We designed, implemented and studied the use of two large display applications to address this goal in two different contexts: first, to study user needs during collaborative work, we designed a digital tabletop application, and second, to explore user behavior in public contexts, we designed and ran a field deployment of a public interactive large display application. Observations from the first study produced a model of collaborative behavior, specifically addressing how users' engagement with the display and collaborator changes throughout the activity. The second study builds on this idea of how users' engagement changes, though in the context of a public interactive large display. Our observations of these users produce a model of bystander behavior, and the information needs of such users. This phase of work, while requiring the development of two large display applications, focused more on the users of the large displays, in particular showing: first, that users' needs change during their use of large display applications; second, that users may each have different information needs of the large display.

Chapter 5 addresses the third research goal, which was to realize a proof-of-concept system that supports transitions. For this phase of work, we designed and built a calendaring tool in the context of an electronic whiteboard application, demonstrating how view changes can support transitions. The tool provides several views, supporting planning and review tasks of calendar scheduling activity. This process involved a paper prototype study, where participants used different views to transition between tasks, illustrating the validity of both the views, and the assertion that these transitions could be supported with different views. We then implement a system prototype to demonstrate the feasibility of implementing view changes with unstructured electronic ink input.

Finally, we address our fourth research goal in Chapter 6 by outlining a design space for supporting transitions with large displays. This design space and vocabulary describe both the different types of transition support (collaborative, temporal, and functional), and different types of mechanisms systems can use to support such transitions (reactive, mixed initiative, and proactive).

Summary of Contributions

We make several contributions in this dissertation, many of which have already been published as peer-reviewed papers as described in the Statement of Co-Authorship.

1. **Development of a new framework for large display interaction focusing on transitions.** This framework suggests a fundamentally different model of interaction for large display application design. This transition framework is informed by four sub-contributions.
    a. **Classification scheme describing activity modes around traditional large surfaces.** By studying how users make use of whiteboards, we uncover the coordinating role of whiteboards in supporting *transitions* in collaborative and independent work, and synchronous and asynchronous activity. These ideas form the basis of a two-axis framework that describes collocated activity involving large surfaces.
    b. **Descriptive classification for the role of surfaces in meeting room collaboration.** From our observations of meeting room collaboration, we develop a descriptive classification of how traditional surfaces are used in collaborative

activity: in particular, we focus on the role that the surfaces play in collaboration (i.e. *how* they are being used).

   c. **Model of collaborative coupling around tabletop displays.** Based on observations of collaborative activity around a tabletop display, we articulate how users fluidly transition between engagement and disengagement with one another's work within the context of tabletop displays. These observations provide designers with a model to build mechanisms that enable transitions between independent and collaborative activity around tabletop displays.
   d. **Design principles to support bystanders around large public displays.** Through the iterative design and deployment of an interactive public display, we show the importance of supporting bystanders' transition between casual bystander to contributor on public displays. We present several guiding principles that aid the design of interactive public displays to encourage public understanding and participation based on a model of bystander behaviour.

2. **A formal operationalization of transitions.** To ground this framework, we develop a formal operationalization of transitions that can be used to describe and discuss transitions using conceptual abstractions. The model serves as a mechanism that designers can use to think about large display application designs.
3. **A system illustrating transitions on large interactive surfaces.** We demonstrate how transitions can be supported in an interactive large display in the context of a calendaring task. We use this system to gain further insight into how transition support can be designed for applications.
4. **A design space for transitions in large display applications.** We bring together the findings from our five studies into a design space that focuses on supporting transitions with large display applications. The design space is a descriptive framework that provides designers and researchers with a vocabulary to describe how transitions can be supported in large display applications.

Transitions: a User's Model, and a User Model

We present two complementary perspectives on transitions: a theoretical user's model as it relates to transitions, and an operationalization of transitions that systems can use to construct a user model. The link between the two models is that the former is an unobservable phenomenon, whereas the latter is the observable outcome of transitions in the user's model. From a system perspective, one can employ these observable aspects to generate and infer a user model that corresponds to the user's model.

Goal, Process, and Task. We begin with the notion of a goal: a desired state of affairs which is likely different from the current state of affairs. When a user has a goal that he or she intends to accomplish, he or she employs a *process* to achieve this goal (Newman & Lamming, 1995). These *processes* are comprised of a set of *tasks*, which are also individually goal-directed. Tasks are often hierarchical in nature, meaning that tasks can generally be broken down into a series of sub-tasks, each of which has a corresponding sub-goal. How these tasks are accomplished may vary: users can usually make choices about to how they accomplish their task goals. At the lowest level, tasks are comprised of a series of immutable *steps*, which are units of decision-free work. These steps are

- Activity: find and send recipe to friend
  - Task 1: find recipe online
    - Task 1.1: launch web browser
      - Step 1.1.1: …
      - …
    - Task 1.2: open internet search portal
    - Task 1.3: search for the recipe
    - Task 1.4: examine search results to find correct recipe
  - Task 2: email recipe to friend
    - Task 2.1: …
    - …

Figure 1.4 A hierarchical task description.

undertaken in a specific order, one after the other, so a task is often just the sequential execution of steps.

We consider the common example of trying to find some piece of information on the Internet (for example, a cooking recipe), and then passing that information on to an interested third party. As depicted in Figure 1.4, this activity breaks down into essentially two tasks: first, find the recipe online, and second, email one's friend. We can further break down the web searching task into a number of sub-tasks: launching the web browser, opening up an internet search portal, searching for the recipe, and examining the search results to find the correct recipe. The launching web browser task may then be broken down into further into a series of mechanical steps. This hierarchical task breakdown is merely an example, of course—other users or designers might break the task down differently depending on the mental model of the system (on the part of the users), or the model of the user (on the part of the designers).

Within the context of HCI, this notion of *task* is important for both design and evaluation: practitioners seek task descriptions to understand what tasks the tools being designed need to support (or be optimized for); when evaluating systems, practitioners employ tasks as tools to measure performance or usability of their systems. In this thesis, we investigate how users accomplish tasks involving large displays or large display surfaces.

User's Model. We observe that a goal is not an observable phenomenon; however, from our conscious experience, we consider our actions to be goal-directed. For the purposes of this dissertation, *a transition is when a user's goal changes*. A brief inspection of the hierarchical *task* description from Figure 1.4 will illustrate the difficulty of separating *task* from *goal*, because tasks are goal-directed, but the distinction is that goals are a desired state of affairs, whereas tasks are typically considered as actionable.

There are at least two issues that are pertinent to this conception of transitions: first, factors that change a user's goals, and second, the problem of magnitude. There are at least two broad sources of factors that change a user's goal: internal and external. Internal factors are those where a user's mental processes lead him/her to change goals (e.g. suddenly remembering that one has a meeting, or perhaps deciding that accomplishing another goal is more important). External factors include changes in the artefacts/tools that are available to be used, or the number of collaborators to work with, and so forth. The second issue pertinent to our discussion here is of magnitude. We consider the task description outlined in Figure 1.4, and in particular, the corresponding goals for each of these tasks. Between Goal 1 (find recipe online) and Goal 2 (email recipe to friend), we may certainly call this a transition. Yet, between Goal 1.1 (launch web browser) and Goal 1.2 (open internet search portal), would this also be considered a transition? After all, at this level, there is certainly a change in goals; however, they are both sub-goals of a higher order goal (Goal 1: find recipe online). For the purpose of this dissertation, we also consider the latter a transition; however, there is evidently the issue of magnitude—some transitions are worthy of support, while others are a matter of course, and do not require specific system support.

When a user's goal changes, there is a trickle-down effect: a user may change the processes he/she is engaged in, and consequently the tasks. Or, a user may change the way in which he/she relates to other collaborators—for instance, changing his/her role in the collaboration from being a listener to a speaker, or a user may add/remove/change artefacts or tools being used to accomplish the task. Notice that these trickle-down effect lead to observable factors. It is these factors that we rely on in a user model of transitions.

User Model. We articulate here an operationalized user model of transitions that serves as a framework that grounds the remainder of this transition. As stated earlier, a user's model of changes in goals cannot be observed; however, the outcome of these goal changes can often be observed. From Figure 1.4, when the user changes from Goal 1 (find recipe online) to Goal 2 (email recipe to friend), the actions that he takes (i.e. his tasks) change in an observable fashion: as he progresses through the different sub-tasks, we can clearly see that the tasks he is engaged in are different.

It is important to note that we derived this user model pragmatically: we were interested in observable factors as they can be used from the perspective of a system attempting to build a user model, and to provide transition support; more to our present purposes however, we needed an operationalization that would allow us to identify *when* transitions were occurring.

We need an additional construct: an *activity state* is an observable state of affairs. Figure 1.5 illustrates a formalization of how we characterize activity state in this dissertation: it is comprised of the *tasks* that are being performed by a *group* of users (or, collaborating

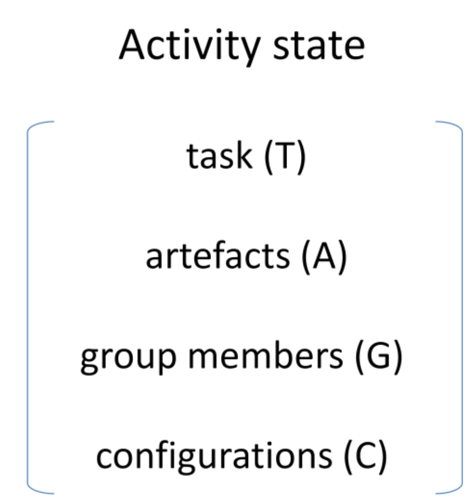

Figure 1.5 A formalization of activity state in our model.

Before (goal: find recipe online)	After (goal: email to friend)
A is the web browser.	A' is A plus an email client.
T is the task of searching for recipes.	T' is the task of composing and sending an email.
C is empty because a single person has no configuration in relation to other users.	C is empty because (again) a single person has no configuration in relation to other users.
G is the singleton user.	G is still the singleton user.

Table 1.1 A simple single-user example of a transition.

individuals) using a set of work *artefacts* (for example, documents, text, etc.), operating with a certain *configuration*. For now, we will assume that the first three of these terms are reasonably self-explanatory—the remaining term, "configuration," deserves some further explication. As outlined in the previous subsection, a user's goal (and by extension a group's goal) may effect certain roles for collaborators: in a pair working with a computer for example, one may be given the role of operating the mouse and keyboard, while the other is responsible for directing the action. These roles are manifest in certain observable configurations of users: in this case, one is holding the mouse and touching the keyboard, while the other is not. Other aspects of these roles may be considered as part of these configurations: who is speaking, whether the group is working together, how are the users positioned with respect to one another, etc. Activity state comprises "configurations," which reflects a configurations of a group of users across the aforementioned range of metrics.

An *activity state* then, is a tuple (A, T, C, G) where A is the set of artefacts in use, T is the set of tasks underway, C is the set of *configurations*, and G is the set of users engaged in the activity. We define a *transition* to a new activity state (A', T', C', G') whenever one or more of the pairs (A, A'), (T, T'), (C, C') or (G, G') is not identical.

Table 1.1 illustrates the transition of our earlier example in terms of our user model.

The transition is a change in the task and the addition of new artefact. After the task is complete, there is a transition back to the original (A,T,C,G) from (A',T',C,G).

A multi-user example of a transition occurs when two people are working together and at some point they decide to work separately. Notice first that there was a change in goal—with perhaps partially in relation to the work, but certainly in relation to the collaborative relationship between the pair. C changes from involving configurations with synchronous collaboration to synchronous independent work. If instead they continue to collaborate, but one leaves and comes back later when the other has left, the change in C is to asynchronous collaboration—a completely different configuration.

This example also serves to illustrate the subtle coupling between some of the factors in activity state as we have operationalized it. Adding a group member to this pair will immediately affect configurations as we have defined it. Similarly, changes in configurations (i.e. working together to working alone or vice versa) changes the low-level, specific *tasks* that collaborators are working on.

Activity Steady States. In our observations of activity and transitions, we tend to observe certain "steady states" where a certain activity state persists for quite some time before transitioning to another. These steady states likely represent local maxima in some metric—efficiency, performance, or collaborative behaviour. A necessary consequence of these local maxima is that there are many activity states that are "unstable" or "unsuitable"—thus, they are not only unlikely to occur, they are likely to not persist for a very long time if they do occur.

Designer's Role. A critical part of a designer's role then is to consider what "next activity states" are likely given a particular scenario. It is usually inappropriate to make this consideration based solely on the model abstraction. Instead, an analysis is best done within the context of the system. A simple example of a design that supports transitions exists on many news sites on the world wide web. Articles on those sites provide links that fast track the process of posting the article to social networking sites, or to email a copy of the article. Such links are an example of a likely change in task (from reading/finding an article to sending it on), and therefore facilitate the transition between these two activity states.

Designers can use this model during scenario construction within the design process for a system. Each transition in an activity-centric design process can be considered as a point where formal support for the transition may be necessary—such points are essentially design choices. In the case of the news article that facilitate posting to social networking sites, various websites make different choices about what other links to provide: print (for transitioning to a different artefact state—paper instead of on-screen), send feedback to the editor (transitioning to a different task), or comment in a forum (transitioning into a conversational task involving a group of asynchronously coordinated individuals).

Some websites do not provide this functionality to post to social networking sites—in many cases, this is the result of an explicit design choice by the designer: on a site reporting on financial data, for example, it may not be prudent or necessary to provide a chat forum because the reading audience is not interested in that transition. Such a site is likely, however, to provide support to transition to a buy/sell activity for the financial products mentioned in the article, or for related financial products.

Magnitude. There should be some notion of the magnitude or scope of a transition: some are more drastic than others. For the case of the hierarchical task description outlined in Figure 1.4, moving between tasks within lower branches is less drastic than a shift in tasks at the top-level. Depending on one's perspective, smaller magnitude transitions may not require formal support compared to large magnitude transitions, which may require coordinating steps; however, smaller magnitude transitions may be much more likely to occur, so if they present a bottleneck, providing smooth support to bridge between different activity states may be beneficial. A one-size-fits-all recommendation is difficult to articulate—it is likely to be context dependent.

Assessing "magnitude" may be difficult to do based on simple metrics: is adding or removing a single member from the participating group of users a large transition? The answer is that it depends. Consider a first year university lecture scenario, where there are 300 undergrads being taught by the professor. If one, two, or a dozen of these undergrads are asked to leave the lecture hall because they are sleeping, it does not affect the presentation given by the professor. Similarly, if one, two, or a dozen members of the public decide to join the lecture (to sit in as audience members), again, nothing is likely to change. As suggested by Gamow (1988), for many problems there is some N that is effectively "infinity" after which it no longer matters if you add a few more, because it's really still the same situation. For a lecture, it is likely that that "N" is somewhere around five to ten audience members—below that the addition or subtraction of a handful of students means the difference between a classroom experience and a gathering of friends. In this particular scenario, not all participants are equivalent, however. If we remove the lecturer, all bets are off—this is definitely a significant "transition."

Supporting Transitions. There is benefit if, when a transition involves changes in the configuration pattern or the group, it can continue to employ the same artefacts to accomplish the tasks. This continuity guarantees a consistent interaction model, which in turn results in an easy-to-learn mechanism. For example, consider writing an email containing a reminder to do something at a later date. Transitioning between the tasks of reminding oneself versus reminding someone else is simple because the mechanism is exactly same (the artefacts are the message and the email application), even if the group and the task are different. Email is a particularly interesting example, because it can make some of transitions transparent, whereas some opacity might be more appropriate. Earlier versions of Microsoft Outlook, for example, made it just as easy to email one's colleagues as it was to email an entire company of thousands of employees. More recent versions of Microsoft Outlook provide confirmation dialogues when attempting to send an email to a large number of recipients.

Summary. We have introduced two perspectives on transitions: a user's model, and an operationalized user model of transitions. In subsequent chapters, we will employ this operationalized model as a way to frame the types of transitions we are observing. Chapter 3, for example, explores transitions between activity states where all three of task, configurations, and group change—and particularly how the whiteboard, as a consistent artefact across these activity states, helps facilitate these transitions. Chapter 4 examines specifically how the configuration—as it relates to physical orientation of a collaborating group on a tabletop—changes as the group's goals change. Finally, Chapter 5 illustrates the design of a calendaring tool that supports transitions between tasks alone.

Dissertation Outline

Chapter 2 provides the research foundations that underlie this dissertation. We discuss prior efforts in designing large display applications, and studies that examine the role of traditional surfaces in collaborative work. This review motivates our interest in the need to bridge between different tasks, and we examine work that has thematically examined the concept of transitions. Together, this survey gives us coverage of existing research upon which our own work is built. Chapter 3 develops the basis for our thinking by cataloging the various roles surfaces play in supporting work. We describe two studies: first, a study of whiteboard use based on survey and interview data, and second, a study of meeting room collaboration. Chapter 4 further demonstrates how users' needs change dynamically in their use of large display applications. We describe two studies that involve such applications: a tabletop application supporting focused collaboration, and a public interactive large display. Chapter 5 realizes a prototype tool that explicitly supports transitions using view changing. We describe the design process that led to this prototype, and demonstrate that view changing using unstructured electronic ink is possible. Chapter 6 presents a design space for supporting transitions in large display applications. We show the utility of this design space by locating existing large display systems within its framework, and discussing areas in the design space that have been underexplored. Chapter 7 presents a summary of the work in this dissertation, and presents a look ahead at the questions that have been raised by the work presented here. We revisit the goals of this dissertation, describing how the work has addressed those goals. Finally, we discuss the implications of this work, and how designers should make use of the research findings in their own efforts to design large display applications.

Chapter 2

Related Work

We begin by broadly surveying large display applications that have been designed to support knowledge work, and also ambient awareness, to lay out the research landscape (Figure 2.1). This review will illustrate that because most researchers building large display applications have (justifiably) narrowed their problem space to support particular activities, they have overlooked the larger context of work that users will expect to do with such displays. We round out this discussion by contrasting the *designs* of these large display applications with findings regarding the *use* of traditional large surfaces. It is from this analysis that our interest in *transitions* arise, because although large display applications are typically well designed for a small focused set of activities, studies of workroom behaviour suggest that users often employ large surfaces for a range of multiple activities.

We then review work in the HCI and CSCW literature that has more broadly explored the idea of bridging between multiple modes of activity. We relate these ideas, in particular, to the concept of collaborative coupling in group work (i.e. the extent to which collaborative partners are engaged with one another). We will see that while the importance of transitions may have been broadly recognized in the prior literature, little has been done in terms of foundational research that identifies the particulars of these transitions, much less realizing support for such transitions.

Large Display Applications

Mark Weiser's articulation of ubiquitous computing (1991), and the PARC group's designs of Tabs (inch-scale displays), Pads (foot-scale displays), and Boards (yard-scale displays) helped define the research space involving large digital displays. This research has generated considerable excitement because large displays allow us to view and interact with data and collaborators in ways that were not possible with standard-sized displays. Within the community of large display researchers, the research programme has been quite varied. We adopt the classification scheme by O'Hara et al. (2003), which conceptually divides the research space: display applications for "knowledge work" (supporting active generation, modification or management of information), and applications supporting ambient awareness information (whose primary role is to provide information in a persistent, ambient fashion).

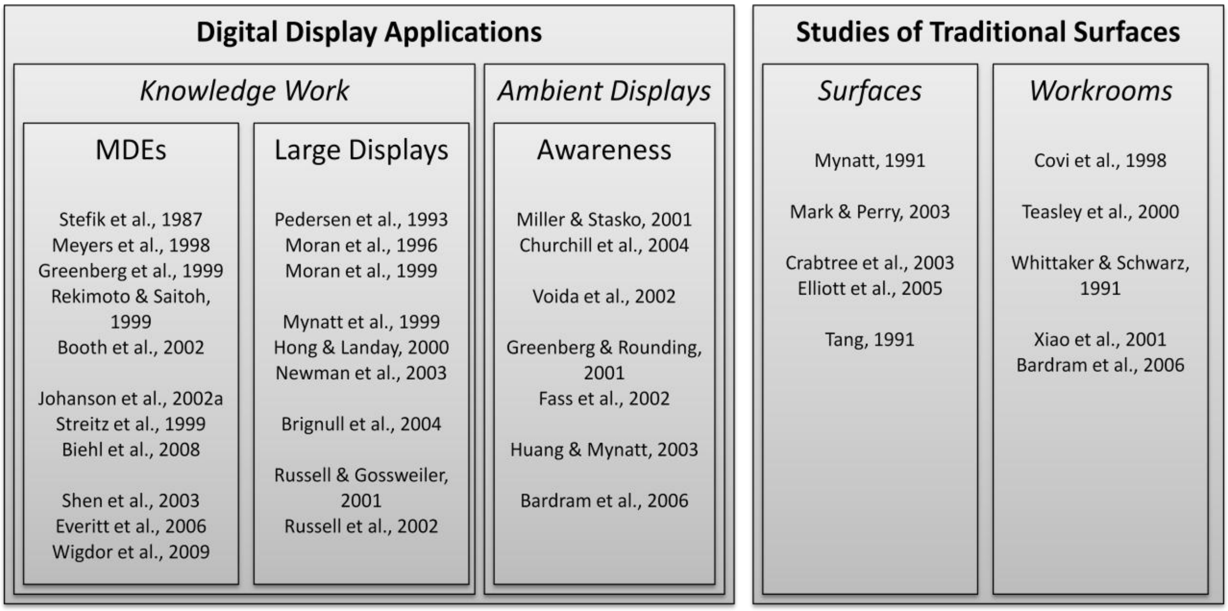

Figure 2.1 A visualization of how previous research that we describe in this chapter fits together.

We contrast studies of these large display systems with research that explores workplace and workroom behaviour as they relate to traditional large displays such as whiteboards. Such studies help to uncover both functions and properties of these surfaces that users rely on to conduct their everyday work.

Figure 2.1 situates the work that we explore in this survey. What we shall see is that in terms of the timing of activity around displays, most technical work has focused on supporting distinct areas of the "temporal" spectrum of activity. The body of work addressing knowledge work, for instance, has focused on supporting collaborative ad hoc or planned use of dedicated display spaces for real-time activity. Researchers building ambient or peripheral displays have instead focused on supporting awareness, coordination or communication between collaborators in an asynchronous fashion. In contrast, we will see that studies of traditional displays have essentially pointed to the use of traditional displays in work across the entire temporal spectrum.

Knowledge Work

Much of the large display research literature can be thematically traced to Colab (Stefik et al., 1987) and Tivoli (Pedersen et al., 1993) from Xerox PARC, where the intent was to explore how large displays could support meeting processes. In these contexts, the large display was used for shared simultaneous viewing of information, or presentation of information. Colab was an multi-display environment (MDE): users were each provided with a terminal connected to a shared, large display. Colab provided simultaneous input to prevent production blocking, and mechanisms to retrieve artefacts and information generated from prior meetings. This work led to development of LiveBoard, one of the original "electronic whiteboard" systems (Elrod et al., 1992), and its companion software Tivoli, which explored pen-based interaction semantics on large displays (Pedersen et al., 1993). Later extensions to Tivoli studied ink-based organization interaction techniques, and integration and interaction with corporate data and documents (Moran et al., 1996; Moran et al., 1999). Thematically, these works introduced several ideas to the research community: (1) centrality of the shared, upright display to the collaborative process; (2) unique and distinct roles for different displays in a multi-display environment; (3) connectedness of data and information flow across displays; (4) the ability to bring in existing artefacts (e.g. documents or applications) for interaction; (5) shared control and interaction with information artefacts; (6) a centrally located, technologically dedicated "computerized meeting room" shared among several groups (though this may be due to pragmatics of deploying such technology), and (7) the transient use of computerized meeting tools.

Multi-Display Environments

Several researchers then began to explore the role of personal digital assistants (PDAs) to interact with large displays, enabling their use as personal input spaces within the broader MDE (Meyers et

Figure 2.2 The Stanford iRoom employed multiple upright and horizontal interactive display surfaces that supported application and input redirection. Users interacted with surfaces primarily through personal devices such as PDAs, mice and keyboards (Johanson et al., 2002a)

Figure 2.3 The i-Land project was a second generation imagining of Stanford's iRoom, incorporating direct-input mechanisms with touch-sensitive surfaces (Streitz et al., 1999).

al., 1998; Greenberg et al., 1999; Rekimoto, 1997). Yet, the optimality of separating display spaces was challenged researchers who suggested designing for continuous display and input spaces which spanned across display boundaries. Augmented Surfaces (Rekimoto & Saitoh, 1999), MightyMouse (Booth et al., 2002), and PointRight (Johanson et al., 2002a) all allowed the focus of control to traverse display boundaries, thereby enabling distributed control (i.e. beyond a single computer) of workspace artefacts in an MDE.

In general, these systems, including the Stanford iRoom project (Figure 2.2 - Johanson et al., 2002b) and i-Land (Figure 2.3 - Streitz et al., 1999) enabled interaction with, sharing (and to some extent creation) of data—primarily media and document artefacts. Beyond moving control and data across displays, however, several researchers began to consider movement of entire *applications* across displays. The ARIS and SEAHorse projects (Biehl et al., 2004) explored these ideas explicitly, providing interfaces for moving entire applications across displays (as opposed to simply document or media artefacts). Application-level mobility became relevant once environments with multiple large upright displays became a reality, because it necessarily de-emphasizes the role of each divided shared upright display, in effect making them functionally equivalent.

To better support collaborative dynamics in a collocated setting, several researchers began exploring the use of tabletop displays. A considerable body of recent work has investigated the use of tabletops as a means to support collaborative activity (e.g. Shen et al., 2003; Tang et al., 2006; Isenberg et al., 2008; Morris et al., 2004; Morris et al., 2010; Scott et al., 2004; Ryall et al., 2004), because the horizontal arrangement facilitates a more intimate seating arrangement (Dietz & Leigh, 2001). Because researchers in this space have focused on the design of the tabletop interface as the locus of interaction, the upright display is often relegated to a "shared viewing" role. The UbiTable

(Shen et al., 2003) and later Multi-Space (Everitt et al., 2006) and WeSpace (Wigdor et al., 2009) environments, in particular, advocated this view. In so doing, these researchers also returned to the notion of providing separate roles and interaction types for different displays (i.e. how one interacts with the laptop in this space is different from how one interacts with documents on the shared table space). Furthermore, these designs also highlighted the use of shared displays for transient, walk-up use by emphasizing implicit login processes.

Large Displays for Knowledge Work

Researchers have built and studied many single large-display systems that operate more or less independently. We sample a few representative systems here:

Figure 2.4 illustrates IBM's BlueBoard project, which explored the design and deployment of a shared upright display in an *informal,* public, transient space rather than a closed, dedicated collaboration space (Russell & Gossweiler, 2001; Russell et al., 2002). BlueBoard's design necessarily emphasized walk-up-and-use usability, and enabled generation of ink-based artefacts, as well as shared (multi-user) exploration of documents and artefacts such as the web or documents from the user's personal drive space.

Flatland explored ink-based interaction semantics for personal electronic whiteboard systems (Mynatt et al., 1999). In particular, Flatland introduced the notion of augmenting ink with computation by interpreting the ink strokes based on "mini applications." Such interpretation

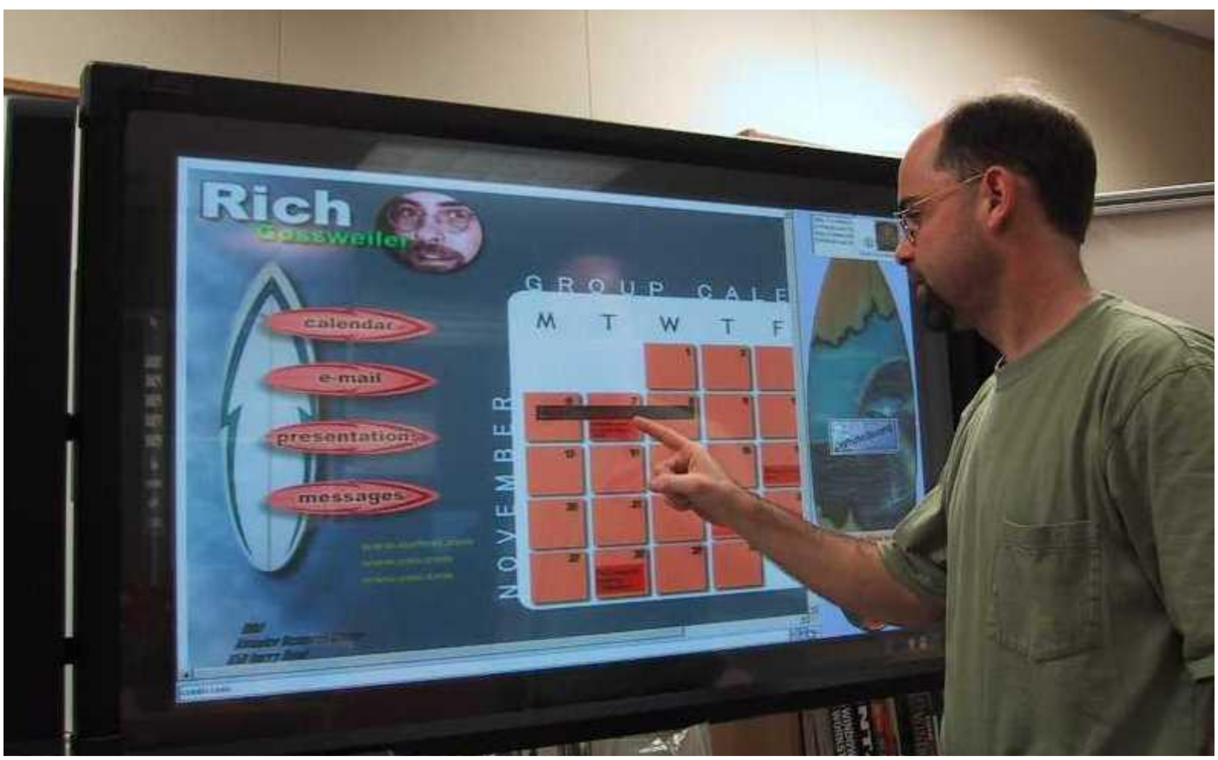

Figure 2.4 The IBM BlueBoard provided users with rapid access to personal information in a variety of communal applications for ad hoc collaboration. (Russell & Gossweiler, 2001)

Figure 2.5 The Dynamo project was deployed as a persistent display where users could collaborate, and leave parcels of media that could be retrieved at a later time.

engines included a simple arithmetic engine, a stroke beautifier, and a list-organizer sub-system. SATIN emerged soon after as a pluggable architecture for the construction of ink-based applications (Hong & Landay, 2000). It provided flexible, generic support for independent engines for strokes, gestures and views. The primary consumer of this toolkit was the DENIM system, a large display application to support visual, architectural design of websites (Newman et al., 2003). The DENIM system was primarily used as a high level prototyping tool during planning meetings.

Figure 2.5 illustrates the Dynamo project, which stands somewhat distinctly in this space due to its use and deployment as a single, dedicated application in a semi-public break room (Brignull et al., 2004). The project dispensed with notions of multiple displays, focusing instead on the ability to move, view and share media artefacts across devices, utilizing the large display as the primary display for interaction (via remote control keyboards and mice). Where the deployment of this system differs from the work outlined above is that the system enabled shared artefacts to *persist* on the display across time. This feature was increasingly used only toward the end of the deployment, underscoring the realities of how work practices develop—they require time to emerge, as users come to understand and appropriate features. This distinction highlights a core theme across these works, which is that the systems, affordances and mechanisms were primarily designed for synchronous, real-time collaboration. Specifically, the design goals of these systems were to support more fluid transfer of ideas, information and control—but in a physically and temporally collocated sense.

Ambient Information Systems and Awareness Displays

In contrast to the large display systems described earlier that were built for synchronous activity, other researchers have been investigating the design of peripheral, or ambient displays. These

Figure 2.6 Plasma Poster was an ambient display surface relying on user-submitted content. (Image source: http://www.flickr.com/photos/xeeliz/)

systems have several unifying behavioural features (Pousman & Stasko, 2006): (1) they display information in the periphery of one's attention rather than in the focus; (2) they enable users to move between peripheral awareness and focused attention and back; (3) they provide subtle, ambient notification of state changes; (4) they display important, but non-critical information; (5) they are tangible and located in the environment, and (6) they are aesthetically pleasing. Of these systems, we are primarily interested in systems that employ a large display (as opposed to those that employ tangible or environmental artefacts), as they provide a point of comparison and contrast from the large display applications discussed earlier.

InfoCanvas is the canonical example of an ambient, peripheral display (Miller & Stasko, 2001). The display was designed and deployed as a picture frame of a cartoon-like scene. Entities in this scene were linked to digital information sources (e.g. weather, stock prices), and updated unobtrusively without user intervention. Thus, the display could be placed anywhere in the user's environment, to provide an up-to-date, at-a-glance, consolidated information source. The deployment of this system made it clear that users' mental model of the ambient display were distinct from the PC (usually a user's work station) driving it. This separation was due in part to the physical separation from the PC (the picture frames were often hung away from the desk, or positioned on bookshelves).

Figure 2.7 The Notification Collage is an ambient display where colleagues comprising a small community post media elements onto a real-time collaborative surface that all members can see. (Greenberg & Rounding, 2001)

Figure 2.6 shows the Plasma Poster network (deployed at FX PAL) which adopted this ambient display model, but was instead geared toward providing users with awareness of the activities of other members of the organization (Churchill et al., 2004). Plasma Posters relied on user-submitted content (images and text) rather than on information from the web (as in InfoCanvas). Kimura was similarly designed as an ambient display, but instead provided an individual information worker with contextually relevant awareness of information pertinent to his/her ongoing tasks (Voida et al., 2002). Thus, related information artefacts could be spatially grouped on a display, and the system could provide notification of events—for example, a collaborator arriving to work, or a collaborator's response to an email. These displays were again designed primarily for asynchronous use, providing users with awareness by displaying information; further, these examples show how researchers became interested in using displays for supporting awareness of people.

The Notification Collage (Greenberg & Rounding, 2001) and MessyBoard (Fass et al., 2002) systems further advanced the use of ambient displays for awareness and communication between individuals. As illustrated in Figure 2.7, the Notification Collage provided support for desktop web cameras, thereby allowing users to see and maintain awareness of collaborators' presence and activities. Both systems supported the "live" addition of textual content, which could then be left as persistent notices on the display. In some sense then, these systems blended the ideas from media space research and placed them on a shared, semi-public ambient display. Beyond simply acting as an awareness display, these systems provided mechanisms to interact and communicate in real-time.

The Semi-Public Displays project focused on providing coordinating mechanisms on a semi-public ambient display, providing presence awareness information as well as an explicit tool for planning attendance to social functions (Huang & Mynatt, 2003). AwareMedia was designed and deployed as an interactive, ambient display for a hospital setting (Bardram et al., 2006). The system provided location (spatial) and schedule (temporal) awareness of medical personnel and operating rooms, allowing users to smoothly coordinate and manage interruptions and communication. The deployment focused on providing a "simple, stable, and predictable display," meaning that users would be able to depend on the display as a persistent awareness resource. At the same time, they could use the system to communicate with other doctors in real-time using a chat interface.

Thematically, we have seen that awareness displays are based on simple core principles derived from work with ambient information systems. Functionally, many of the awareness displays we have discussed are deceptively simple—users cannot use them to create sophisticated documents or other digital artefacts. Yet, it is this simplicity that makes them useful as ambient, awareness displays—they can rest in the periphery of one's attention. We have also seen that they are increasingly being designed to provide awareness of collaborators, along with explicit support for real-time communication, thereby allowing collaborators to transition into real-time collaboration when desired.

Studies of Traditional Displays and MDEs

In parallel, several researchers have studied work practice of collaborating groups—many with an explicit interest in the use of traditional display surfaces (e.g. non-digital whiteboards) and how information artefacts support and coordinate collaborative activity. Although such display artefacts perhaps have well-understood affordances from our day-to-day interactions with them, what becomes clear from these studies is that many of these affordances have been lost in digital display designs due to our implicit assumptions about the nature of work that takes place around these displays, and how this work actually takes place. It thus serves to re-sensitize ourselves to these concepts, as a deep understanding of these concepts brings these assumptions to light.

Everyday Traditional Surfaces

The common office whiteboard has been an oft-studied artefact—largely due to its ubiquity, flexibility and utility. Mynatt's study of personal office whiteboards revealed their use for multiple parallel pre-production tasks (Mynatt, 1991). Activities involving the whiteboard spanned a range of temporal ranges: using the whiteboard for thinking activities for instance would be synchronous activities, while using the whiteboard for quick capture or reminders are asynchronous uses of the whiteboard. Mynatt also highlighted the spatial organization of information on whiteboards themselves, showing that marks were typically organized into clusters called "segments"—some of which were persistent long-term, with other regions being "hotspots" of constant, transient activity.

Finally, Mynatt points to the use of the whiteboard as being highly contextualized—meaning that information on many whiteboards are only useful or relevant in the place where it is located.

Mark & Perry (2003) expand on this view of whiteboards, studying their general use in the workplace. Their findings emphasize whiteboards' role in supporting visually persistent artefacts for quick reference (for prospective and retrospective reminding), and as coordinating devices. For instance, they are used for awareness and communication: informing others about one's absence on a schedule board, or as a conversational resource. Further, they are also used for explicit coordination in planning and developing schedules. Together, these findings underscore the role of self-evident properties of whiteboards: they are typically physically fixed (information is immobile), and information on these displays is persistent until explicit action is taken to remove it (information is persistent and visually available). Without these properties, the activities described above would not occur as they do.

An interesting departure from this artefact-centric view of traditional displays is the focus by Crabtree et al. (2003) on the term "display" in its verb form—that is, as an action rather than as a noun. This articulation brings focus to the *intentionality* of the act, rather than an actual physical, tangible artefact. In so doing, we see that many things in our environment are actually "displays" because someone arranged objects or artefacts in such a way that they would be "seen" at a later time (Crabtree et al. (2003) provide the example of a post letter that has been placed on one's dinner placemat). In this sense, they were "displayed." Elliott et al. (2005) explore the construction and use of displays in the home, articulating so-called "contextual locations" that home inhabitants use to maintain awareness of each others' activities, and to coordinate activity. This work shows how locations in the home are evolved to become important due to the flow of inhabitants through the home and their activities. Thus, in the traditional environment, *location* is of central thematic importance. Further, this line of thinking points to an interpretation of physical "displays" in a traditional environment as being spatially demarcated locations where information is placed to be displayed and to be seen.

Traditional Surfaces in the Workplace

Focused investigations of what we consider to be "traditional" multi-display environments again emphasize the role of location and spatiality (at different levels of granularity), as well as the utility of storage and history displays (Teasley et al., 2000; Covi et al., 1998). Studies of teams' use of dedicated project rooms (where teams operated primarily in a single room for the duration of a project) reveal fundamentally different uses of traditional displays than what is typical in meeting rooms. For instance, these teams relied on the use of paper flip-charts, where charts could be ripped off the pad, and placed around the room as a permanent record of information for ready-reference. This visible, permanent record of activity then became physically located in the room, thereby allowing cognitive associations to be formed between the spatial areas of the room or walls and the information itself. In many cases, information chosen for this type of permanent display was created with the intention of displaying it (e.g. task lists, schedules, etc.), which is consistent with the notion that the artefacts created during problem solving are really boundary objects (Star & Griesemer, 1989)—artefacts that are only useful for conveying concepts. Nevertheless, some of these boundary objects were still sometimes posted. These work practices are somewhat peculiar

Knowledge Work	Ambient Display
• Design work • Focused collaboration • Artefact creation/modification • Presentation	• Awareness • Asynchronous communication (notices) • Coordination (scheduling) • Status

Table 2.1 Types of tasks supported by large display applications in O'Hara et al. (2003)'s classification.

given our common experience with meeting rooms; however, they show that in dedicated project rooms, where it is acceptable to post things on walls, collaborators will employ persistence and spatiality to retain information in their environment.

Whittaker & Schwarz (1991) provide a compelling example of how the "locatedness" of a traditional display and the information on it (in this case, a schedule) enabled coordinative awareness and communicative action beyond the information contained on the display itself. The comparative study (traditional display vs. an online schedule) found that users of the traditional display took greater efforts to ensure the schedule was accurate and current, and also took more steps to ensure that it would be accurate in a prospective way. Further, it facilitated group processes, such as awareness of when the display was being updated (when someone walked toward the location, there was a chance something was to be changed on the schedule). Thus, with traditional displays used in a persistent way, the information and coordinative function of the display becomes bound with the location.

The utility of this type of contextual location has also been observed in nurses' use of whiteboards in hospital wards (Xiao et al., 2001; Bardram et al., 2006). In this context, a single whiteboard often acts as a central, coordinating resource for dozens of nurses, doctors and technicians. This whiteboard is commonly located in a central, well-trafficked area, and is maintained at all times, providing an up-to-the-minute awareness display of the location of patients, their state, and doctor and nurse assignments. The specific means through which this is accomplished varies from situation to situation, with each group adopting its own conventions; however, the purpose of such coordinating whiteboards is surprisingly common. The importance of the coordinating role of a display cannot be understated (Bardram et al., 2006). These displays are fixed function—they are *dedicated* to a specific set of tasks. Further, the inhabitants of the space know the role of the whiteboard and its function, and do not manipulate its contents inappropriately, thereby enhancing the visual persistence of the display. Again, we notice that the function of these displays relies on the deceptively obvious properties of traditional displays: they are spatially fixed, visually persistent, and in this case, dedicated to a fixed function.

Summary of Prior Work on Large Displays and Surfaces

We have reviewed representative research from three related threads of inquiry in the CSCW literature regarding large displays: environments and large displays for knowledge work, ambient displays, and studies of traditional work surfaces and work environments. Table 2.1 summarizes the types of tasks supported by the large display applications based on the classification by O'Hara

	Independent	Collaborative
Sync	*Worker* • Word processor • Spreadsheet • CAD software	*Real-time interaction* • Telephone • Video conferencing • Instant messaging
Async	*Personal Management* • PIM, schedule, agenda, task list • Reminders, post-it notes	*Ongoing tasks* • Team rooms • Bulletin boards • Email

Table 2.2 A modified groupware matrix that emphasizes modes of activity, and the tools that support them.

et al. (2003). Broadly, most applications were built *either* to support knowledge work or ambient display of awareness information.

In contrast, traditional surfaces such as whiteboards are used for multiple activities, and to some extent, simultaneously: meeting room whiteboards support the generation and brainstorming of ideas, yet at the same time, they can also provide awareness information, such as a schedule (e.g. Whittaker & Schwarz, 1999). Although large display applications provide added *depth* and richness in support for specific tasks, they fail to provide the same *breadth* of task support as whiteboards. This is likely less of an oversight, and more likely a deliberate choice to narrow the scope of given research projects. Nevertheless, this mismatch should be a core concern. It is an example of the gap between users' existing work practice and the design of most large display applications.

Our review suggests that traditional surfaces are being used as a bridging mechanism between multiple tasks. In Chapter 3, a study of whiteboard use confirms that indeed, whiteboards are used for a variety of activities, and that they are used to facilitate transitions between these activities. While our particular context of study (large displays) brings a new perspective to this idea, we have seen that conceptually, researchers have considered this issue in prior work. In the next section, we explore how researchers have addressed transitions in groupware more generally—a concept that has sometimes been called "seamlessness."

Modes of Activity and Seamlessness

The standard groupware matrix's primary axes, same/different time and same/different place, define four modes of activity (Johanson, 1988). Of these, the vast majority of both research and commercial groupware tools have been primarily designed to support geographically distributed workers. Yet in the case of situated large displays, the "place" of activity is fixed—users are collocated. As a theoretical grounding, we have found that Thimbleby et al. (1990)'s articulation of *reflexive-CSCW* (the use of CSCW systems for personal or independent work) to be most pertinent to our interests in the present study. Table 2.2 illustrates this conceptualization.

Thimbleby et al. (1990)'s definition arises from the observation that groupware has sometimes been appropriated for *independent activity* (e.g. Whittaker & Schwarz, 1999): for instance, when a user sends email to himself for reading at a later date (i.e. asynchronously), he is performing *personal management*, coordinating his activities by setting a task list or reminder (Cockburn & Thimbleby, 1991). Thimbleby et al. (1990) argue that researchers should explore how to support reflexive-CSCW: with well-designed mechanisms, users can rely on existing or known practices to smoothly *transition* between independent and collaborative activity (e.g. Cockburn & Thimbleby, 1991; Geyer et al., 2003; Greenberg & Roseman, 2003). Existing email systems are a good example of such a tool because they do not distinguish between independent and collaborative use: people can apply the same mechanism to asynchronously communicate with others or themselves (Whittaker & Sidner, 1996).

Some systems have realized the inverse: enabling users to transition from independent to collaborative activity while preserving existing work practices. For instance, TeamWorkStation (Ishii, 1990) and ClearBoard (Ishii et al., 1992) overlay a video image of a remote collaborator's drawing workspace on one's own, thereby fusing the two workspaces and allowing collaborators to interact with the workspace as if they were working independently while facilitating joint activity. These systems illustrate a design philosophy termed as "seamlessness" that enables users to smoothly transition from independent work to collaborative work.

TeamRooms (Greenberg & Roseman, 2003) introduces a "room" metaphor for shared visual workspace groupware, where people can enter and leave rooms at any time, and rooms can be populated by persistent groupware artefacts (see also Geyer et al., 2003). This simple model enables all quadrants of Table 2.2, where transitions are afforded by how people use the room rather than by explicit technical means. If two or more people are in a room at the same time, they are doing synchronous collaborative work. If one leaves an artefact for another to view later, it is asynchronous collaboration. If a room is used by an individual who works a bit, then leaves and comes back to continue where they left off, it covers both modes of individual work.

Finally, some systems offer a form of data-centric transition, where digital artefacts may be transferred to a shared system, allowing users to share work completed in an independent fashion in a collaborative space (e.g. Geyer et al., 2003). For instance, MessyBoard (Fass et al., 2002) and Notification Collage (Greenberg & Rounding, 2001) allow users to post information from their personal PCs to the shared display. Yet, like TeamRooms, both systems were designed for distributed work, where the interaction capabilities of the client (used for posting information) are different from the shared display (which primarily acts as shared context for conversation); here, we are interested in how these transitions can be afforded by a situated large display.

Collaborative Coupling

Researchers have typically described tasks and group work as being "tightly" or "loosely" coupled (e.g. Baker et al., 2002; Gutwin & Greenberg, 2002; Gutwin & Greenberg, 2004; Morris et al., 2004; Salvador et al., 1996; Scott et al., 2003; Tse et al., 2004). Very generally, coupling refers to the dependency of users on one another—when users cannot do much work before having to interact, the work is tightly coupled; conversely, when users can work independently for long periods of time, the work is loosely coupled (Salvador et al., 1996). Considerable evidence suggests that in both collocated and distributed shared workspaces, group activities cannot be neatly dichotomized into "independent" and "shared" activity. *Mixed-focus collaboration* is a term used to describe this type of activity, which involves both independent and shared tasks (Gutwin & Greenberg, 1998).

We use *collaborative coupling* to refer specifically to *the manner in which collaborators are involved and occupied with each other's work*. While researchers recognize the value of using collaborative coupling as a way of describing group activity, we do not have a systematic understanding of mixed-focus collaboration beyond recognizing the end points: individual work and shared work. Most computer interfaces are designed for one end point or the other. Yet, individuals do not instantaneously shift between independent work and group work. Instead, as we will see in Chapter 4, a group's *collaborative coupling style*, or the manner in which collaborators are involved and occupied with each other's work, frequently changes during work (Baker et al., 2002; Salvador et al., 1996). For instance, an individual might work on an idea alone before presenting it to the group, and then later work with the group to jointly manipulate the idea (Scott et al., 2004; Tang, 1991).

The conceptualization of activity in Table 2.2 actually shows the two end points of independent and collaborative work. A more systematic model of the space between these endpoints would provide a deeper understanding from which to base design. In Chapter 4, we explore collaborative coupling within the context of collaborative activity over a digital tabletop, articulating a classification of collaborative coupling styles. This model illustrates the changing information and interaction needs of users, providing a more nuanced perspective on the types of transitions that should be supported in large display applications.

Chapter Summary

In this chapter, we have provided a survey of prior work involving large displays. This survey covered the design and study of applications for knowledge work and as ambient displays, and provided treatment of studies that investigated traditional contexts to inform large display application design. Our survey showed that while most applications have typically been designed for a fixed subset of activities and tasks, traditional surfaces are often used across these task spaces and activities. We then discussed prior literature that has also considered transitions between different modes of activity: first in the context of distributed groupware, and then in the context of collaborative coupling in mixed-focus collaboration.

The review in this chapter motivates and sets the scene for the studies that we describe in the next two chapters. In Chapter 3, we focus on the question of how traditional large surfaces are used for multiple tasks, and how they support transitions between those tasks. Chapter 4 develops models of how users' needs change during their use of large display applications. The findings from these studies form the basis for a new framework for large display application design emphasizing the transitions between different tasks and changing user needs.

Chapter 3

Multiple Roles of Traditional Large Surfaces in Work

In the previous chapter, we outlined how researchers had explored the design of large display applications to support collaborative work (e.g. Russell et al., 2002; Pedersen et al., 1993; Greenberg & Rounding, 2001; Churchill et al., 2004; Snowdon & Grasso, 2002). Most of these designs, like desktop applications, support a specific set of tasks as they relate to an activity. In contrast, traditional large surfaces such as whiteboards are used for a host of activities (e.g. Xiao et al., 2001). What are users' models of these large surfaces, and are these consistent with how we are designing large display applications?

A user's model of a tool (e.g. large surface) pertains mainly to his or her understanding of how that tool fits into his work, his expectations of the tool, and his goals in using the tool. HCI researchers gain insight into this model with both direct (e.g. asking users), and indirect methods (e.g. studying users' work practice with the tools). These methods complement one another: direct methods give insight into conscious expectations, experiences and goals, while indirect methods provide evidence that reflect unconscious (or assumed) expectations and goals.

In this chapter, we describe two studies that we conducted with the goal of classifying the roles that traditional large surfaces take on in work—reflecting part of a users' model of large surfaces. The first study examines whiteboard use through survey and contextual interview methods. It shows that users employ whiteboards for multiple tasks and activities (Tang et al., 2009). It further shows how spatial partitioning facilitates transitions between activities, while ink-based semantic encodings particular to those activities support transitions between tasks.

The second study explores meeting room collaboration using a naturalistic observation technique. It examined collaborative activities involving traditional surfaces, and how the surfaces supported collaboration. A descriptive classification was generated that describes *activities* surfaces are involved in, and the *roles* surfaces play in those activities.

Together, these studies of traditional activity suggest that the desktop application model may not be consistent with users' mental models of large surfaces. While the desktop model predominantly relies on a single application focus, users' behaviours around large surfaces reflect a model where they fulfill multiple roles simultaneously.

Section 1: A Study of Transitions in Everyday Whiteboard Activity

Prior studies of whiteboards have explored their use as a collaborative medium (Teasley et al., 2000; Xiao et al., 2001; Covi et al., 1998) and in personal office spaces (Mynatt, 1999; Perry & O'Hara, 2003), examining how the artefacts on such whiteboards support work. The present study results in a detailed, focused examination of how whiteboards support *transitions between different related tasks and modes of activity*—effectively, low-overhead task switching To begin, our findings will show that most whiteboard tasks can be categorized into a simple 2×2 matrix (we introduced this matrix in Table 2.2), where the primary axes are synchronous *vs*. asynchronous work (i.e. same/different time) and independent *vs*. collaborative work.

In considering many of the large display applications outlined in Chapter 2, we see that most applications support activities only within a single quadrant of this matrix. In contrast, we have seen that traditional whiteboards enable a broad set of activities from multiple quadrants. It was this flexibility that motivated our current study: what affordances of whiteboards enable their use for these activities? Further, what work practices do users develop to take advantage of these affordances? To address these questions, we designed a study that would allow us to gain a broad-spectrum understanding of users' activities involving whiteboards, as well as in-depth understanding of how users employed these whiteboards to support their tasks.

Our analysis will show how the affordances of everyday situated whiteboards support transitions across different modes of work and tasks. The purpose of this chapter is to articulate the nature of these transitions (as they occur in the traditional whiteboard context), and we illustrate, through a number of examples drawn from our study, the role that the whiteboard plays in supporting these transitions. The study underscores several basic but important affordances of whiteboards that support this practice, including visual persistence, flexibility of interaction primitives, and their situated physicality.

Study Method

Our study began with a broad base survey, targeting people who regularly used their whiteboard—defined as "at least once a week" in our advertisement (we felt this would avoid being unduly influenced by incidental users). The purpose of the survey was to gain a broad understanding of users' whiteboard activities as a basis for our classification scheme. Based on these ideas, we developed questions that were asked during a set of in-situ interviews with "heavy" whiteboard users to understand how their practices with whiteboards fit into the context of their work and workspaces.

Survey. We deployed a web-based survey using a snowball recruitment sampling technique, beginning with email ads posted on computer science and engineering graduate student mailing lists. This technique asks that recipients complete the survey if they meet the requirements, and to forward the email to correspondents that might fit the requirements. Ultimately, the reach of our survey encompassed primarily industry (i.e. non-academic) users with a wide variety of backgrounds: graphic artists, software designers, engineers, business analysts and communications specialists. We received 167 survey responses, of which we discarded 32 due to incompleteness: we therefore report on 135 complete responses.

Participants were entered in a draw for a $100 prize. The survey consisted of 53 items, asking them about their whiteboard behaviour: What activities did they engage in (derived from Mynatt, 1999), and how frequently? Were these activities independent or collaborative? We asked users about two whiteboards important to them, where the whiteboards were located, what they were used for, what was currently on them, and who else used the whiteboard.

In-situ interviews. To add further richness to our understanding, we conducted in-situ interviews with 11 users (3 females) selected from our survey pool. These users were self-identified "heavy" whiteboard users, and we selected them primarily based on geographic convenience, but also aimed for a broad variety of occupations (including academics, managers and engineers). Of these, two participants were selected from the United States for interviews as a check against geographic bias. This general approach of using in-situ interviews is consistent with prior work investigating the role of traditional large surfaces (e.g. Mynatt, 1999; Perry & O'Hara, 2003).

The one hour interviews were conducted in front of each user's "most important" whiteboard, and were audio recorded for transcription (interviews with the two international participants were conducted using the phone and with the aid of digital photos of their whiteboards). We also collected photographs of users' whiteboards and their physical context, and used the whiteboard as a grounding artefact for discussion. We developed a list of questions based on theme areas from the survey results, but allowed the flow of the interview to guide the dialogue, referring to the list only to ensure that all themes had been addressed. Interview participants were given $20 remuneration.

Interview Analysis. We conducted an inductive analysis of interview data, using an iterative open coding technique on the interview transcripts to categorize the roles of surfaces, and to derive a thematic understanding of our participants' activities. This analysis process, based on the work of Strauss & Corbin (1999) ensured that our hypotheses and findings were grounded in actual data rather than our assumptions about users' practices.

Methodological Justification. We employed a survey to establish a general basis on several fronts: to understand general whiteboard activity, to understand how these users appropriated whiteboards, and to learn about the whiteboards themselves. We relied on a snowball sampling method to achieve a high number of respondents. This technique typically results in a sample that is not necessarily representative of the population. However, because we were interested only in developing a classification of the types of activities (and not precise numbers on the prevalence of these activities in the population), this sampling technique was appropriate as it ensured a broad reach of users.

The in-situ interviews helped substantiate our ideas about whiteboard use given the questionnaire data. Our choice to use "heavy" whiteboard users was to capitalize on the "expertise" of these users.

User type (n)	Uses per week	Boards per week	Boards per month	# Important boards
Heavy (22)	8.4	2.9	5.2	2
Medium (69)	4.7	2.4	3.9	2
Light (43)	2.5	1.5	2.7	1

Table 3.1 Relationship between mean self-rated frequency of use to uses/week, mean boards used per week and month, and median number of whiteboards considered "important."

Location	#1 (n=129)	#2 (n=110)	Total
Home	19	15	34
Work/Personal	54	14	68
Work/Shared	51	63	115
Work/Coworker	3	11	14
Public/Other	2	6	8

Table 3.2 Location of users' most important (#1) and second most important (#2) whiteboards. Notice half the whiteboards (shaded) are used in fairly personal spaces.

Given that they had chosen to make use of whiteboards as an integral part of their work process, we were likely to see these users appropriate whiteboards in rich and interesting ways compared to only casual users.

The number of participants we interviewed was not pre-determined. Instead, we conducted interviews until we felt we had exhausted the diversity of uses and were no longer learning anything new. We employed this method to ensure we would not artificially limit our study—there would have been no rational *a priori* reason for deciding to interview two or eight participants rather than ten.

Findings

We begin by describing our survey data, which frames everyday whiteboard environments, and whiteboard use within the four-quadrant model introduced earlier (Table 2.2). Drawing on our contextual interviews, we then discuss how several users appropriated a whiteboard in ways that allowed them to transition between multiple tasks and modes of activity. We then show how the physical context and social practices around situated whiteboards support this practice in general, and then further illustrate the importance of the location of the whiteboard as a "place of work" beyond its function as a sketching device.

Whiteboard use. Table 3.3shows how this self-rated frequency of use (as "heavy", "medium", or "light" users) relates to reported use of whiteboards in terms of usage frequency, and the number of perceived "important" whiteboards. In spite of relatively large differences in frequency of use, users tend to only use a small number of whiteboards overall. Table 3.3 compares self-rated frequency of use to a variety of tasks, showing that heavy users appropriate the whiteboard more broadly for independent use than do light users.

We asked each user to report in detail on up to two whiteboards that were "most important" to him or her, including information about where the whiteboard was located, who used the whiteboard, the number of segments on the whiteboard (Mynatt, 1999), the age of the content on the whiteboards, and so forth. We collected data on 239 such whiteboards. Table 3.2shows the location of these boards. Of particular note is that while many of these whiteboards are located in

Independent	Brainstorm	Task list	Reminder	Storage	Other
Heavy	4	4	4	3	2
Medium	3	3	3	3	0.5
Light	2	1	1	2	0

Collaborative	Brainstorm	Conveying Ideas	Task list	Reminder	Storage	Other
Heavy	4	5	3	2	3	4
Medium	4	4	3	2	2	0
Light	3	3	1	2	2	0

Table 3.3 Relationship between users' self-rated frequency of whiteboard use, and median rating for frequency of whiteboard tasks (6 pt Likert scale: 0=Never, 1=Very rarely, 2=Rarely, 3=Occasionally, 4=Frequently, 5=Very frequently).

"collaborative" contexts (i.e. shared workplace area), about half (shaded) are located in personal spaces.

Characterizing Whiteboard Tasks

To corroborate users' reports on the frequency of whiteboard use for various activities, the survey asked users to immediately examine and report on the content of their whiteboards, describing what the content was for. While about half the descriptions lacked necessary detail (i.e. only describing content without intent), it was possible to characterize whiteboard content from about half of the whiteboards (n=122) along the dimensions introduced earlier (independent *vs.* collaborative, and synchronous *vs.* asynchronous). Table 3.4 provides examples of our classification which we elaborate on next. Strikingly, over half of the whiteboards contained deliberately un-

	Independent	*Collaborative*
Sync	**Thinking** • "Flow (boxes and arrows) of a presentation I am about to give" • "A mind map of my current largest project"	**Communication** • "Two different design diagrams, drawn by me to illustrate points for coworkers" • "I need to be able to convey ideas and brainstorm with other faculty and students"
Async	**Planning** • "Six different to-do lists, for each project I'm working on, and several small post-it notes with ideas or sketches I don't want to forget, stuck to the to-do list for that project" • "Project milestones and the different modules that need to go into the game for each milestone"	**Coordination** • "All active projects and their schedules" • "Action items (tasks for team members) from the team meeting"

Table 3.4 Examples of user reported whiteboard contents, classified in our modified groupware matrix.

erased persistent content for later, asynchronous use. While this data gives us evidence of the way the information was used in one case, we cannot, for example, know whether data used for independent synchronous activity was used in a later episode.

Independent synchronous (15% of whiteboards contained *remnants* of this content type): These activities involved a person making use of the whiteboard to help him or her think in some way. The primary value of this activity was at the time of creation, where it helped the user address a problem in the immediate term. Examples included working out problems visually, organizing information spatially, or simply using it as a "large writing surface."

Independent asynchronous (61% of whiteboards contained *remnants* of this content type): This type of activity involved a user deliberately putting or leaving information on the whiteboard with the intent of using it at a later time for his/her own use. This information was used for planning types of tasks, for example to help the user recover context, or to remind the user about something. Examples included task lists, notes, reminders, and reference sketches.

Collaborative synchronous (30% of whiteboards contained *remnants* of this content type): These activities involved groups of users employing the whiteboard to accomplish a task, for example to communicate information, or to work out ideas. Examples included brainstorming, collaborative design, or presenting ideas.

Collaborative asynchronous (26% of whiteboards contained *remnants* of this content type): When users deliberately placed information on the whiteboard with the intent of others either seeing or re-engaging with it at a later time or in an ongoing basis, we labeled the activity as collaborative asynchronous. Examples included collaborative task lists, schedule boards, action lists, etc.

Consistent with Mynatt's reports of the use of segmentation to delineate different activities (Mynatt, 1999), our data also suggest that spatial partitioning was often used to separate activities. As evidenced above, many whiteboards had content remnants intended for multiple different activities.

Beyond our categorical definitions, however, it became clear that for some participants, this four-quadrant view insufficiently represented their use of the whiteboard. In many cases, their use of the whiteboard content transcended our conceptual boundaries, suggesting that the work artefacts supported users' transitions between modes of activity. For instance, *"Ongoing project sketch/notes",* suggested both that the content was being used as reference for asynchronous activity, and for ongoing thinking.

Using Whiteboard Artefacts to Transition across Modes

Our in-situ interviews were designed to more deeply understand this phenomenon: if our four-quadrant view was insufficient for classifying some whiteboard activities, what was the nature of these activities, and how was the whiteboard being used in these cases? We came to understand that these unclassifiable activities were actually sets of related activities belonging to different quadrants, and that the representations users generated on the whiteboards allowed them to easily transition between these activities. Of the 11 interviews we conducted, 5 users had created representations (an integrated, related collection of marks) used in multiple activities/work modes;

another 4 used the whiteboard for multiple tasks, but employed a spatial partitioning strategy for each task (e.g. Mynatt, 1999). We illustrate how users employed the whiteboard to transition across multiple modes of activity by describing how three users used their whiteboards. These vignettes are real reports from three users from the former group; however, the names have been changed to protect the identity of the participants.

Vignette 1: Ongoing reference on a semi-public whiteboard

Larry is an engineering graduate student working in a shared lab with his peers. The main whiteboard in the lab is shared between the lab mates and their supervisor (whose office is elsewhere in the building, but who comes to the lab occasionally), and its location is such that it can be easily viewed from most areas of the lab. *For Larry, this shared whiteboard is used both to brainstorm and discuss ideas, and the same content is deliberately persisted, allowing it to be used as a reference for ongoing discussion, and as a personal reference for independent activity*.

One region of the whiteboard contained remnants from a recent brainstorm discussion with another student regarding a new project (collaborative/sync). This sketch was deliberately left on the whiteboard because it was incomplete. In the meantime, Larry and his collaborator had transitioned into a reflect-and-elaborate mode on the sketch (collaborative/async) so that when their supervisor returned from his week-long trip, they could, "*restart the discussion from this point,*" and resume discussing the ideas as a group. Thus, this single representation generated from collaborative brainstorming could be used later both by individuals and by the group for a brainstorming session.

Another region of the whiteboard contained a similar set of elements (a mix of sketches and text), but related to another project. Larry reports that this region is also partially the product of discussion, but that it is *continually maintained and used*. This content provides a number of functions: first as storage, so they can "*recall what we have discussed without much trouble*"; second, as a tracking mechanism for "*decisions from the previous week to… match [this week's] progress to what we decided last week*", and third, as an ongoing reference: "*One thing I did last week was a 'lit review' related to this discussion, so I kept coming back to see, to remember the points of discussion… like that sketch or plot there.*" Notice that the whiteboard content's representation facilitates its use each week for synchronous collaborative work, and through the week for asynchronous independent activity.

The lab whiteboard is used completely differently than Larry's meeting room whiteboard, where sketches never last beyond the duration of the meeting: there, sketches are only drawn as communication aids before being erased. On the lab whiteboard, the same content functions as a grounding mechanism for later discussion and further refinement, as a tracking mechanism for agreed upon goals, and as an ongoing reference for later personal use. The same information representation is used to enable Larry and his coworkers to transition between distinct modes of work, even though the *role* of the content is different in each use-context.

Vignette 2: Lo-fi ideation, deferral and storage of personal activity

John is a researcher for a small telecommunications start-up, responsible for delivering architectural designs that link together hardware and software components with customers' systems. For John, generating these designs is an iterative problem solving process that deeply involves his whiteboard. John's office whiteboard (visible from his desk) is used for generating, capturing and storing his design ideas, which he calls "brain states." *These "brain states" help John "think" with the whiteboards, and their persistence supports his ongoing activity as an organizing resource.*

> *In fits of ingenuity, I may come up with "this may solve the problem", and I want to capture [it] because it's important, but I don't want to capture it [formally]... The ideas are sketched out… and I have some key ideas to solving the problem, but [they] may not have been rigorous: I haven't thought of every situation, or cases where that solution may not work, so I have to think through those, or cases where I made assumptions that were erroneous.*

At the time John was interviewed, about 70% of the whiteboard content pertained to three such "brain state" sketches. John generates these "brain states" (independent/synchronous) to represent his current, up-to-date understanding of each problem he is tackling, and the space devoted to each design sketch is stable for fairly long-term (e.g. two or three months), informally capturing decisions and ideas.

> *Putting it on the board, it gives me these things I have to process... so I have to go research, [and] these ideas will send you on different work to prove them out. By having them on the board, when I start going down those tangents, if I don't write it down, I'll forget what it was. At least if I have it on the board—aha— this is what I was trying to do when I put this on the board.*

The sketches structure transitions in his ongoing work: they help John transition into "seek-and-understand" mode, persistently reminding him of unresolved issues or uncertainties in designs (independent/async), directing him to engage in communication with others, or to resolve them on his own. As John gathers more information or resolves these issues, he transitions back to thinking mode, updating and working on the design sketches (independent/ sync), so that the brain states are always up-to-date. The whiteboard and the brain state sketches ground John, reminding him of the tasks he was engaged in, or needs to be engaged in. Thus, since his activities often entail gathering information from others about questions or issues, the sketches function as a task list for John.

Once the ideas become more stable, and are captured with formal documentation, the whiteboard space is reclaimed. This example shows us that the whiteboard supports John's ongoing thinking process, helping him to transition from planning modes where he organizes his activities, to more active thinking activities.

Vignette 3: Persistent team scheduler

Jill is the project manager for a small web development company, and is responsible for a team of six designers and developers. Planning, managing and coordinating this team's schedule is Jill's primary challenge: at any given moment, Jill's team is working on up to six different projects (members contribute to just about every project), with personnel working simultaneously on different projects, and each project having dependencies on other team members and clients. Jill manages her team's schedule primarily from the whiteboard in her shared office (Figure 3.1). *This whiteboard, dedicated to the team's schedule, is used for multiple tasks: both for Jill's planning and reference activities, and for the rest of the team's awareness and discussion.*

Figure 3.1 Part of Jill's scheduling whiteboard, which is visible from her desk.

The schedule on this whiteboard is a six-week overview, organized into six vertical columns, with each column representing a week. Projects span across the columns, and each team member's tasks are colour-coded. Jill updates the schedule throughout the day, and once a week, Jill removes last week's column, and shifts over the other columns. Through the day, Jill receives requests from clients for new work. Because of the organization of the whiteboard and its location relative to her desk (it is visually accessible and only steps away from her seat), Jill can use it as a ready reference to rapidly assess the state of her team in the upcoming weeks and give immediate responses to clients (independent/async). If Jill decides that the team can take on the new work, she transitions into a planning mode, using the whiteboard to decide how the team's schedules in the next six weeks will be adjusted to accommodate this new work (independent/sync). The whiteboard allows Jill to try different versions of the schedule spatially, and to spot immutable deadlines and dependencies in the schedule.

Team members also use the whiteboard to maintain awareness about their schedule (collaborative/async) and communicate with Jill about their constraints (e.g. vacation time). Each Monday, the entire team meets in front of the project schedule whiteboard, and Jill can transition into a presentation mode, using the whiteboard as a presentation aid to communicate changes or updates to the schedule (collaborative/sync).

This whiteboard and the schedule representation are powerful: Jill reports that it acts as "ground truth" on her overall understanding of the team's progress, meaning that it also acts as an awareness display. The schedule representation facilitates transitions across multiple modes of work: when Jill uses the whiteboard to plan, she engages in primarily independent synchronous activity; she uses the whiteboard asynchronously to check on her team's status when clients ask about new work; her team regularly looks at it to maintain awareness of their schedules, and finally, when the entire team convenes weekly, the whiteboard functions as a shared display.

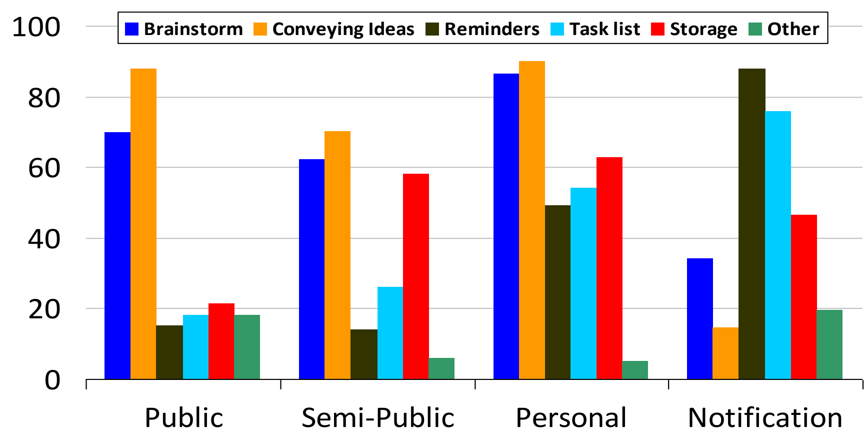

Figure 3.2 Location of whiteboards reported to be used for various activities.

Summary: These vignettes illustrate three instances of how users employ whiteboards to facilitate activity in several modes of work, helping them to transition between different activities with ease. Our analysis brings two themes to light: first, the whiteboard is useful primarily because users can *flexibly generate representations* of knowledge; second, while the representation may be static, their *role* and *function* in these different modes of work can be fundamentally different: the representation of Jill's whiteboard schedule, for instance, operates as thinking space, ready reference, awareness display and presentation aid. Thus, these users go beyond using the whiteboard for a specific activity, such as "information sharing" or "awareness", and can instead use it to fluently move between activities.

Role of Location and Social Practice on Transitions

Just as we found differences between users and their use of whiteboards for different activities, we suspected that there were different "types" of whiteboards in that they would actually be used differently from one another. Our data suggest that indeed, the physical location of whiteboard, its users, and the social practice that develops around the use of that whiteboard work in concert to shape its role in an environment.

With our sample of 239 whiteboards, we had also collected data about the frequency of their use for different activities (as in Table 3.3), where they were located (Table 3.2), as well as who typically made use of the whiteboards (self, close co-worker, co-worker, family, other, unknown). We analyzed this data using a *k*-means cluster analysis using Hamming distance as a similarity measure since some data was categorical (e.g. whiteboard location). The cluster analysis produced four stable clusters (Figure 3.2) which we labeled *post hoc* based on an analysis of the whiteboards in each cluster. The number of clusters, and the names of these categories is not important as they are likely to differ between samples (based on the way clustering algorithms behave); instead, the important observation is that while the whiteboard *artefact* is the same across contexts, it will have *different roles* in different physical and social contexts.

Public whiteboards (18% of the sample) were whiteboards located in public places that seemed to belong to no one, or were shared with anonymous, or "unknown" individuals. These whiteboards were primarily used for synchronous activities, such as brainstorming or conveying ideas in meetings, and are often wiped clean after being used. Lecture hall or boardroom whiteboards are examples of this type of whiteboard.

Semi-public whiteboards (27% of the sample) tended to be in shared locations (such as a lab), but the users and viewers of the board were typically known to one another. They were used for similar tasks as public shared whiteboards, but in addition were occasionally used for *storage* of information or shared knowledge. Storage is made possible because the user pool was known and fairly fixed—as a consequence, a common social practice or expectation about the role of the whiteboard could be developed over time. Whiteboards in workplace common areas, "war rooms", or labs are a good example of this type of whiteboard (e.g. Streitz et al., 1999; Whittaker & Schwarz, 1999).

Personal whiteboards (32% of the sample) were located primarily in users' personal workspace, and were therefore primarily used by the user in question (e.g. Mynatt, 1999). Only a small set of close co-workers were sometimes invited to use these whiteboards. It is on these whiteboards that content is used for the largest variety of tasks (Figure 3.2). This likely stems from its location (almost always being nearby and visible), and the limited set of users of these whiteboards (i.e. they are largely only used by the owner), so a fixed practice could be developed around the whiteboard itself.

Notification whiteboards (22% of the sample) were similarly often located in personal workspaces (and in the home). The users of these boards were almost exclusively the owner, and the boards were *primarily* used for asynchronous activities such as posting reminders, or task lists. The main distinction here is the relative dearth of synchronous activities on these whiteboards (e.g. brainstorming). These notification whiteboards were often *dedicated* to a specific purpose (e.g. a fridge whiteboard for messages or grocery list).

We would expect that with a smaller or well-known set of users around a whiteboard, a practice would evolve that allows them to develop expectations about the nature of the content on the whiteboard, whether it could be erased, what should be left on, and for how long. For example: "*If someone did erase [my whiteboard]… I would be upset. Maybe I should put a "do not erase" thing, but it's never been erased*." Notice the user's expectation of persistence and the lack of need for explicit signs on his personal whiteboard.

Another benefit of a small user group is the ability to develop a vocabulary or practice about how information is encoded on the whiteboard. Commonly used phrases may be transformed into abbreviations, concepts into symbols, and so forth. Indeed, during interviews, whiteboards contained many instances of short-hand or abbreviations—many of which were incomprehensible to the interviewer, though they were readily interpreted by the interviewee (e.g. "*It only makes sense to me because I use a shorthand notation for these kinds of things*."). Figure 3.1 provides a visual example where number-letter combinations represent vacation schedules, and colour is used to represent task type. We see then that the ink affordance is used to semantically encode information. Observations of "inside jokes" and differentiated usage between groups with MessyBoard accord with this interpretation (Fass et al., 2002). We therefore only expect the transitioning practice described earlier to function when a whiteboard is used by a closed set of users—these users develop a common practice and encoding schemes for information that is displayed. For instance, it would be foolhardy to expect content on public whiteboards to stay persistent without explicit requests on the whiteboard itself (e.g. "Please do not erase").

The distribution of activities across the different whiteboards (Figure 3.2) shows us that both the *physical context* (where it is located, and what it is nearby) and *social context* (who uses this whiteboard) shape how the tool is used and perceived by its users (Perry & O'Hara, 2003).

Beyond a Sketchpad: Whiteboard as a Place of Work

Throughout our interviews, it became clear that the whiteboard, beyond being just the *medium* for activity, was also a *place* where work was accomplished. We come to this interpretation because of the way information resources are brought and placed on or around the whiteboard. In many cases, we saw users placing information on the whiteboard for asynchronous purposes (as in Perry & O'Hara, 2003): either to remind themselves later of work that still needed to be conducted, or to support work that would be conducted later on the whiteboard. We illustrate the latter asynchronous case with two vignettes, again examples from individuals in our study whose names have been changed to protect their identity.

> *Lisa uses her office whiteboard as a project list, with "next steps" for each of the items in the project list. Of interest was a printed photo of another whiteboard that was affixed to the whiteboard next to one of the project items. The photo was of a different (lab) whiteboard on which Lisa and her students had engaged in an extended brainstorm. By positioning the photo on the office whiteboard, Lisa could not only maintain the existing use of the whiteboard as "project overviews" display, but also use the photo as a "window" to another, prior meeting. Lisa's use of this photo was temporary (it was removed in a week); however, here, the whiteboard functioned as a "storage device" to remind her of earlier work during a later meeting with her student (which took place in front of her whiteboard).*

Users place information next to the whiteboard often when they recognize that work is to be relevant for ongoing discussions. In the case of Fred, such work was formalized, and placed at the side of the whiteboard. It thus operated as a ready reference when engaged in later discussion, helping to convey those core ideas. Its placement next to the whiteboard was deliberate, allowing it augment collaborative sketching activity (e.g. design/brainstorming) occurring on the whiteboard.

> *[The paper attached to the whiteboard is] a high level architecture of a system what we're working on that we sometimes come back to. It's a project we worked on, and a lot of thought and energy went into it… I keep it up there so if I encounter other projects that are similar to that…, I use the same terminology. [I use it] especially for helping [when] we talk to some of our clients... and they want some capabilities from us, and want to know what's available in general for these things. So we can go through and re-use the structure we have for that.*

Thus, the whiteboard, beyond operating as a medium to support various sketching activities, is also a *place* where anticipated activities are expected to occur. Users take advantage of the fixed nature of the whiteboard to place and accumulate resources for these future activities.

Discussion and Implications for Design

We have seen that activity on a whiteboard can be usefully classified along the asynchronous/synchronous and independent/collaborative dimensions (e.g. Table 3.4). In relation to the model of transitions in Chapter 1, we see that the whiteboard facilitates several different types of transitions. As illustrated in Figure 3.3, whiteboards facilitate transitions between different types of tasks, different group members, as well as configuration patterns. Of particular interest in this particular case is that the artefact (the whiteboard) remains the same across activity states. As we discussed earlier (and continue to in this section), the whiteboard therefore helps to coordinate and smooth these transitions.

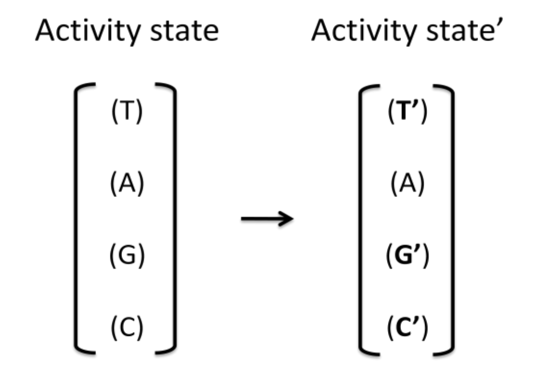

Figure 3.3 Whiteboards facilitate transitions in the task, group members, and configuration pattern components of our transition model.

We now synthesize our findings and existing literature to explore how we can design technologies that support these transitions. In so doing, we contrast a whiteboard's affordances with existing large display technologies, discussing how: (a) it is a *container* for task and coordinating information (Greenberg & Roseman, 2003), where (b) information is easily *revisitable*, (c) information is readily *updatable*, and (d) the flexibility *allows users to build representations of information suitable for many modes of activity*.

Whiteboard as a container. Building on Greenberg & Roseman's articulation of the "room metaphor" to support transitions, we also see the whiteboard as functioning as a *container* (Greenberg & Roseman, 2003). Information placed on many whiteboards is expected to be *persistent*. Similarly, the container is *permeable*, and readily provides access to that information. Whiteboards are typically constantly *visually available*—unless information has been deliberately obscured. This simple property has been difficult to replicate with large display applications. For instance, Huang et al. (2006a) report on how concerns over energy consumption (i.e. for projectors) often result in displays being turned off, thereby breaking the persistence of the information contained within. The fact that these displays often need to be explicitly turned on (e.g. Fass et al., 2002; Greenberg & Rounding, 2001) means that the information is not reliably visually accessible in the same way.

Beyond the virtual metaphor in Greenberg & Rounding (2001), the whiteboard has physical embodiment, and is *contextually located* near or in a place where action takes place (Hutchins, 1996). Thus, the whiteboard *limits access* to people who would likely be in the context, and *aids interpretation* by being in the same context (Mynatt, 1999; Perry & O'Hara, 2003). Fass et al. (2002) provide an instructive example: two MessyBoards were deployed nearby one another, but used by different groups. Each board was used differently, but the contents of the displays could be readily interpreted *because each MessyBoard was located near where the group using it sat*. In contrast to whiteboards, where physical access to the display itself engenders group processes (e.g. Whittaker & Schwarz, 1999, where a person updating a whiteboard-based schedule would result in a conversation about the reason for the update), MessyBoard enabled remote access (meaning that it was unclear who made updates to the display). Consequently, conventions around physical access (e.g. restricting access based on location) were lost, but new ones created (restricting login to a closed set of users). This example illustrates the tension between physical and remote access to traditional surfaces *vs.* large display applications. We see then that ***whiteboard practice is largely***

enabled by the conception of whiteboards as contextually located containers for visually accessible information.

Evolved meaning through representation. Beyond the artefact itself, meaning, as has been alluded to by several authors (e.g. Mynatt, 1999; Cherubini et al., 2007; Xiao et al., 2001), is created by the *users* of the whiteboard: information can be organized, drawn, written in any way the users like. This meaning can be embedded in spatial organization (e.g. via partitioning, as in Mynatt, 1999), and also via encodings and representations that users choose to use (e.g. through colour or shorthand). As illustrated by the vignettes, these representations can *evolve* over time as needs change (as in Larry's lab whiteboard), they can be *diverse* (as in John's brain state sketches, some of which are written, others of which are drawn), or employ space meaningfully (as in Jill's whiteboard, where columns of space represent weeks). *Users can mold the task-agnostic whiteboard with representations using ink primitives that are consistent and meaningful for multiple tasks.* ***Providing users with expressive primitives will allow them to flexibly generate meaningful applications themselves.***

This latter aspect of whiteboards presents a difficult design tension for designers of interactive whiteboard or large display groupware: how can we build and enable meaningful, powerful and flexible visual primitives without dictating their use? On a traditional whiteboard, primitives such as layout, color and partitioning allow users to construct meaning. Many systems similarly provide semi-structured primitives, where how the artefacts are used is not prescribed by the system itself. Notification Collage, for example, provided several widgets (e.g. text, URL, image), and the text widgets were appropriated for a variety of purposes (Greenberg & Rounding, 2001): notifications, reminders, peri-synchronous and synchronous conversation. In contrast to other systems that provided more sophisticated and integrated interaction (e.g. Trimble et al., 2003; Russell et al., 2002), we see that in many cases, simple, understandable metaphors can be easily appropriated by users for other unintended purposes (e.g. for posting personnel schedules (Churchill et al., 2004), or for simple game play (Fass et al., 2002)).

Flexible representations enable appropriation. Finding suitable middle ground in this design tension is difficult, but important: designers taking a careful application-centric view of groupware are likely to inhibit unusual or creative uses (perhaps deliberately) that allow the tool to be appropriated in other, or across work modes. On the other hand, by focusing on building suitably powerful primitives, users will be able to more flexibly appropriate the technology to their uses. In the context of interactive whiteboards, for instance, Flatland's approach allows users to create meaning with ink primitives, and provides functionality to specific segments in an on-demand basis (Mynatt et al., 1999). Flatland retains the 'one fixed page' metaphor of a whiteboard, providing "scaling" capabilities rather than relying on a file-based storage metaphor or switching interfaces for different applications altogether (e.g. Huang et al., 2006). An alternative approach is to organize interaction around shared artefacts (e.g. Geyer et al., 2003; Greenberg & Rounding, 2001), though this idea is perhaps better suited for distributed systems. In all of these cases, the focus is not on designing for a specific application or activity, but instead focusing design attention to core primitives, returning the meaning-making to the user while still providing powerful digital functionality. *The traditional whiteboard supports transition between work modes and activities*

Observation of Whiteboard Use	Implication for Large Display Application Design
"Ink" allows users to generate rich, flexibly meaningful representations of information	Designers should consider building functional *primitives* rather than applications to allow users to construct meaning
Whiteboards are contextually located containers for visually accessible information	Situated nature of large display can help determine which primitives are appropriate for the context

Table 3.5 Summary of the findings of the whiteboard study and implications for large display applications.

through informal ink primitives rather than structured interaction. ***Supporting transitions on interactive whiteboards means designing functional primitives rather than applications***.

Location and context of use. In the case of the whiteboard, we saw that *location* had a strong effect on the types of tasks it was used for. While intuitive, this has several implications for large display groupware. First, the affordances and functionality needed in different contexts is different—what is suitable in one context (e.g. in a personal workspace) may be wholly inappropriate in another (e.g. in a public area)—this may apply to input technologies, information sources, and so on. Second, users will employ primitives in unique configurations to support different types of activity depending on the location of the large display—thus, *location commutes meaning to displayed primitives*. Fortunately, only a few primary locations exist—the vast majority of important whiteboards were located in two places: personal workspaces, or shared spaces. ***Designers can rely on the situated nature of interactive displays to determine which primitives are appropriate for that context***

Section 1 Conclusions

Based on our data, we generated a classification scheme (Table 3.4) that describes users' activities involving whiteboards (synchronous vs. asynchronous work, independent vs. collaborative work). More generally, it is a model of activity that can reasonably be used to describe the type of supported activities in the large display applications described in Chapter 2. Beyond framing work in terms of the temporal and participation structure of activity, the scheme brings focus to the *transitions* between independent and collaborative work, as well as between synchronous and asynchronous work as suggested in our transition model of Chapter 1.

The examples we described in this section point to the dynamic role and function of traditional large surfaces to support tasks—even within the context of a single activity. As in the example with Jill, the whiteboard functions in a variety of roles for a multitude of related tasks. Our observations helped us to derive an understanding of whiteboard affordances that support this work practice of transitioning: visual persistence, flexibility of the ink primitives, and its situated social and physical context.

Section 2: A Study of Large Surface Use in Meeting Room Collaboration

In this study, we aimed to develop a descriptive classification of the roles of traditional surfaces in meeting room collaboration. As discussed in Chapter 2, prior work provides some insight into the general question of how whiteboards are used to collaborate. Studies of radically collocated groups (who make use of dedicated project rooms for independent and collaborative activity) have revealed that traditional large surfaces (e.g. whiteboards) are used to provide an active shared surface to work and develop ideas upon, and also to provide ready reference to coordinating information (Covi et al., 1998; Teasley et al., 2000). Medical personnel also make use of large shared surfaces such as whiteboards (Xiao et al., 2001). In this context, the workers *physically* work in the same general area, but are *temporally disjoint*: many nurses, for instance, will only pass by each other a few times during an entire shift. The whiteboard thus functions as a coordination centre, storing status and scheduling information of staff and resources (e.g. operating rooms), thereby putting this awareness and coordinating information on display at all times.

While this prior work provides detailed descriptions in specific case studies of whiteboard and meeting room activity, it has not provided a generalized model of this activity to inform design of large display applications—particularly within the context of meeting rooms. To inform the design of such applications, it is important to not only understand the nature of activities in such contexts, but to view these activities abstractly. Based on the study described below, we develop a classification of both the *activities* large surfaces are used for, and the *roles* these surfaces play in such activities. This scheme provides a basis for understanding the activity states that large display applications need to support. Furthermore, just as the teams in this study employed the surfaces for multiple activities, we would expect that such a large display application to support transitions between activities and tasks.

Method

We targeted our study to address two questions: first, what are the activities that large surfaces are involved in, and second, within the context of these activities, what role does the large surface actually play?

Participants. We recruited three existing teams of six undergraduate engineers (5 female, 13 male), all enrolled in a year-long team-based learning program. Teams completed four projects for the program, each taking five weeks. Our study focused on their third project, where teams were building magnetically propelled trains.

Environments. Teams worked in two time-shared workspaces assigned by course instructors: a meeting room (containing two whiteboards, a large table, two computers, and a filing cabinet) and a laboratory workbench (with two computers and electrical equipment). In a given week, each team was allotted two days in a meeting room, and two days in the laboratory space. Teams generally spent at least four hours per day working in the assigned space.

Method. We observed each team for at least four of their work sessions (each session lasting three to four hours), taking field notes and photos of the workspace, paying particular attention to the use of large surfaces such as whiteboards, tables, desks, and places where information was posted, such as doors, walls, and even the sides of filing cabinets. We augmented our field notes with opportunistic unstructured interviews with participants for clarification of the team's activities. Finally, we videotaped 60 hours of this workspace activity (across the three teams) for further analysis. This type of approach is consistent with prior research, where researchers have conducted ethnographic studies of workrooms (e.g. Teaseley et al., 2000; Covi et al., 1998). These "close quarters" studies of users in their working environment gives researchers a deep understanding of work practices within the specific context rather than relying on self-report.

Time	Type	Comment
8:15	Context	Ian & Thomas are working through a capacitancy problem (Ian initiated)
	Person	Thomas was trying to print something
8:22	Person	Thomas got up and is now whiteboarding something
		** Q: why don't they use the table for this?
	Activity	Thomas is basically sharing his idea with Ian
8:28	Positioning	Ian's position allows him to get up and point out things to Thomas
		** Ian took Thomas's pen --> how?
		Just by taking it -- it was his "turn"
		**NOTE: Watch pens and hands during this section
		** Explore how turntaking of ideas works here
8:29	Person	Thomas picked up different pen --> why?
	Person	Needs to write at the same time
8:31		Switch pens
		Hyp: How do pens instrument floor control & organization?
8:33		More switching of pens and ideas
		** Why --> colour coding
8:37	Person	Ian: "I'd like you to write that down" --> what does writing it down mean for permanence? Does that mean the whiteboard is all temporary?
		** Because room is going to be used by a different team later; board can't be depended on to be permanent
8:40	Activity	Tim has left to get pens to do math problems
	Activity	Ian is still working on design

Table 3.6 A sample excerpt from one of the analytic transcripts from the sessions. Multiple passes on the video data resulted in several annotations (identified by asterisks or italics). This transcript was also sometimes augmented by knowledge taken later from informal interviews.

Analytic Framework. We took a qualitative, ethnomethodological approach in our analytic method (e.g. Blomberg et al., 1993; Hughes et al., 1994). Field notes, observations and recorded activities were iteratively grouped into similar classes in an open coding process (Strauss & Corbin, 1990). We performed a modified interaction analysis (Jordan & Henderson, 1995) on our video data in multiple passes, using the field notes to sensitize our efforts, focusing on the structure of activity, spatial and temporal organization of activities and information, and participation structures in these activities. Table 3.6 provides an excerpt from a transcript in this analytic process.

Methodological Justification and Limitations. Our choice of study context was based primarily on convenience. We were interested in collaborating teams making use of traditional large surfaces. We used teams of engineering students because they were easily accessible, and beyond university ethics approval, required no additional provisions for protecting intellectual property. We chose to employ a naturalistic observation technique in order to immerse ourselves in the work of the participants: beyond merely examining the artefacts that appeared on the whiteboards, for

Activity	Description
Ideation	Activities involving focused, relatively synchronous interaction with the surface. Examples include brainstorming and design.
Explication	Activities where a users employs the surface is to support the explanation of an idea.
Notification	Activities where information is left deliberately on the surface to be used/read at a later time.

Role	Description
Present	Where the surface shows information for others.
Retain	Where the surface temporarily holds information.
Store	Where the surface collects and holds information long term.
Organize	Where the surface helps the spatial arrangement of elements.

Table 3.7 A summary of the classification scheme articulating collaborative activities traditional surfaces are used for, and the roles they play in these activities.

instance, our technique allowed us to see how the progression of activities in the space developed to include the surfaces, particularly over time.

However, our analytic orientation and choice of participants may limit the generalisability of our findings. For instance, our participants were students, and may not have worked on large projects before—thus, they may not be using the most efficient ways of collaborating or interacting with one another. Secondly, while our participants were using the surfaces to work together, they were also learning course material, and the fundamentals of group work. Finally, the scope of their work likely did not encompass all possible activities in project or meeting rooms given their overall engineering project requirements (e.g. Teasley et al., 2000).

In spite of these limitations, the study context is still useful: first, the students still do find ways to work together and to complete their tasks—thus, they are most likely to find the most obvious ways to use tools (such as whiteboard surfaces) to collaborate; second, we were able to watch the entire project cycle from conception to completion, thereby allowing us to see different collaborative needs throughout the entire project cycle, and finally, unlike other studies of collocated teams (e.g. Covi et al., 1998; Teasley et al., 2000), our students time-shared their workspace with other teams, meaning that students could not leave information or work artifacts out—instead, they had to develop ways of retaining this knowledge for later.

A Descriptive Framework of Traditional Display Surfaces in Collaboration

We developed a classification scheme by basing our analysis from two perspectives: (1) a *user-centric view*, where we draw attention to the activities (or tasks) that surfaces were used for, and (2) a *surface-centric perspective*, which identifies the specific interactional roles that surfaces played

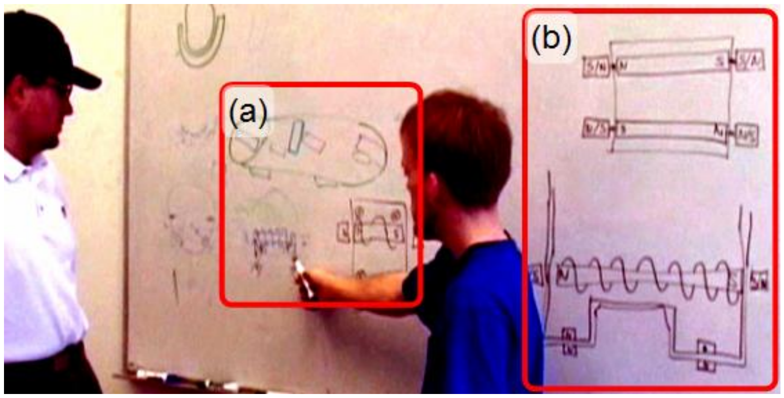

Figure 3.4 This frame comes from the transcript in Table 3.6, and illustrates idea locales: (a) the group of three sketches on the left belong to Alex's design, (b) Bob's design is grouped in two sketches.

in these activities. In the next subsections, we describe the conceptual components of this framework as outlined in Table 3.7.

Collaborative Activities involving Surfaces

Our initial analysis focused on the kinds of tasks or activities in which surfaces were involved. In particular, we were interested in the spatial and temporal organization of information and activity, as these would have significant implications on how and by whom that information could be used. We distilled three major categories of collaborative activities that used surfaces: *ideation activities*, which involve the generation and development of ideas; *explication activities*, which make use of the surface to explain ideas, and *notification activities*, where the surface is used to communicate asynchronously.

Ideation. A frequent category of activity involving surfaces such as whiteboards or paper on a table is ideation. Here, the surface is used as a dynamic work surface as a working memory store to work through ideas (e.g. when detail is generated around an idea) or to record generated ideas (e.g. in a brainstorm). In groups, this activity is often carried out on larger surfaces such as a whiteboard (Figure 3.4), whereas in smaller groups, sheets of paper on a table were often used. Speech frequently co-occurs with drawing during this type of activity, what we call *draw-alouds*, indicating that the drawing process is somehow linked to the generation of the idea itself (Tang, 1991).

Frequently, information and content placed on the boards during ideation activities seemed chaotic; however, we also observed some forms of implicit organization, which we call *idea locales* (Figure 3.4). Figure 3.4[1] illustrates a meeting where a pair was engaged in heavy brainstorming about two possible designs. Students worked through each design in turn, and investigated shortcomings and filled out detail by drawing atop an existing design, or by using callouts. These sketches were implicitly organized into *idea locales*: related ideas were frequently deliberately placed nearby. This locality helps to dynamically coordinate space on a large surface (e.g. Mynatt et al., 1999).

These idea locales are functionally similar to, but conceptually distinct from the *segmentation* described by Mynatt (1991). Segments, as described by Mynatt, separate semantically distinct

[1] Names of participants have been changed to preserve anonymity.

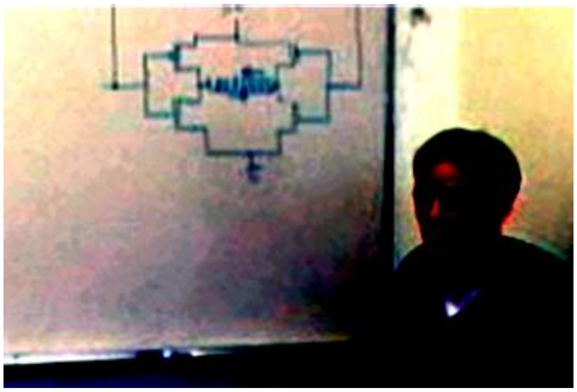

(a) Alex turns and asks for confirmation before continuing.

(b) Alex continues to draw while speaking (a *draw-aloud*).

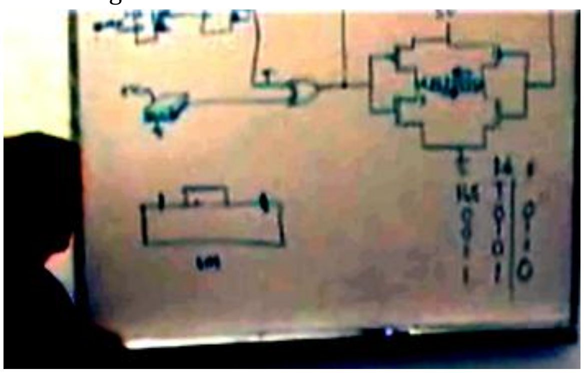

(c) Notice the parenthetical on the bottom left. This is erased and replaced in the next frame.

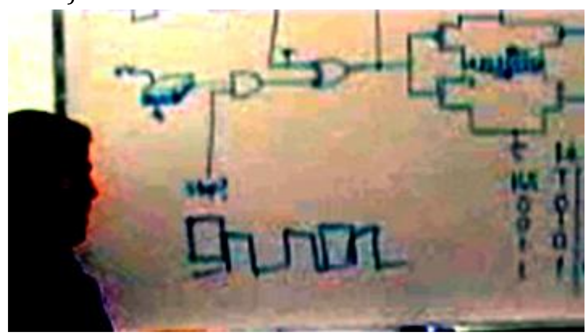

(d) Becky asks for clarification and Alex obliges (the bottom left now has a different sketch now).

Figure 3.5 This sequence illustrates parentheticals and stepwise use of the display, characteristic of explication activities.

groups of information (e.g. a phone list in the top right of the whiteboard, brainstorming in the centre, and a task list on the top left). Idea locales are naturally occurring groupings in a short temporal span, and may be quite related to one another.

Explication. This category characterizes activities where the surface is used to explain an idea. Generally, there are fairly clear role divisions in this kind of activity (presenter and audience), although the role of presenter is sometimes swapped between group members. Figure 3.5 illustrates such a sequence. Here, Alex was explaining the design of his circuit to his team. To do so, he placed the circuit on the table, and used the whiteboard to draw attention to particular design decisions in specific areas of the circuit. His use of the whiteboard was surprisingly methodical, and he added additional information only after brief negotiation with his audience (e.g. Figure 3.5a where Alex turns and asks, “Does everyone get this?”).

The careful, deliberate use of the whiteboard distinguished explication from ideation: presenters seemed to essentially know the starting and ending points of the discussion (and therefore have a sense for how content might be laid out on the surface). For instance, it was rare for an entire surface to be wiped clean: changes were generally evolutionary. On whiteboards, *idea locales* were typically wiped clean as units, and the rest of the surface not disturbed. Similarly, clearing an entire surface of its contents indicated that the task focus was to change drastically, so it was quite rare in the collaborative flow.

Some board content was not planned. For instance, we also saw drawings and sketches we call *parentheticals*, which were used to explain some concept (e.g. a piece of background knowledge). These *parentheticals* were typically located non-centrally on the surface (Figure 3.5c,d), and were often erased once a concept was clear (Figure 3.5d). Their transience reflects their relative importance to the explication activity: they are temporary asides that are intended to support the explanation, but are not intended to detract from the explanation's central flow.

Notification. This category of activity typically involved information being left on a surface for long periods of time, edited rarely and subtly, and often used to provide awareness of a group's activities or intentions (e.g. action or TODO lists). The primary motivation seemed typically to typically be to communicate with others (or oneself) asynchronously. For instance, we saw instances of messages being left for one-time use (e.g. for tardy group members: "Larry, we are in the wood shop!!"). Thus, surfaces used in this way provide information in an ambient fashion (e.g. Moran et al., 1996; Teasley et al., 2000).

Summary. Our observations helped us to come to a three-activity descriptive model of the collaborative work around traditional surfaces in this context. Figure 3.6 illustrates how these activities relate to the temporal axis of the model we introduced in Table 2.2. This model also follows the types of activity supported by most large display applications (as described in Chapter 2): the "ambient" uses maps onto the asynchronous activities, while the "knowledge work" uses map to the synchronous activities that we have described.

Roles of Surfaces in Collaborative Work

During collaborative activities, individuals play specific roles in the process (e.g. moderator, transcriber, critic, etc.); similarly, we found that the surfaces seemed to play very specific roles in collaborative activities. We arrived at surface-centric role classifications by coding how surfaces were used within the context of these activities, and grouping thematically related observations together: in the *present* role, information is placed on the display explicitly for use by others; in the *retain* role, the surface "stores" information for immediate use; in the *store* role, information is preserved on the surface for later use, and in the *organize* role, the large surface is exploited to facilitate semantic spatial organization.

Present. When information is placed on a surface to exhibit information to others, it is said to be fulfilling a presentation role. In our observations, when a surface was being used in a present role,

Synchronous	Asynchronous
Ideation	Notification
Explication	

Figure 3.6 Activities from this framework in relation to the synchronous/asynchronous model of activity from Table 2.1.

distinct roles of presenter and audience frequently emerged, and these were related to the kinds of interactions taking place with the surface: presenters changed content of the surface, and the audience merely viewed (Figure 3.6). The information on the surface was therefore the focal point of discussion, and so during any discussions that might occur, the information on the surface was often referred to explicitly (by pointing gestures and so forth).

In our observations, the presenter generally stood closer to the surface (typically a whiteboard) itself, and so was generally tasked with any interaction with the surface. In contrast, the audience was generally further away from the surface, and so their interactions with the surface were restricted to viewing and pointing at the surface. Thus, one's proximity to the surface conveys differential social roles in collaborative activity: being close to the surface means that one also controls surface content (e.g. Rogers & Lindley, 2004).

In several instances, we saw presenters working from notes, suggesting that information being presented on a surface was often prepared in advance.

Retain. Surfaces being used in the retain role are quite dynamic, with constant addition, editing, removal of information occurring, and at times co-occurring (e.g. a whiteboard in an ideation activity). Individuals are within close proximity of the surface itself, all able to essentially touch, point, and make changes to the surface (e.g. Figure 3.4). To some extent, we saw interaction was mediated by certain interaction artefacts (e.g. an eraser or pen), but this did not appear to adversely affect the flow of activity—the transitions were quite smooth.

The use of a surface, as in the retain role, was generally limited to the number of people who could comfortably use the surface, but the number of participants ranged from a single individual right up to that upper limit. In this role, the surface provides support as a form of working memory, storing information for immediate use (i.e. within a few minutes). Ideas are often modified in place, and in the form of changing sketches, words, and at times, the addition of detail. The surface therefore lends itself to a form of transient storage, since information only persists for the extent of the discussion. A good example of this transience is the *parentethicals* described earlier. The retain role therefore supports rapid exploration of ideas, allowing ideas to be removed in part or wholesale at any time.

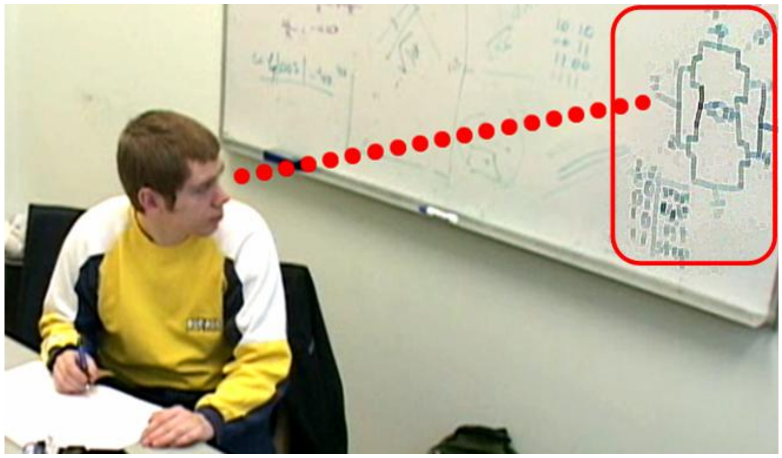

Figure 3.7 Larry makes use of information the team had created earlier on the whiteboard.

Store. The storage role appears when a group or individual's activity changes dramatically from when information was placed on the surface to when the information is used again. For instance, Figure 3.7 shows an example of a student making use of a sketch the team had made earlier on the whiteboard. Use of the surface in this way is not always planned: often, the information being used as a reference was the result of prior work using the surface in the scratch role; in other cases, information is explicitly placed on a peripheral surface for later use. In another example, we observed a student engineer returning to a chalkboard to review a sketch drawn earlier by an instructor, stressing again the importance of storing information in context for use at a later time.

Organize. Surfaces such as tables are often used as large workspaces where the spatial relationships between units of information and content can be interpreted semantically. Depending on the particular type of surface and the content that is used, this information can be organized in an ad-hoc fashion (while the information is being created or added), or in a post-hoc fashion (after the information has already been created or added, the location of the information is changed). This role is evident both when teams are attempting to structure ideas and information, as well as when tools and raw materials are being organized. Traditional physical surfaces such as tables support this organization role well by providing fine-grained means to locate and orient information for coordination (e.g. Luff & Heath, 1998). Idea locales on whiteboards (e.g. Figure 3.5) also reflect how spatiality and location of information can contain or maintain semantic metadata about the information.

Section 2 Summary

The classification scheme categorizes the breadth of activities traditional surfaces are used for in meeting room collaboration, and the roles these surfaces play in those activities. While the scheme is abstract, we have attempted to provide a user model from which designers can base their large display application designs. It acts as justification for certain types of applications for this context, as well as a mechanism through which designers can discount certain applications (i.e. if they do not meet the requirements and needs of the users as set out in this classification). The main insight provided by the scheme is that it shows the wide variance of information needs and activities within even the context of a meeting room.

Chapter Summary

The studies described in this chapter provide a rich perspective on how traditional large surfaces are used in work.

In the first study, we provide a classification that describes users' activities with whiteboards (Table 3.4). This scheme classifies activity in terms of participation (independent vs. collaborative) and temporal (synchronous vs. asynchronous) structures. We then illustrated that whiteboards were used across these modes of activity, supporting users' transitions between activity modes. The occurrence of these transitions was explored through descriptive case studies of users' practices around their whiteboards, and we identified how spatial and semantic ink encodings were used to support these transitions (Table 3.5).

In the second study, we developed a second descriptive classification for the roles of traditional surfaces in meeting room collaboration. This scheme further illustrates how traditional surfaces support multiple activities, and provides the basis for understanding the types of roles large display applications should likely support, and the types of activities that applications should support transitions between.

Together, these studies inform our understanding of work practice involving traditional large surfaces, identifying the role of transitions in shaping users' practices with whiteboards, as well as describing the types of activities large surfaces are used for.

In the next chapter, we explore the design space of large display applications, and the role transitions play in these applications: we design, build, and observe the use of two systems (alternately through an observational lab study and a public field deployment). In the first case, we explore the transitions of users between highly coupled and loosely coupled work in the context of a digital tabletop application, calling these transitions "collaborative coupling." In the second, we study and explore the transitions of users around large public digital displays, articulating their transition between bystander and contributor. These studies help to broaden our definition and understanding of transitions, relating them to the changing needs of users as they make use of large display applications.

Chapter 4

Large Display Users' Changing Needs

While Chapter 3 focused on the multiple roles and activities supported by traditional large surfaces, there is also need to understand how *users'* roles change during work. That is, how do their information and interaction needs change? Furthermore, because large display applications are likely to be used by multiple users, what are the needs of these users, and how can we go about supporting their differing and independent needs? The goal of this research phase was to develop models that capture users' needs as a function of their use of large display applications.

To build these models, we wanted to develop breadth in our understanding, so focusing on a single type of activity context (i.e. whiteboard-like applications) would not be sufficient. Consequently, we established two requirements for choosing the functionality of these applications: first, it was important to create applications that were functionally distinct from traditional whiteboards; second, we wanted to sample from disparate regions of the design space as a means to explore different types of user needs (i.e. we needed to study more than one system). In this chapter, we report on two such systems that we designed, implemented and studied.

The first case study investigates the design of an interactive tabletop for mixed-focus collaboration (where parts of the task require independent work, and other parts of the task require collaborative work). Based on two observational studies of pairs using this system, we articulate a model for describing how users fluidly engage and disengage with one another's work, and how they transition between these collaborative coupling styles (Tang et al., 2006).

The second case study describes the iterative design and deployment process of an interactive public large display deployed in a public space on campus (Tang et al., 2008). We develop a classification of bystanders, articulating how these users transition between different types of bystanders and contributors, particularly with regard to their changing information needs. Finally, we discuss three themes relevant to the design of systems for bystander use.

		Collaborative Tabletop Routing Application	*MAGICBoard*
Description		Tabletop-based application intended for collaborative route finding.	Interactive public display that collects opinions of users.
Application Type		**Knowledge Work:** Involves focused interaction with data and other users through the application.	**Ambient Display:** Primarily provides awareness about some information source.
Interactive Display Design Parameters	*Orientation*	**Horizontal**: Tabletop surfaces are of interest because users' configuration around such a display lends itself to face-to-face orientations.	**Vertical**: Upright surfaces can be more readily observed from a distance, and by more users.
	Input	**Direct**: Multi-touch using DViT (infrared camera system). A direct input mechanism is one where the interaction takes place in the same area as the response.	**Indirect**: SMS sent via cell phone to the display. An input mechanism that is not direct is considered as an indirect input mechanism.
Application Design Parameters	*User-User Relationship*	**Collaborators**: Other users are well-known.	**Strangers**: Other users are not likely to be known.
	Coupling Style	**Mixed-Focus Collaboration**: Each user's actions affect others, both in terms of the workspace, and in terms of success.	**Not Coupled to Loosely Coupled**: Users' activities are effectively independent events, and are mainly unrelated to one another.
	Temporal Interaction	**Synchronous**: Users are working and interacting with the surface simultaneously.	**Asynchronous**: Users' interactions with the surface are at different times, and are considered as independent events.

Table 4.1 Summary of the design choices we made in the two large display systems we designed and studied in this chapter to explore transitions.

Table 4.1 briefly summarizes the characterizing design parameters of the large display systems we built. While it is beyond the scope of this work to examine each of the design parameters in detail, other researchers have demonstrated that many have an effect on users' interactions with both the system and other users (e.g. direct vs. indirect input mechanisms (Ha et al., 2006), horizontal vs. upright displays (Rogers & Lindley, 2003)).

Together, our studies illustrate that users' needs change during their use of large displays, which has direct implications for the design of large display applications. We develop two models describing these changing needs, and discuss their implications for design—in particular, that effective large display application design should provide mechanisms to support users' changing needs.

Section 1: An Interactive Tabletop for Mixed-Focus Collaboration

Many collaborative activities are characterized as being *mixed-focus collaboration*, which involve both independent and shared tasks with the group. Mixed-focus collaboration presents many challenges for system design because such systems must support both individual and group needs, which are often in opposition (Gutwin & Greenberg, 1998): for instance, should individuals be able to partition the workspace's view, or should the group share a single view? While independent views of a shared workspace may better support parallel execution of individual tasks, they may interfere with others' tasks, or shared tasks. Furthermore, they may also negatively affect the group's ability to coordinate its activities and manage shared resources. How should a large display application support activities that involve mixed-focus collaboration?

We develop this understanding within the context of a tabletop application designed to support a knowledge work activity, where users are asked to perform collaborative tasks on a shared visualization. We present two observational studies that examine how three viewing techniques affect coupling: (1) *lenses,* which show information in spatially localized areas, (2) *filters,* which show information globally, and (3) *ShadowBoxes,* which allow spatially localized areas to be displaced (Stone et al., 1994; Ware & Lewis, 1995). As we will see, these tools affect how groups coordinate their use of the shared space and their engagement with one another—their coupling style.

Based on the findings from this study, we develop a descriptive model of these coupling styles, which describe how users' level of engagement with one another change as they make use of the large display application.

We conduct two observational studies of collaborative activity with our tabletop application. Our first study provides insight into how groups coordinate activity over a spatially fixed visualization. The second study reveals six distinct styles of coupling and how they relate to factors such as task, physical location around a tabletop, and interference management. These results motivate several design implications for the design of fluid, collaborative tabletop interfaces that support transitions between collaborative coupling styles.

Techniques for Collaborative Exploration of Fixed Spatial Data

In some mixed-focus collaboration tasks, groups must share fixed spatial areas such as those involving maps. In such tasks, the visualized data set takes up the entire display and cannot be moved from its location, potentially introducing physical and visual interference issues when individuals need to work independently in the same area. For example, in a meteorology application, one person may need to examine wind patterns while another studies temperature and pressure, both in the same geographic location.

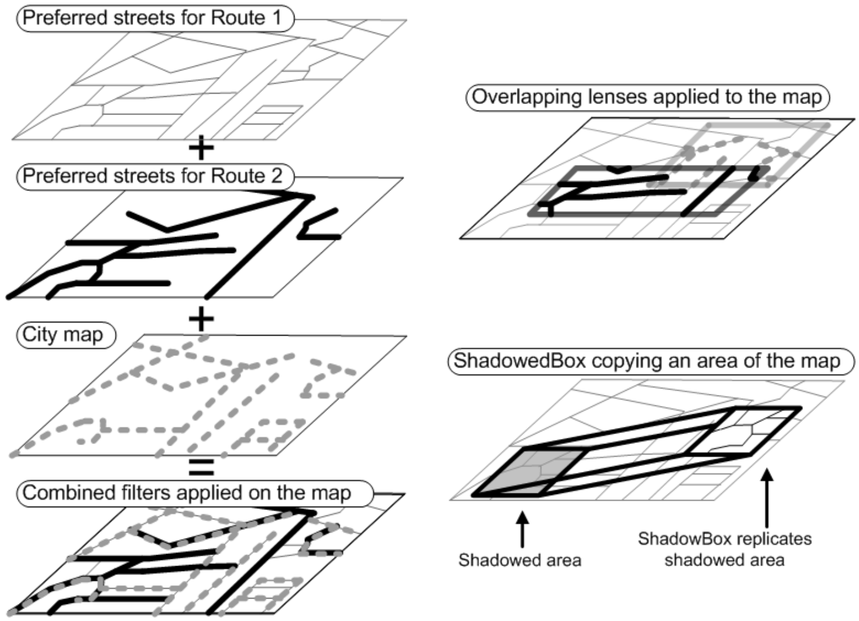

Figure 4.1 Three display techniques for coordinating space: filters being applied to a map (left); lenses being applied to the same map (top right), and a ShadowBox relocating a region of the map (bottom right).

Figure 4.1 shows three techniques for exploring fixed spatial data that potentially support different working styles by providing view-based partitioning of the data: *filters* (left), *lenses* (top right), and *ShadowBoxes* (bottom right). Filters are common in commercial mapping applications (e.g. Google Earth), allowing users to selectively view multiple "layers" of visual information by simply stacking them atop each other. This global approach provides a single view of the workspace. In our studies, filters could be displayed or hidden using a set of graphical buttons. Lenses are mobile, resizable windows providing the same set of data visualizations as the filters, except in a localized view (Stone et al., 1994). Several data layers can be shown simultaneously by overlapping multiple lenses. Lenses facilitate local view changes without affecting the global space. In our studies, lenses were created by using a set of graphical buttons, and were moved and resized by dragging their borders or corners, respectively. ShadowBoxes allow users to select an area of the display and copy the underlying information to a moveable viewing window, similar to the DragMag visualization technique (Ware & Lewis, 1995). Interactions in the viewing window are "shadowed" in both regions, meaning that drawing and erasing activities in either location is immediately reflected in the other. This displacement allows multiple individuals to work on the same part of the data simultaneously.

In our observational studies, we were interested in determining how these tools supported mixed-focus collaboration and coupling. By providing a single view, filters were expected to ease communication by facilitating gestural and deictic references (Gutwin & Greenberg, 1998). Independent work was expected to be disadvantaged because view changes were global. Lenses were expected to support spatially distinct, local views of the data, allowing individuals to view and work on parts of the workspace independently. Lenses were not expected to solve the problem of physical interference that might occur when two individuals want to work in the same physical space. We believed ShadowBoxes would provide a solution to this problem by allowing individuals to work in the same part of the data in physically distinct locations.

Overview of Observational Studies

We conducted two observational studies to better understand mixed-focus collaboration. In the context of a collaborative visualization task, we wanted to understand how tools such as filters, lenses, and ShadowBoxes would be used for both the independent and shared work aspects of mixed-focus collaboration and how these affected coupling.

In our first study, participant groups created routes by connecting multiple end points on a fictitious city map using our three tools to reveal data on the underlying map. This exploratory study was designed to understand how participants would use the tools to coordinate their activities over the workspace.

In our second study, participant groups created routes connecting multiple end points on a fully connected graph. The purpose of this second study was to confirm the presence of certain coupling patterns observed in the first study and to characterize the role of coupling in these activities. We included specific roles for individuals within a group, and created independent and shared tasks to tease out the transitions between individual and group work.

Study 1: Exploring Group Work

To explore how groups work over spatially fixed data sets, we designed a map-based route creation task requiring collaborative visualization. Pairs created two separate bus routes in a map of a fictitious city (Figure 4.2, left). Beyond simply creating routes to connect designated end points, participants were to optimize their routes based on a set of constraints, ensuring that created routes: (1) were reasonably direct, (2) traveled along preferred streets, (3) passed through residential and commercial zones while avoiding industrial zones, and (4) avoided overlapping with each other. Various data layers including the street map, "preferred streets," and locations of residential, industrial, and commercial zones were provided to groups to help them construct routes (Figure 4.2, left). These data layers were accessible to participants via combinations of filters, lenses and ShadowBoxes, depending on the study condition they were completing.

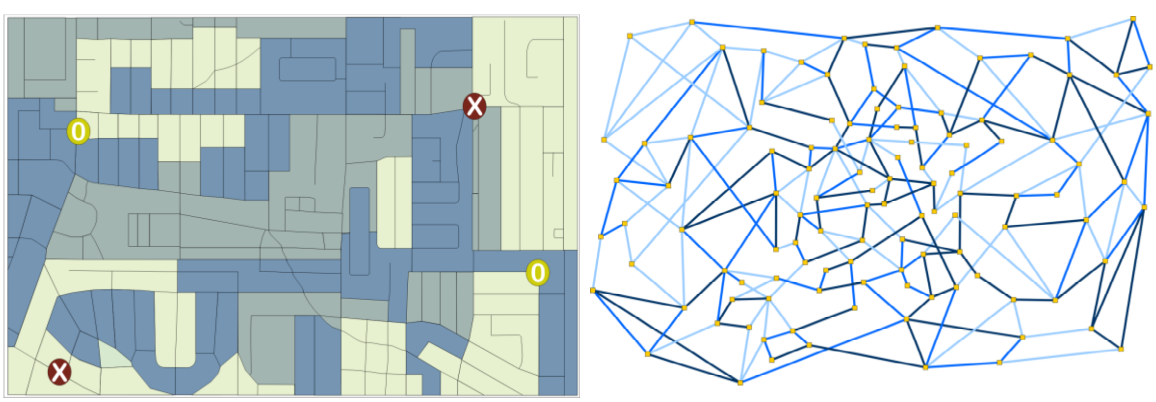

Figure 4.2 The street map provided to participants in Study 1 (left) and one of the fully connected graphs provided to participants in Study 2 (right).

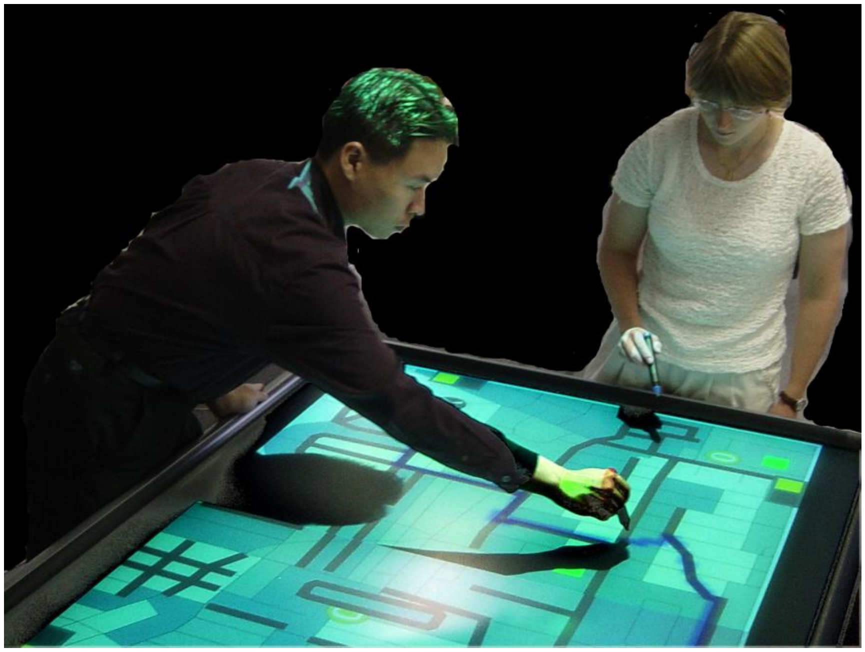

Figure 4.3 Tabletop users completing the task from Study 1.

Based on prior work, we expected groups to exhibit certain kinds of behaviour:

- *Divide-and-conquer.* We expected participants to use a divide-and-conquer approach to the task, with each individual working on a separate route (e.g. Gutwin & Greenberg, 1998).
- *Individual work would be better supported by lenses and ShadowBoxes*. We expected participants to prefer lenses and ShadowBoxes for individual work because these tools would allow individuals to work independently without disturbing the view of others. We also expected filters to induce interference because one person's view of the space would affect the entire group's view.
- *Group work would be better supported by filters*. We expected participants to prefer filters when engaged in group work because we believed the single view provided a shared context for discussion, and interference would not be an issue.

Design

Our exploratory study used a 2 (filters vs. lenses) × 2 (with ShadowBoxes vs. without ShadowBoxes) within-subjects design. The presentation order of the conditions was counterbalanced across groups using a balanced Latin square because pilot testing suggested presentation order affected work strategy. Thus, every group participated in four different conditions: (1) filters with ShadowBoxes, (2) filters without ShadowBoxes, (3) lenses with ShadowBoxes, and (4) lenses without ShadowBoxes. Every group received a unique presentation order of these conditions.

Participants

Eight paid participants (four pairs: five males, three females) with normal or corrected-to-normal vision were recruited from the general university population. Each participant group was made up of two people who knew each other well. All participants reported being right handed. Four had previous experience with large displays, two of those had experience with tabletop setups, and five had experience with Web-based mapping software. Mean age of the participants was 29 years.

Apparatus

We used a large front-projected tabletop display (5 × 4 feet) with high resolution (1534 × 1024 pixels) supporting simultaneous two-touch interaction via SMART Technologies' DViT. Participants could interact with the table by directly touching the table with pens or their fingers, though most participants only used pens. We placed two chairs within easy reach of the tabletop display and told participants that they could use them, but no participants chose to use the chairs in this study.

Our custom-build groupware application was built with C#, and Direct3D using the Trans2D library[2]. This software ran on a dual-Xenon 2.8 GHz Windows XP PC. Dragging a pen or a finger over the tabletop display would draw routes on the custom-made map while dragging a digital eraser widget would erase drawn routes. The application also provided widgets to control the display of additional data based on the study condition. Because our tabletop display supported two touch interactions simultaneously, both participants in a group could draw, erase, or manipulate widgets on the tabletop at the same time.

Method

Participants first filled out a questionnaire to collect demographic information and to assess their experience with mapping applications. They were then given a short tutorial on how to use the table display and general instructions on the task. Prior to each of the four conditions, participants practiced with the tools they would be using for that condition. Participant groups generally felt comfortable about using the various tools with less than five minutes of practice before each condition.

During the conditions, we instructed the participants to use a "talk aloud" protocol, and videotaped their interactions with each other and the tabletop for later analysis. On average, groups completed each trial in approximately 15 minutes. Once all four trials (one for each condition) were complete, participants took part in a semi-structured feedback session, which allowed us to gain valuable insight into their impressions of the task, the different interaction techniques, and their own performance during the study.

Three observers were always present during study sessions with groups, though only one directly interacted with the participants; the remaining two were passive observers during conditions. Observers collected field notes of group behaviours as they worked, which were later combined with the video recordings for a full analysis.

Results

Collected field notes and video were analyzed using an open coding approach similar to that used in other research (Kruger et al., 2003; Scott et al., 2004; Tang, 1991). Field notes were used to inform initial coding categories, such as whether participants were working independently or together at a given point in time. A video analysis was supplemented by a descriptive statistical analysis. We present our most salient findings below.

[2] Trans2D. http://mail.rochester.edu/~mabernet/trans2d/.

Tendency to work together

Contrary to our expectations, pairs worked together across all conditions, visibly working independently for only 24% of the total time. This was surprising because we had predicted that participants would prefer to work independently in the presence of lenses or ShadowBoxes. In only 6 out of the total 16 study conditions (4 groups × 4 conditions each) did pairs even attempt to divide up tasks. Groups generally worked together to find one route before finding the other route. Groups were highly mobile, with individuals frequently moving around the table to gain a shared perspective of the area of interest. Groups also worked in tandem. Often, one person would control the widget (either lenses or filter buttons) while the other would draw the route on the display. In some sense, this division of labour could be considered as divide-and-conquer; however, the pairs were working closely together on the same problem as opposed to working independently on different aspects of a problem.

Group 3 was a notable exception. In the filter conditions, this pair worked in parallel on different routes. To facilitate this parallel operation, they used the filters in a "time sharing" mode: when one needed to see a given data layer, he would tap and view his layer for as long as he needed while the other worked from memory. Group 3's working style suggests that some groups may desire to work independently. Group 3 found an awkward way to support their independent working style since the interaction widgets we provided did not provide fluid support for them.

This observation suggests that the mechanisms that we designed to support individual activity were not, on their own, sufficient in preventing visual or interaction interference. As a consequence, users modified their behaviour to facilitate a more collaborative style, where interference would be less of a problem, and working together rather than independently.

Maintaining context

In addition to working together most of the time, pairs did not use the data widgets in the ways we had expected. Most strikingly, participants overwhelmingly preferred global filters. In lens conditions, users essentially mimicked the functionality of the filters (even when the filters had not yet been presented) by creating table-sized global lenses, moving them in and out of the workspace to cover the working area. ShadowBoxes were simply moved out of the way; widgets that affected the global space were preferred.

Participants reported that the lens widgets suffered from several usability problems. First, they were somewhat cumbersome—resizing and moving the lenses required a switch from the route planning task to a widget manipulation task. Second, lenses did not support the way in which participants worked (i.e. as a group rather than independently). Finally, lenses could not meaningfully partition the space because each lens needed to be larger than half of the table to provide enough information to plan each route. Since the task involved planning global routes, participants preferred global filters, which provided global rather than simply local information.

Discussion

Pairs were mobile and *non-territorial* when working over the spatially fixed data. For the most part, they moved together, worked together on the same route, and did not work independently. The entire workspace was therefore group territory and the tools for establishing personal territories (i.e. ShadowBoxes and lenses) were not used. Groups preferred visual mechanisms (i.e. filters) that allowed them to view the space together, even frequently standing in close proximity with one another, which was surprising since the workspace had no orientation cues.

From Study 1, we began to think about how to describe the group activity. For instance, many groups spent time working very tightly coupled, but in different ways: at times, they would draw routes together, and at other times, they would simply point alternatives out where one individual was more active than the other. Group 3 exhibited loose coupling, and often worked in parallel. To understand this issue of coupling further, we modified the task to explicitly include independent activity in Study 2.

Study 2: Transitions in Group Work

Based on the outcomes of Study 1, the following objectives were established for Study 2:

- *Giving participants independent roles*. Participants often worked in tightly coupled fashion in Study 1, but this may have been because they were not given independent roles.
- *Explicitly introducing independent and group tasks.* By imposing activity at the extremes of mixed-focus collaboration (independent work and shared work), we hoped to observe a range of group activity.
- *Multiple sub-problems.* Study 2 had three sub-problems that could be spatially partitioned (i.e. a person could work on each sub-problem without requiring the entire work surface). To induce instances of interference, one of these sub-problems slightly overlapped with the other two. Study 1's sub-problems covered the entire space, so spatial interference may have precluded independent work.
- *Completely conflicting data layers.* Data layers in Study 1 overlapped only in certain regions, so participants could often work with all filters turned on. In Study 2, we used completely occluding data layers to preclude this strategy, and to simulate situations where there are so many data layers that displaying all the information needed by one person will necessarily interfere with the other.
- *Redesigned lens widget.* Based on Study 1 feedback, we redesigned the lens to include filter buttons that could selectively apply layers in a local space.
- *Removal of the ShadowBox condition.* To focus our efforts on the effects of local and global views on independent and group tasks, we removed the ShadowBox condition. Including a ShadowBox condition in this study would have prevented us from practically using a within-subjects design.

Pairs found routes in a fully connected graph (114 nodes, 218 edges; Figure 4.2, right) covering the entire workspace. This task represented an abstract route planning task (such as airline routes). Two independent data overlays provided edge weight information ("travel time" and "financial cost"), where the weights could be 1, 2 or 3. Participants generated routes to connect four specific nodes on the graph. Depending on the condition, each participant was responsible for generating one of two *independent routes* (one for travel time, one for financial cost), or the pair was responsible for a single, group *compromise route* (taking into account both travel time and financial cost). We also varied the visual tool pairs used: global filters, or the redesigned lenses.

Design
Study 2 used a 2 (filters *vs*. lenses) × 2 (individual routes *vs*. compromise route) within-subjects design. The presentation order of the conditions was counter-balanced across groups using a balanced Latin square.

Hypotheses
Based on the results from Study 1, we had two major hypotheses.

- *Individuals will work independently with lenses.* Since lenses allow people to work in different parts of the table and some sub-problems were spatially distinct, we expected independent work to occur for those independent sub-problems. We expected this to occur even when participants were working on a compromise route, since participants could use lenses to work on different areas of the route at the same time.
- *Perspective sharing during tightly coupled work*. When working together on the same sub-problem, we expected groups to stand in close proximity to each other, thereby allowing groups to share the same perspective view of the problem space.

Participants
We recruited eight paid participants (four pairs: four males, four females), different from those in Study 1, with normal or corrected-to-normal vision from the general university population. Seven were right handed, two had previous experience with large displays, none had experience with tabletop setups, and six had experience with mapping software. Mean age of the participants was 28 years.

Apparatus and Method
We used the same apparatus and setup as Study 1, except that we replaced the map with a custom-made, fully connected graph, and that participants used the global filter and redesigned lens widgets. Study 2 used an identical protocol to Study 1.

Results
Video was analyzed using a multi-pass, open coding approach similar to (Kruger et al., 2003; Scott et al., 2004; Tang, 1991). Field notes were used to inform initial coding categories, such as individuals' positions around the table and which sub-problem each was working on at a given point in time. Subsequent coding passes were driven by iteratively refined coding schemes based on further study of the videos. This methodology facilitates an intimate familiarity with the intricate, subtle mechanics occurring in the sessions, providing a very rich understanding of the underlying collaborative processes.

This study imposed a variety of activities ranging from independent to group tasks, allowing us to explore a range of collaborative behaviour. Our analysis revealed six different types of collaborative coupling. These coupling styles were related to a range of other factors, including the experimental condition (i.e. task type and tools being used), collaborators' physical positioning around the table, and how interference was handled, providing strong support for our coding scheme. We begin by describing the six coupling styles we saw, and then describe other factors and how these related the coupling style.

Styles of coupling

Based on field notes, we iteratively refined a coding scheme for the videos in Study 2 to describe and capture the dynamic styles of coupling for each group. Individuals fluidly transitioned between styles, for example, moving from tight coupling, actively working together, to "medium coupling," where they worked somewhat independently on the same task. We identified six coupling styles in Study 2; of these, we consider the first three (identified with round parentheses) to be "working together."

- **(SPSA)**: (Same problem same area): Collaborators are actively working together to evaluate, trace, or draw a route (e.g. one person points at landmarks while the other connects them with a pen). Often, this is accompanied by conversation.
- **(VE)**: (View engaged: One working, another viewing in an engaged manner): The pair is working together, but only one is actively manipulating the display. For instance, one may be showing a route to the other, or one may just be watching the other's actions very carefully. In the latter case, the individual is watching closely enough to suggest corrections. Conversation often accompanies this style.
- **(SPDA)**: (Same problem, different area): Collaborators are working simultaneously on the same sub-problem, but are focused on different parts of the table. For instance, participants may be evaluating alternate solutions of the same sub-problem. This style is not accompanied by conversation. Instead, conversation and gestures often transition groups to more tightly coupled work.
- **[V]**: (View: One working, another viewing): One collaborator is working on the task, and the other is watching, but is not sufficiently involved to help or offer suggestions. The person watching only reacts to high-level activities, such as when the active person stops working or needs resources (e.g. a widget).
- **[D]**: (Disengaged: One working, another disengaged): One collaborator is completely disengaged from the task, not paying any attention to the task or partner.
- **[DP]**: (Different problems): Collaborators are working completely independently on separate sub-problems at the same time. Each person's interactions with the workspace are not related to the other in any way. In this style, participants often peeked at one another to maintain an awareness of the other's activities.

After coding each of the 16 sessions (4 groups × 4 conditions), we ran a set of statistical analyses to understand how coupling related to the study conditions. Total time spent working in a particular coupling style was broken down by study condition. These coupling styles were subjected to a two-way, within-subjects ANOVA with repeated measures (filters and lenses × individual and compromise routes).

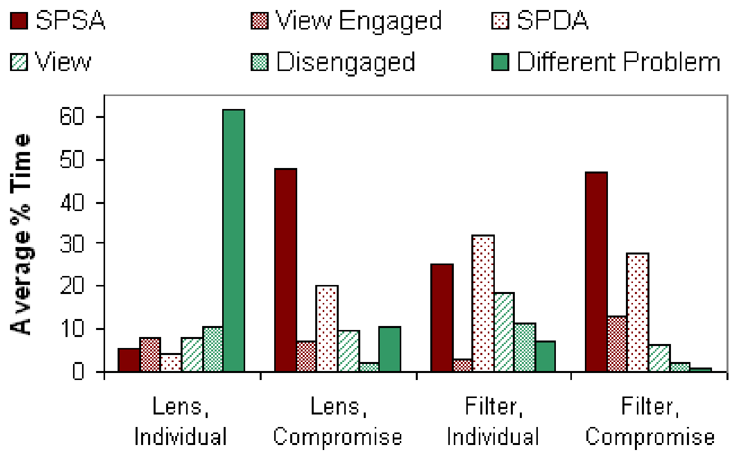

Figure 4.4 Observed coupling styles in each study condition.

Relationship between coupling styles and tool/task type

Consistent with our hypotheses, when creating compromise routes, pairs' were more tightly coupled than when creating individual routes. They also worked more tightly with global filters than with lenses. Figure 4.4 shows the mean proportion of time participants spent working in particular coupling styles as broken down by condition.

The ANOVA revealed a borderline significant interaction between interaction technique (filters and lenses) and route type (individual and compromise) in the amount of time participants spent working on different problems (DP) [$F(1, 3) = 9.5$, $p = .054$, $\eta^2 = .76$]. Additional main effects for interaction technique [$F(1, 3) = 14.3$, $p = .032$, $\eta^2 = .827$] and route type [$F(1, 3) = 14.9$, $p = .031$, $\eta^2 = .833$] were also present for different problems (DP). These effects collectively suggested that participants spent the most time working on different problems in the lens+individual route condition.

A main effect was also present for route type and the amount of time participants spent working on the same problem and same area (SPSA) [$F(1, 3) = 159.6$, $p = .001$, $\eta^2 = .982$]. This indicated that participants spent more time working together on compromise routes than they did when working on individual routes. This was unsurprising given the differences between the two task types.

In all groups, we observed that participants worked independently and loosely coupled on the two problems that could be spatially separated, and then transitioned into more tightly coupled work, working closely on the problem that overlapped in the lens+individual condition. We were surprised by the activity in the lens+compromise condition, where we expected all participants to work in parallel on separate sub-problems. Instead, we found that three groups worked together in this condition about 96% of the time. Group 2 was an exception: they worked in a parallel, independent manner to generate the best individual routes, and later worked in a more tightly coupled manner to find the best compromise based on the individual solutions. They only worked together about 51% of the time.

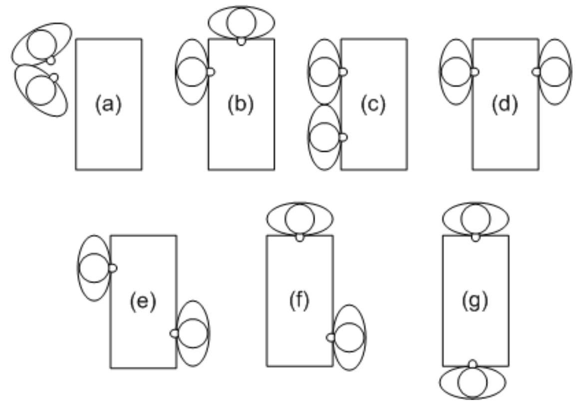

Figure 4.5 We coded seven position arrangements around the table (based on relative positions): (a) together, (b) kitty corner, (c) side by side, (d) straight across, (e) angle across, (f) end side, and (g) opposite ends.

Conversely, we found that participants usually worked together when using global filters. They worked together 79% of the time on individual routes and 94% of the time on compromise routes.

Arrangement

We suspected that with tighter coupling, participants stood physically closer to one another. To examine this relationship, we first video coded participants' location changes, thereby providing arrangement information, and then cross-tabulated this data with coupling style. Our coding scheme (Figure 4.5) considered the relative positions of participants and not their absolute positions.

As we expected, when collaborators worked more closely together, they stood physically closer, and when they worked independently, they stood further apart. This can be seen as a dark diagonal trend from the top left to bottom right of Table 4.2. Although this effect is complicated by the fact that participants were physically closer when working on the same sub-problem, it corresponds

White: < 1% Light grey: 1-4 % Dark grey: > 4%	SPSA	VE	SPDA	V	D	DP
(a)Together	7.8	1.6	3.4	0.5	0.2	0.5
(b)Kitty corner	9.4	1.9	5.2	2.4	0.9	1.9
(c)Side by side	2.5	1.0	2.3	0.9	0.9	3.1
(d)Straight across	9.2	2.3	8.7	3.3	2.3	1.0
(e)Angle across	3.8	1.4	2.4	2.3	1.4	6.2
(f)End side	0.5	0.1	0.1	0.3	0.3	4.9
(g)Opposite ends	0.0	0.0	0.0	0.1	0.0	3.1

Table 4.2 Percent time working in each coupling style and physical arrangement. Arrangement categories are in increasing order of average distance between participants. Coupling styles range from working closely together (left) to working independently (right).

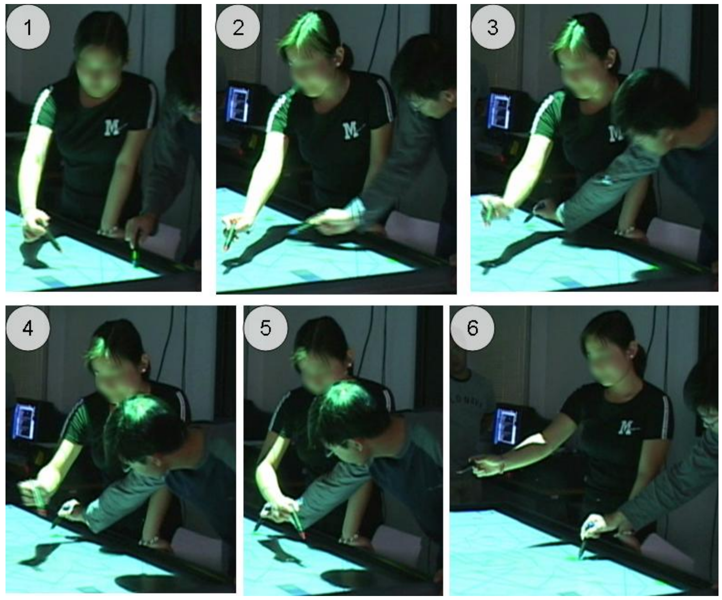

Figure 4.6 A series of frames representing a particular instance of interference with a loosely coupled pair (note how they are focused on different areas of the workspace in frame 1).

with results from our first study, which did not have spatially separated sub-problems.

A notable exception to this observation is that Side by Side arrangements were physically closer than Straight Across, yet Straight Across was a very common arrangement for group work. This result is likely the consequence of the particular collaborative ergonomics of our table: working Straight Across the table yielded a good position to work on the same problem while providing smooth face-to-face communication. Similarly, when working independently, standing Side By Side rather than Straight Across may have reduced visual distraction.

Consistent with prior work (e.g. Scott et al., 2004; Tse et al., 2004), physical positioning appeared to be related to territorial behaviour. Individuals tended to explicitly interact only with areas physically close to them, and avoided interacting with areas physically close to their partner (an exception is shown in Figure 4.6). Yet, these "territories" were transient. As individuals moved, others were no longer restricted from operating in those areas. Similarly, when a pair worked closely together on the same problem, we often observed one person taking on the other's perspective. In these instances, the second person would never displace the first: even if the second person was to gesture toward the table, he would move to a different location around the table before doing so.

Handling interference

We also saw many instances of interference, where one collaborator either blocked another's view or ability to physically interact with the workspace (Figure 4.6). When more tightly coupled, interference was less frequent, and was handled more gracefully, with one person moving out of the way just as another moved into the space. When collaborators worked in a loosely coupled fashion, we saw more frequent instances of one participant waving the other away, and in some cases, physically pushing or grabbing each other.

Not all interference was detrimental. While interference often interrupted independent work, interference (intentional or otherwise) often signaled or aided groups to transition to more tightly coupled working styles. For instance, any activity covering workspace (e.g. one person counting or drawing) often signaled that one collaborator "had a good solution." Since many pairs liked to validate their routes together, this activity would act as an invitation to closer collaboration (i.e. tighter coupling).

Discussion

The original observations we made in Study 1 were largely consistent with the results of Study 2. Participants preferred tools that supported their particular working style. Despite explicit independent and group tasks in Study 2, participants generally preferred to work together when they had joint responsibility for the outcome of the task, as was true when they were constructing compromise routes.

From Study 2, we identified six collaborative coupling styles to describe the workspace activity. These styles were closely related to other factors such as physical arrangement, the task and tool being used, and the incidence of interference. The coupling descriptors are useful in that they contribute to the description of significant aspects of group activity by characterizing the nature of mixed-focus collaboration: namely, that groups frequently and fluidly transition between several stages of working closely and working independently.

General Discussion

In studying participants as they work over tabletops, and in attempting to communicate those observations to a wider audience, we have come to appreciate that collaboration is a highly complex and multifaceted construct—even when we constrain the investigation to real-time tabletop collaborative knowledge work within a task in a study. The term "coupling," and associated terms "tightly coupled" (that entities work closely together), and "loosely coupled" (that entities work fairly independently), for example, have been used to describe one dimension of collaboration. Yet we have found that there exist many points along the spectrum between the two endpoints, and suspect that the nature of collaboration may vary in other orthogonal dimensions.

We describe six collaborative coupling styles, though we do not believe this is an exhaustive list. The styles we observed were likely limited by the specific parameters of our study. For instance, if we had limited ourselves to the methodology of Study 1, we would not have seen the extent of varied styles of independent work that we described above. We expect that additional styles may be uncovered through studies with different user groups, tools, and tasks.

Although we attempted to order the coupling styles from tightest to loosest, the detailed ordering of all styles is not necessarily obvious or finalized. We are unsure whether these coupling styles even fall along a single dimension. For instance, same-problem-same-area, view-engaged, and same-problem-different area appear to fall along a continuum of "degree of involvement in the other's task." However, it is not clear whether coupling is tighter when working on different problems or when one person is disengaged. To account for these observations, a promising approach may be to consider how collaboration might be described as a dynamic and fluid stateless system (Hancock & Carpendale, 2006).

Nevertheless, these collaborative coupling styles accord with the user model of transitions that we developed in Chapter 1. In particular, the tabletop studies helped elucidate different types of configuration patterns in the context of synchronous collocated collaborative tabletop activity (coupling styles). As illustrated in Figure 4.7, these were the primary driver in most of the transitions that we saw in this chapter. To a lesser extent, one might consider there to have been subtle changes in task (depending on the experimental trial condition), or changes in group members (e.g. independent vs. shared activity); however, these would likely be considered as secondary (or lesser in magnitude) compared to the major transitions we saw as a function of changes in coordination patterns. In direct comparison to the whiteboard studies, we see that the artefact largely remains the same (even though we had hoped that the mechanisms we provided would help smooth transitions).

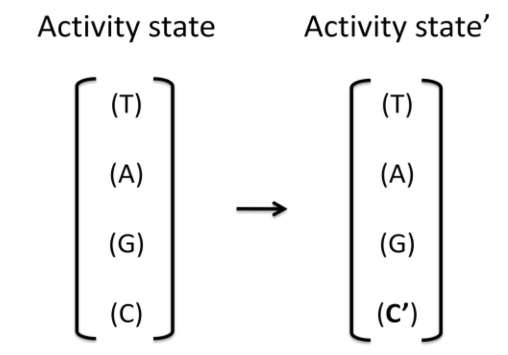

Figure 4.7 The transitions we saw in the tabletop studies were primarily changes in configuration pattern (coupling styles).

Limitations

Because our interest in this study was in mixed-focus collaboration, we chose an observational laboratory study method, which allowed us to carefully construct the task so that users would need to explicitly move between independent and shared activities. The specific tasks that we designed lack ecological validity in the sense that they require multiple shifts in a very short period of time, and that they were perhaps of an "unrealistic" scenario. However, we have maintained the core defining feature of such activities by allowing users to manage their work (i.e. the tools did not mandate a particular work process), meaning they were free to shift between modes of work as they chose. This method is consistent with prior work examining tabletop collaboration (e.g. Scott et al., 2005).

We are limited in our ability to generalize (strictly speaking) to the nature of mixed-focus collaboration beyond pairs, however. Studying collaboration in of itself is tricky: groups tend to behave in extremely idiosyncratic ways, and often, their group dynamic overrides treatment effects of tools. For instance, dominant members of a group will often override the suggestions or actions of other, more submissive members. Our choice in studying pairs rather than larger groups was deliberate, and mainly as a simplification mechanism. Nevertheless, we suspect that while the coupling classification may not be directly applicable to (for example) a group of three, the model that we describe of users' changing engagement with one another would still be valid in larger groups.

Implications for Tabletop Design

Groups move frequently and fluidly between many styles of coupling. Each style is accompanied by different behavioural mechanics. For example, in tight coupling, individuals work in close proximity even when the workspace has no implicit orientation. Furthermore, different display techniques support different styles: independent views support individual work because they reduce interference, and global views support group work because they provide common ground.

1. Support a flexible variety of coupling styles.

Mixed-focus collaboration encompasses many coupling styles: even in our own studies, there was a wide variance in the approach groups took. Most systems fail to provide support for multiple coupling styles, falling back on social protocols to effect different coupling styles (e.g. Kruger et al., 2003; Ringel et al., 2004; Scott et al., 2004). Since tabletop displays are dynamic, we can provide a variety of tools to support different coupling styles. A promising future direction may be the convergence of physical and digital media on tables as an alternative means for providing individual and group views.

2. Provide fluid transitions between coupling styles.

Supporting mixed-focus collaboration requires supporting the *transitions* between loosely coupled independent work and tightly coupled group work (Baker et al., 2002). Providing only a single view of the workspace limits individuals' abilities to work independently (Gutwin & Greenberg, 1998), yet using separate copied workspaces may prevent many group collaborative dynamics, such as being able to see what others are doing, from emerging (Scott et al., 2003). Our results do not suggest mitigating interference altogether since some forms of interference signal transitions between coupling styles and benefit group coordination. Furthermore, the recognition of interference can be used as a means to fluidly transition between coupling styles. For instance, the act of rotating an object toward a fellow collaborator temporarily signals the desire for attention (Kruger et al., 2003): the system may also use this cue to transition the workspace to match the tighter coupling.

3. Provide mobile high resolution personal territories.

The interference we observed was a direct result of individuals' desired working areas overlapping. Creating usable and useful personal territories could take several avenues, including a higher resolution workspace, or mobile regions of high resolution, or even using distinct displays for personal work (such as Tablet PCs or PDAs).

4. Support lightweight annotations.

Tabletop task spaces should support mobile, unobtrusive, and transient annotations. One of the affordances of the tabletop form factor is the ability to conduct independent work unobtrusively (Tang, 1991). Annotations help to generate and track independent work, and may be moved to be shared with the group (Kruger et al., 2003). In our studies, participants surprised us by frequently annotating the map space with both spatially-relevant and spatially-invariant annotations. These annotations sometimes helped and other times hindered the other participant. We recommend supporting the easy creation, mobility and modification of annotations.

Section 1 Conclusions

These studies illustrate the complex nature of users' interactions with one another and with a shared interactive display surface during collaborative knowledge work. In particular, we have developed a model for this behaviour, describing a set of coupling styles, and how users transition

Figure 4.8 The MAGICBoard only comprises a small space in the overall deployment location (d), and bystanders comprise the majority of individuals near the display (a), (c). Only a single user is actually making use of the display (b).

between them as necessary to accomplish the task at hand. In relation to the transition model described in Chapter 1, we have demonstrated that changes in coordination patterns (i.e. coupling styles) are related to other variables such as preferred tools, physical arrangement, and the incidence and handling of interference. For example, groups use tighter coupling styles when working together closely, preferring common, global views. By developing this model, we have provided design opportunities for tabletop researchers to support collaborative knowledge work. For example, we have further developed this model in the context of collaborative visual information analysis (Isenberg et al., 2008). Others have explored how the artefact can be manipulated to help smooth these transitions in collaborative visual analytics applications (e.g. Isenberg & Fisher, 2009; Tobiasz et al., 2009).

We now turn to our second case study, which investigates a completely different context. Whereas the first case study explored collaborative knowledge work in a closed laboratory study (perhaps mimicking a meeting room scenario), the second system, MAGICBoard was deployed in a public space as an ambient display. As outlined in Table 4.1, MAGICBoard differs from the tabletop system on a wide variety of factors, including interaction mechanics and user-to-user relationships. Of particular interest is in how the design and deployment of this system helped reveal a pattern of bystander transitions: in particular, that as bystanders transition from passers-by to contributor, their information needs change. This transition is of particular importance when attempting to engage bystanders in an open system.

Section 2: Design and Deployment of MAGICBoard: an Interactive Public Large Display

Large public displays are typically used for broadcasting a stream of location-relevant information, but most deployed displays of this nature are not yet interactive. This lack of interactivity may change with the increasing proliferation of high-power handheld devices (mobile phones, PDAs, MP3 players), which enable new forms of use (e.g. Fass et al., 2002; Huang et al., 2006; Vogel & Balakrishnan, 2004; Myers et al., 1998). Despite the emergence of new technology that could allow users to interact with large displays, past research has found that motivating people to interact with these displays in a public space remains a real challenge (Brignull & Rogers, 2003). An oft-cited deterrent is the potential for social embarrassment when interacting with a public display.

In designing MAGICBoard (shown in Figure 4.8), a public digital forum, we sought to address this challenge by using SMS messaging as the primary means of interaction with the large display, thereby allowing users to interact with the system from the privacy of their own personal devices—a concept we call *supporting covert engagement and interaction*. The core functionality of MAGICBoard was simple: users post text-based items on the display, which persist until newer items pushed them off-screen. In designing this interactive display application, we found that many of our design choices ultimately focused on individuals who might not be actively engaged with the display itself: *bystanders.*

We situate our work in the context of using public displays as social catalysts—or artefacts/events that focus the attention of diverse inhabitants (Karahalios & Donath, 2004). Brignull & Rogers (2003), in studying people's activity patterns around a similar large display applications, described three classes of users based on their patterns of activity: (i) those engaging in *direct interaction* with the large display; (ii) bystanders whose activities indicated a *focal awareness* of the display, and (iii) bystanders whose activities implied a *peripheral awareness* of the display. The tabletop system described earlier in this chapter would be an example of a system whose users would be engaging in direct interaction with the large display. The bystanders of Brignull & Rogers' framework arise only in the context of ambient display applications, or public installations. To motivate these bystanders to interact with the system, Brignull & Rogers advocate designing applications to support *transitions* between these thresholds.

The findings from this research support this conceptual framework, and we develop a model of bystanders that considers both their information needs, and how to support their transition from a bystander role to a contributor role.

In this section, we first describe MAGICBoard and its deployment, which allowed us to investigate and categorize different *types* of bystanders. From there, we re-examine several design heuristics from Huang et al. (2006) and arrive at three thematic design implications to support bystanders' use of public displays: *supporting graduated proximal engagement*, *lowering barriers for interaction,* and *supporting covert engagement and interaction*.

Figure 4.9 The MAGICBoard's two displays have different functions. The right display is intended to be viewed from a distance, and functions as the "overview." The left display is the "detail" display, and intended to be viewed up close.

MAGICBoard: A Digital Public Forum

MAGICBoard is a public forum for trivial but amusing topics (see Figure 4.9). Two side-by-side projectors present the current topic, the votes and opinions of those who have commented on the topic, and a summary of the votes on the topic. The right display allows passers-by to easily glean the overall opinion of the community on different topics. Interested bystanders can engage with the system by stepping closer to view the comments themselves. They can then interact with the display by either: (1) sending an SMS message from a mobile phone, or (2) using a kiosk next to the display. The kiosk provides a basic form-based mechanism of interaction, and the SMS gateway supports more "private" entry and preparation of content (e.g. Myers et al., 1998; Greenberg et al., 1999).

Figure 4.9 shows each display in action: the left display shows "overview" information while the right display is the "detail view." The overview display (containing the topic and overview of the tallied votes) is intended to be viewable from a long distance: font size is large and viewable from 20 meters. The detail display is intended to be viewed from much closer, and shows the last 16 submitted comments.

MAGICBoard was constructed using the MAGIC RESTBroker, an HTTP-based toolkit intended for the rapid prototyping of large display applications (Erbad et al., 2008). The RESTBroker enables lightweight message passing using state-based channel semantics. The toolkit allowed different parts of MAGICBoard to be built and run on different client machines: the kiosk, SMS gateway, and display application are all completely separate applications communicating through this lightweight protocol.

We deployed the MAGICBoard in a common study/social hallway of the applied science building at our university (Figure 4.8). This corridor is a common area with a small coffee shop to the side, and a small alcove where students frequently meet to study. The two displays themselves measure about 6m × 2m and were positioned to be visible from the front door of the building throughout the day.

Our interest in MAGICBoard is unique from prior work in two respects: first, our focus on SMS interaction enables participation by users who might otherwise not partake due to the potential for *social embarrassment*, and second, MAGICBoard was deployed in a public setting with bystanders who are unlikely to know one another, whereas prior work frequently deployed such displays in social event settings (e.g. Brignull & Rogers, 2003), in a distributed setting (Karahalios & Donath, 2004), or in contexts with known users (e.g. Brignull et al., 2004; Huang & Mynatt, 2003).

Design Lessons from Deployment

We deployed MAGICBoard for a week near the beginning of the school year, collecting field notes, photographs and video of users and bystanders making use of and observing the display. In this exploratory study, our goal was to observe users, and to catalogue and understand their interactions with the MAGICBoard. We use the term "interaction" here in a broad sense: while we were interested in users that made active use (i.e. contributed) to the display itself, we were also interested in how other, more "passive" users, simply behaved around the displays.

To gather observations, we would use time sampling, taking note of different types of users in the space, and their behaviours around the display at different points within the hour during the deployment. We would also take special note of interesting events or behaviours around the display outside of these time slices. To organize our observations, we used a method of provisional verification, where we would iteratively theories and models to describe the data, and then further collect data to determine whether the model was appropriate (Strauss & Corbin, 1999).

This method of data gathering and analysis was necessary, and consistent with prior exploratory work that has examined how public interactive large displays are appropriated (e.g. Brignull & Rogers, 2003). A specific challenge was that we wanted to ensure that our presence did not impact users' reactions to the display itself; consequently, we could not approach users explicitly to ask for their feedback about the display.

We report the most salient observations from our study relevant to design here.

Classifying Three Types of Bystanders

Our interest in bystanders began during the design stage of MAGICBoard in our discussions with our focus group (comprised of primarily engineering and computer science undergrads): What would someone see on the large display? How would one understand what was going on? How would one interact with the display? How would one know *how* to interact with the display? It became clear that our design focus, which typically centers on "users"—those already interacting with the display, needed to be balanced with an equally concerted focus on *bystanders*—potential contributors who may not yet be engaged with the display, but "users" of the display nonetheless.

Figure 4.10 Examples left-to-right of (a) a passer-by, who is en route to another location, and does not linger; (b) a stander-by, who is sitting in the space, and therefore somewhat coincident with the display; (c) an engaged bystander, who is reading the detailed comments and was about to pull out his cell phone, and (d) a *contributor*, who is actively engaged with SMS on his cell phone.

Our initial observations of MAGICBoard's use validated the ideas raised by our focus group, and revealed three different types of bystanders: passers-by, standers-by, and engaged bystanders. We differentiate bystanders based on their behaviour and engagement with the display (illustrated in Figure 4.10).

- *Passers-by* (Figure 4.10a) were *in-transit*, passing through the area en-route to another location. Thus, the amount of time and effort they expended toward looking at the display was extremely limited—those that looked at the display typically gazed for no longer than 10 seconds. And although these passers-by may have glanced at the display, most did not typically stop to interact with it.
- *Standers-by* (Figure 4.10b) were actually *spending time in the environment* itself (akin to those with peripheral awareness in Brignull & Rogers, 2003), be it at a nearby table to study, in the line-up or condiment area of a nearby coffee shop, or simply waiting for someone. While they were not in the environment primarily to interact with the display, they had more time to actually read the content and understand the display.
- Finally, *engaged bystanders* (Figure 4.10c) were interested enough in the display (with focal awareness) that they were actively staring at the display and "making use" of the content on the display.

This classification scheme has strong similarities to those in Brignull & Rogers (2003) and Vogel & Balakrishnan (2004), and supports the notion that bystanders have differing awareness levels (and hence differing information needs) of the display.

Support Graduated Proximal Engagement

Bystanders cannot be expected to be standing near the display: instead, bystanders' proximity to the display is extremely variable, affecting their ability to view the display's content. To support distal bystanders, the first approach might be to increase the size of all fonts; however, this solution is not only a suboptimal use of the display space, it also compromises the possible interactive complexity of the display. Our design approach was to support *graduated proximal engagement* where *the display can be engaged with from a variety of distances*. This design approach assumes that one's proximity to the display correlates with one's interest with the display, and aims to “reward” users for being closer to the display by providing those users with an improved experience.

- *From far away (20m),* users can see and make out the topic question (and associated picture if present) on display. Graphics summarizing the votes also show that there is a vote going on, even though it is unlikely that the details of the chart is visible from such a distance. These large visuals are intended to provide awareness of the display’s purpose to *passers-by*.
- *From closer (10m),* users can make out the details of the summary charts to see the opinion of the community on the topic. Further, it is possible at this visual distance to read the last comment that was made (presented in bigger font). It is clear from this distance that comments have been posted on the display; however, one cannot read these comments. *Standers-by* capable of reading this information can make a decision about whether to engage with the display further.
- *From up close (5m),* all content on the display is visible. At this point, the user can read all of the detail on the display, and in particular, see the comments of prior users of the display and instructions on how they can vote and comment. Our hope is that *engaged bystanders* will become contributors when they are close enough to see all of this content.

Although we realize this concept of graduated proximal engagement by varying the size of visual elements on the display, it should be emphasized that rewarding users for transitioning one type of bystander or contributor to another can occur in a variety of ways. For example, Brignull & Rogers (2003) “rewards” users close to the display by providing them a method of interacting with the display. Similarly, Vogel & Balakrishnan (2004) provide increasingly personal and explicit interaction for users of ambient public displays based on their tracked proximity to the display.

What is pertinent to our discussion here is that that these different types of bystanders/users have different needs of the MAGICBoard—their *tasks* involving the MAGICBoard are different. In relation to the transition model introduced in Chapter 1, these changes are therefore transitions where the *task* changes. As illustrated in Figure 4.11, group membership stays the same (i.e. it is typically a single individual), and the configuration patterns at work (asynchronous) remain constant. In the MAGICBoard, the artefact also stays the same. Work described earlier by Vogel & Balakrishnan (2004) examines how the artefact can also change to facilitate the transitions; however, in the MAGICBoard, the artefact stays the same.

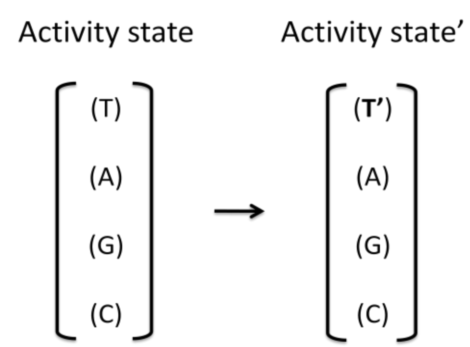

Figure 4.11 The MAGICBoard allowed us to explore transitions where the task changes.

Lowering Barriers for Interaction

Because large interactive public displays are uncommon, *bystanders may not be aware that they are able to interact with the display*. Beyond this initial knowledge barrier, there is the problem that *bystanders may not be aware of how to interact with the display* and also that *users may be embarrassed to use the display* (Brignull & Rogers, 2003). Once bystanders have overcome these two barriers, and have begun interacting with the display, we are faced with the usual problem of *providing feedback in a timely and meaningful fashion*. In consideration of these issues, we focused on providing knowledge and mechanisms to lower barriers to interaction (Huang et al., 2006b). This theme raises the design tension between lower fidelity input *vs*. feasibility of complex interactions with the display.

It was important to communicate to bystanders how to interact with the system. Thus, our instructions were designed such that from a medium distance, one could see a cell phone as a cue that the display had something to do with cell phones. We felt that from this cue, interested bystanders could decide to approach the display, thereby becoming *engaged bystanders*; thus, the instructions could be placed in comparatively small font.

Since SMS is already widely used, we chose to support interacting with the display using SMS messaging from the phone rather than another input mechanism (e.g. web-based forms, downloadable mobile applications, etc.). The trade-off here is clearly evident: we chose to lower the barrier of entry to mobile phone users to increase the number of potential users, but in so doing, sacrifice rich interaction possibilities (e.g. Fass et al., 2002; Vogel & Balakrishnan, 2004). We also provided a form-based interaction mechanism with a laptop right at the display, and we briefly discuss its impact on participation patterns in the next subsection.

Support Covert Engagement and Interaction

Many authors have suggested that a core deterrent to users making use of large public displays is the potential for social embarrassment (Brignull & Rogers, 2003). This is likely to occur for several reasons: (1) the display is large, so actions (and errors) are made more obvious to others (compared to a laptop-sized screen); (2) it is likely the display employs an obvious input device (so users are easily identifiable), and (3) it is likely the display system employs novel or one-off software (so users are unfamiliar with how it behaves). Thus we suggest *supporting covert engagement and interaction* (though not necessarily exclusively) to draw in curious onlookers who may be understandably shy.

With MAGICBoard, we support this covert interaction using SMS messaging from users' mobile phones. In general, however, this "covert interaction" approach introduces two new design tensions: the problem of feedback *vs*. identifiability, and the problem of learnability *vs*. privacy. The first problem is providing users, who may be dealing with a novel interface (as they were with the SMS mechanism), with feedback in a timely and relevant fashion without revealing their identity. We address this issue by showing only part of the user's semi-unique phone number on the display itself, using a dedicated "Most Recent Post" area of the display to highlight recent contributions (Figure 4.9, left), and by responding to users' contributions with a text message in return. This SMS response was direct, and "in-context"; any errors would not reveal their identity to the public.

Many authors have observed that bystanders often learn how to use a large display because it provides useful *feedthrough* of interaction (e.g. Brignull & Rogers, 2003; Huang et al., 2006a). Clearly, this mechanism for learning is lost with covert interaction. We address this problem by providing easily visible instructions and a straightforward interaction mechanism. Vogel & Balakrishnan (2004) provide a video of an actor on the display itself to show bystanders how to use the display.

Nevertheless, the covert interaction mechanism (SMS messaging) produced visibly different participation patterns compared to the overt interaction mechanism (the laptop). Parallel to Brignull & Rogers (2003), the laptop tended to produce a "honeypot effect", drawing in other bystanders when users made use of it; however, users making use of their cell phones to interact with the display tended to leave longer, more thoughtful messages.

Initial Observations

Our deployment also produced a number of interesting observations about users' behaviour around the display. We report on those here.

SMS users seemed more engaged than kiosk users

One surprising observation that came to light was that SMS users typically entered more content than kiosk users. Based on server logs, SMS users keyed more characters and words, and clearly seemed to more carefully craft their contributions to the large display compared to kiosk users. There are likely several reasons for this type of behaviour. First, SMS users have more time to think about and compose contributions to the display because they do not necessarily experience the same *social embarrassment* as those users at the kiosk (who are, in contrast, very visibly interacting with the system). Second, SMS users are likely more committed to contributing to the system because they actually invest effort into retrieving and setting up their own devices. We would expect this to be true if more users made use of the lower-barrier kiosk, and indeed, we saw a 5:2 ratio of kiosk to SMS users. Third, there is some reason to believe that the personal device is simply more conducive to reflective thought compared to a visibly public input device.

Allow relaxed SMS interaction

The core difficulty of using SMS is the relative lag between submission of an SMS message and response by the system. This lag is imposed by the device (via menu systems, for example), and potentially in bottlenecks of the network service. Nonetheless, this lag suggests that user input via SMS should be somewhat lengthy (thereby making up for the lack of responsiveness by providing a long stream of input at once), thus implying a user's interaction with his/her SMS device is also somewhat lengthy. Ironically, it is this lengthier interaction with one's own SMS device that makes it likely that there will be "formatting errors" in the resulting input stream to the large display.

We suggest designers use a relaxed syntax when using SMS interaction for two reasons: (1) it is already difficult to contribute via SMS, and (2) rejecting a user's initial interactions with the system can be devastating.

Although we initially provided mechanisms to provide users with feedback on how to correct their contributions (via an SMS error message), we later simply relaxed the "formatting requirements" of

SMS contributions. Thus, ill-formed SMS contributions were simply shown on-screen, thereby providing users with positive feedback that their contribution was valued. Better approaches may be to interpret users' SMS strings, and to infer intended commands.

Kiosk users garner more attention that SMS users

Akin to Brignull & Rogers' observation of a honey-pot effect around the keyboard (Brignull & Rogers, 2003), we found that bystanders more frequently congregated around a kiosk once a user was standing and making use of the kiosk. This effect was extremely noticeable, and users therefore seemed to appear in groups around the kiosk before disappearing. In contrast, we only prominently noticed one SMS user that clearly had a group gathered around him. It is difficult to say whether this effect was difficult to detect because we did not know where SMS users were interacting from, or whether it is an effect of the input device itself.

Regardless, it seems likely that bystanders are more likely to be interested in what a stranger is doing at a public input terminal versus a stranger using an SMS device.

Passers-by are unlikely to participate

As we alluded to earlier, passers-by are typically goal-directed in the sense that they are en route to a location or task. Thus, while many passers-by clearly gazed intently at the display to interpret it, they did so while continuing on in the direction they were headed—that is, passers-by had no intention of stopping. It is unclear whether these passers-by did not participate because they: (1) were unaware that the system was interactive; (2) were unaware of how to interact; (3) were not interested in interacting, or (4) simply had no intention of stopping while in transit. Given the number of users who were able to make use of the display, and the fact that some passers-by did stop to engage with the display, the first three possibilities are put to question.

Regardless, it should be clear that there is another threshold that needs to be overcome from passer-by to stander-by. This threshold may not have been detected in the past (e.g. Brignull & Rogers, 2003), because displays intended for the "public" in these contexts were deployed where all bystanders were standers-by by virtue of the setting (e.g. at a party). In subsequent deployments we investigated additional approaches to encourage this transition (Finke et al., 2009).

Discussion

In the research literature, interactive ambient displays have typically been deployed in semi-public environments, where users are somewhat well-known to one another. In these scenarios, these awareness displays have typically functioned as a type of *social catalyst* (Karahalios & Donath, 2004)—artefacts intended to stimulate social engagements between individuals. McCarthy (2002) designed and deployed a range of such displays, focusing on the ability of such displays to provide appropriate information for those attending to the content, and on how interaction with such displays may occur. Karhalios & Donath (2004) focus more on the interactions that may occur around such artefacts across distributed sites, displaying deliberately abstract/ambiguous representations of remote sites. Churchill et al. (2004) and Huang et al. (2006b) employ ambient awareness displays to connect disparate research teams using web-based forms and email as contribution mechanisms.

Derivatives of these ideas have also appeared in lab environments, where the smaller set of dedicated core users facilitate the use of far richer interactions. For instance, Messyboard (Fass et al., 2002), Dynamo (Brignull et al., 2004) and Notification Collage (Greenberg & Rounding, 2001) are large shared displays that allowing individuals to post information snippets from their individual devices/clients. In addition, they facilitate the transfer, display and manipulation of multimedia content.

The particular focus we bring here, however, is on the use of ambient displays by individuals who are largely unknown to one another. In this sense, the deployments of Blueboard (Russell et al., 2002), Opinionizer (Brignull & Rogers, 2003) are more closely related to the work presented here because they focus on the social aspects of the interaction between strangers. Russell & Sue (2003) point to several notions of social behaviour around such displays: learning interaction through observation, learning etiquette around such shared displays, and turn-taking. In the particular context of BlueBoard, many of these issues arise because of the need to be within close physical proximity to others when interacting with the display. In a similar way, interacting with the Opinionizer necessitates being in a focal location (at the keyboard), drawing attention to the interactions between these users.

With MAGICBoard, we explored a different tact by allowing users to interact from (technically) anywhere via a wireless link to the board (through SMS). In this sense, some of the notions of physical and control etiquette and behaviours discussed by Russell & Sue (2003) are no longer relevant; instead, the classifications and discussions of bystanders from Brignull & Rogers (2003) and Vogel & Balakrishnan (2004) are more pertinent to design. These works are interested in casual "use" of large displays by bystanders, and the present work brings additional attention to that issue.

The deployment of MAGICBoard and our design process has produced some interesting research contributions: first, we have developed a model of bystanders, identifying their needs with respect to the display; second, we have described the behaviour of bystanders around such a display with regard to two different input mechanisms; finally, we have articulated a set of design tensions and themes for design of similar types of systems.

General Discussion and Summary

Our goal at the outset of this chapter was to investigate the large display design space in order to outline models for how users' needs change with large display applications. Our investigation of two systems allowed us to explore different kinds of transitions as defined by our transition model in Chapter 1. In the case of collaborative coupling, or changes in coordination patterns, we see these changing needs as a matter of course for collaborators working on mixed-focus collaboration tasks: they will necessarily need to engage and disengage with one another. In the case of MAGICBoard, we saw not only that different users had differing information needs simultaneously, but that some users would transition from being a bystander to becoming a contributor—that is, their *task* changes.

Yet, even with this insight *designing to support transitions* may be ambiguous. How should this support be provided? Should it be automatic, or should it be user-controlled? Should the support be explicit? Or should it be subtle? To summarize, Chapter 3 provided evidence that traditional large surfaces supported multiple tasks—often simultaneously fulfilling several roles. In this chapter, we have demonstrated that users' needs are dynamic: that is, their interaction and information needs change as a function of their interaction with the display. The challenge, it seems, is matching users' changing needs with the appropriate application functionality.

In the next chapter, we design a tool that supports transitions between tasks in a given activity using view changing. This tool is intended to support scheduling in the context of an electronic whiteboard, and interprets electronic ink so that the ink can be viewed in different ways by the user. We provide multiple views of the interpreted data, allowing the user to easily toggle between them to support related tasks. If our theory about transitions is correct, then users will toggle their view of the data so that the view matches the particular task they are trying to accomplish.

Chapter 5

Enabling Transitions Using View Changes

In this chapter, we explore how transitions can be supported in an electronic whiteboard-like application that has a scheduling component. Our goal was to provide a proof-of-concept that would support transitions between tasks as suggested by our findings in Chapters 3 and 4. The three sub-goals of designing and building this system were: first, to demonstrate one mechanism to support transitions in a large display application (i.e. changing the visual representation of data), second, to illustrate that users would be able to use this mechanism easily as they completed tasks, and third, to show that it is feasible to construct these alternate visual representations based on ink interpretation.

We focus on the set of tasks surrounding the use of a whiteboard to schedule contractors in a hypothetical home renovation project: for instance, planning the activities of the contractors, re-reading the schedule to resolve conflicts, communicating the schedule to others verbally, and or sharing the schedule visually with others. The first phase of our research was to design and study a paper prototype of a system that would allow users to dynamically plan events, and toggle between views of this scheduling data to review information. Our general hypothesis was that different views of the scheduling data (e.g. a calendar view, a list view, or a timeline as in Figure 5.1) would be useful for different types of tasks involving the schedule data, and that users would actively toggle between them as they transitioned between tasks in the overall scheduling activity. We treat the study as a set of design sessions where participants were given tasks to complete, with our goal being to learn as much from their behaviour as possible. This method allowed us to validate the underlying principle that view changes were useful in allowing users to transition between tasks involving schedule data.

In the second phase of our research, we design and implement a prototype electronic whiteboard scheduling application based on the findings from the first phase. In this system prototype, we allow the use of unstructured electronic ink, and show how the use of simple rules can be applied to provide structure to that ink, thereby facilitating the generation of different views of schedule data.

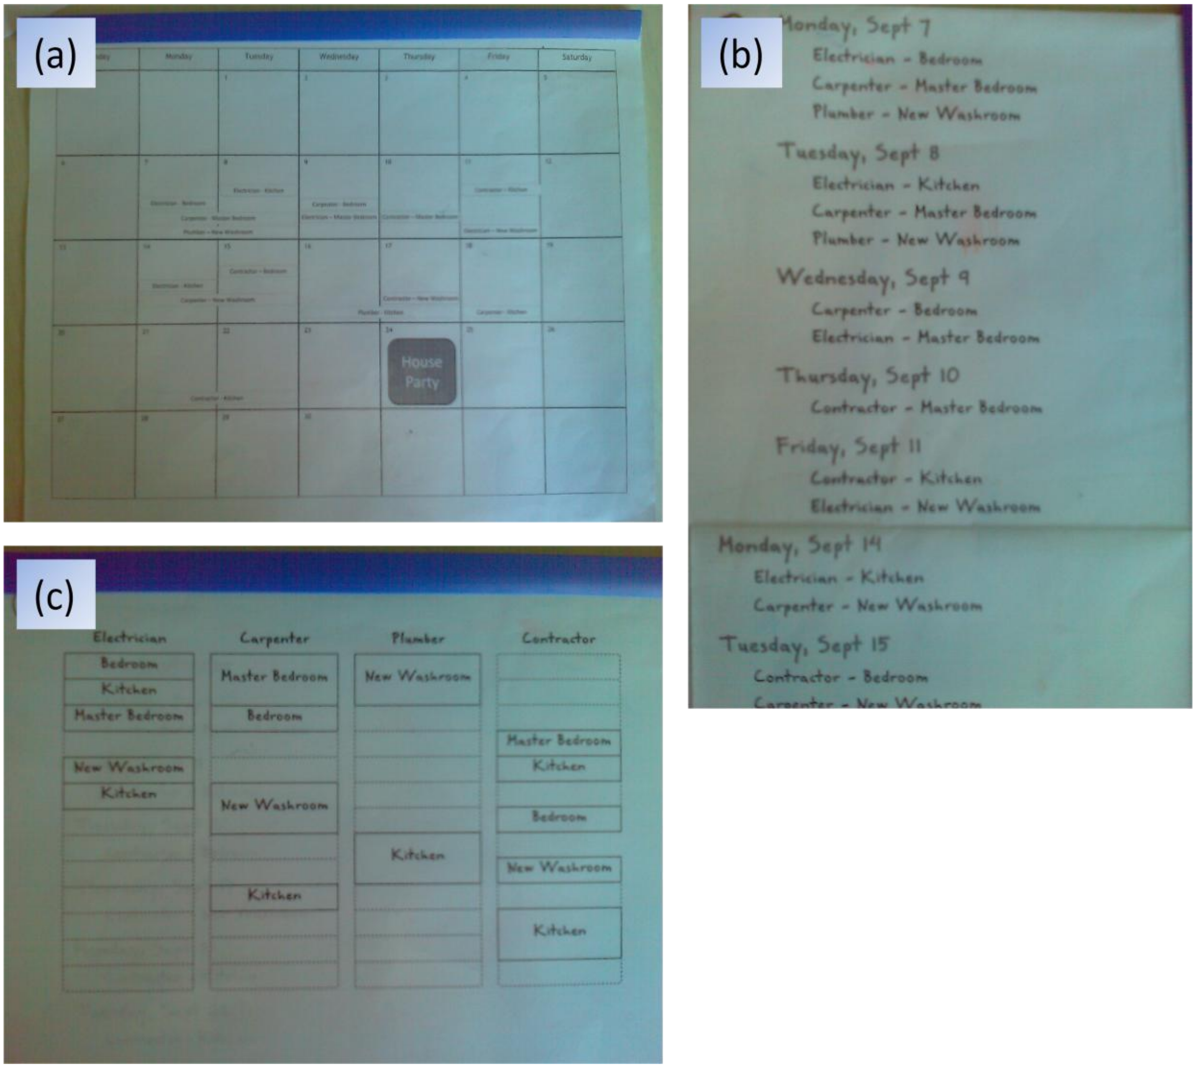

Figure 5.1 The three views of scheduling data we used in our study. All views represent the same data: (a) traditional calendar view, (b) agenda/list view (slightly cut off), (c) "timeline" view similar to that constructed by some participants.

The work here shows how *enabling the view changes on interpreted electronic ink* can be used as a mechanism to support these transitions on large display applications. In effect, we are proposing that different *views* of data may be more or less appropriate depending on the user's needs during a given activity: for instance, when attempting to get an overview of the schedule data, the calendar view might be more appropriate, while when attempting to find scheduling conflicts, a timeline view is more useful. Consequently, allowing users to easily change their view of data better facilitates their at-the-moment needs.

In terms of the transition model outlined in Chapter 1, we are attempting to pair changes in *task* with changes in the *artefact* with explicit, user-initiated transitions. In general, of the four components of activity state in our model, as designers we really only have control over the artefact (i.e. the application) itself. Our goal in this chapter is to illustrate how a designer might apply this idea in practice in the context of an electronic whiteboard-like application for large displays.

Many commercial systems already employ this principle of providing different views of data. Microsoft Word, for instance, provides four such views of text documents: print layout, web layout, outline, and draft. Similarly, many tools have been designed to visualize the content of hard drives in various forms (e.g. as a TreeMap, pie chart, treeview, or as a list). Each of these views supports different tasks with the underlying data, but the ability to toggle between views allows users to show important or hide irrelevant information at the time. Thus, the concept of providing different views on data has already been demonstrated to be useful. Here, we are tackling the concept from the perspective of everyday data that we might conceivably observe on a large display, adopting views that one would find in a conventional personal information manager application (e.g. Microsoft Outlook).

Prior work such as DENIM (Newman et al., 2003) and Cognotor (Foster & Stefik, 1986) similarly provide these transitions, but the research goals and context are slightly different. In DENIM, the system provides a zoomable canvas, and so view changes are effectively different "levels of zoom." Here, we interpret and generate a different view for the electronic ink in-place. Cognotor also provides transitions between the different tasks (brainstorming, organizing, evaluating of ideas), but these are done on highly structured data. Where our interest in information presentation differs from this prior work is that we demonstrate that the idea can be accomplished using ink as the primary interaction mechanism, whereas in prior work, structured input is requisite. Here we demonstrate that digital ink can be interpreted to support these view transitions, and the utility of the concept in a broader application.

We conclude this chapter by reflecting on our design process, the results of our study, and some of the challenges of attempting to infer structure from inherently unstructured data. In so doing, we suggest how the design ideas from our original prototype can be iterated upon, and we reflect on avenues to address the many difficulties of designing to support transitions.

Phase 1: Formative Study of View Changing to Support Transitions

We designed and conducted an initial formative study to explore whether users would use view changing using a paper prototype system. Our goal in designing this study was to develop an understanding of the utility of view changing given a semi-realistic task, and to elicit feedback on how such a system should be designed. To address this need, we adopted a paper prototyping methodology (Snyder, 2003).

Paper prototyping is commonly used as a technique to elicit design feedback in a user-centric design process (Snyder, 2003). Using the paper prototyping technique, a low-fidelity prototype of the interface is created using traditional materials like paper, tape, markers, and so forth. The interfaces are roughed out (and often simply drawn) so that they have a decidedly informal look. For example, a typical Windows desktop screen might be mimicked by a simple collection of post-it notes placed on a tabletop. Interactivity with the system is provided using a wizard-of-oz technique: the experimenter interprets the actions of a user (e.g. a user may "click a button" by tapping on it with his/her finger), updating screen elements (e.g. placing another post-it note on the table to represent a pop-up dialog, or updating one of the post-it notes to reflect some interaction), and so forth. Because paper prototypes do not have to deal with the particular constraints of a

Figure 5.2 Materials used for the paper prototype study. On the left are the timeline, calendar, and agenda views. These were contained in transparencies (typically used for overhead projectors) so that they could be written on. On the right are the post-it notes that were used for UI elements. Not shown are the scissors and tape that were sometimes required. The materials are sitting atop a traditional whiteboard.

working system, the technique is often used as a means to quickly garner user feedback about system design before engineers are committed to building a particular design. Users are apt to provide a great deal more feedback when using a paper prototype (compared to a system prototype) for several reasons: first, the prototype appears incomplete, so there appears to be more room to grow; second, because the prototype is made with such rudimentary materials, it does not feel beyond the reach of users to suggest improvements; third, studies with system prototypes are more subject to demand effects (where participants attempt to "please" the experimenter).

By way of clarity, the paper prototype that we used in this study was a combination of a traditional whiteboard with some paper elements; however, we use the term "paper prototype" for consistency with Snyder's (2003) terminology.

Study Method

The purpose of this study was to investigate our model of transitioning. In particular, we expected that when transitioning between tasks, the ability change the view of the data (to suit the task) would be useful. Our intention was to have participants experience and understand view changes, and to explore whether they would make use of that feature in the paper prototype.

As we expected, users provided us with a great deal of feedback on the design. We immediately made use of this feedback, and iterated on the features of the prototypes as the study progressed. Using this feedback, we were able to address and understand the utility of the view changing mechanism as a means to support transitions on a large display application.

Apparatus

Participants completed two trials: one trial using a combination of a standard whiteboard and calendar (whiteboard+calendar), and one trial using the paper prototype.

With the whiteboard+calendar, participants were given a whiteboard, a set of wet-erase markers, and a calendar that had a thin sheet of plastic laid atop (so that it could be written on and erased).

Figure 5.2 shows the materials used for the paper prototype, consisting of the same whiteboard and calendar, except that it had a number of features that could be accessed. These features were simulated by the experimenter in using a wizard-of-oz technique. Any changes made by participants in the calendar would be reflected in other views so the data was consistent across views:

1. The calendar had two alternate views (close-ups appear in Figure 5.1): an agenda view (which listed the information in a text format), and a timeline view (which visualized the information in such a way to help resolve resource conflicts). Each of the alternate views were still interactive: information added/removed/modified from any of these views would be reflected on other views.
2. The calendar also had the ability to filter the information being displayed. This filtering mechanism was available in each of the other views, as well.
3. Items could be written on any part of the whiteboard, and then dropped onto the calendar using a drag-and-drop-type interaction.
4. Elements of the interface (i.e. the calendar) could be easily replicated.

We added the timeline view (close-up in Figure 5.1c) to the paper prototype condition after running the first six participants. Of these, four generated timeline-like views in the whiteboard+calendar condition, which they told us followed a standard practice in management. Since we were interested in whether users would find utility in having different views of the data, we added this third view for subsequent participants. The original prototype did not have the timeline view.

Tasks

For each trial, participants were given two types of tasks to complete: planning tasks, and review tasks. The planning tasks were always performed first, and the review tasks were completed based on the schedule produced from the first set of tasks.

Planning: We gave participants a modified version of the job shop scheduling task (Tan et al., 2008). In this task, participants are asked to create and optimize a schedule for jobs in a hypothetical hardware shop containing only limited resources (e.g. one saw, one hammer, etc.).

Each job makes use of a subset of the resources for varying lengths of time. There are three primary benefits to this task: first, it is non-trivial, meaning that users need to spend a good deal of time coming to grips with the problem, and then finally solving it; second, it is sufficiently complex to require users to require visual aids—there are several sub-problems that need to be addressed; finally, while the objective of the task is stated at the outset, there are many strategies and approaches to the problem, meaning that users have flexibility in how they approach the task. We modified the original job shop scheduling task's scenario to a home renovation scheduling task to better motivate the study, and simplified the extent of the task (the original job shop task was designed to be completed by groups of users, whereas here, only one user would be performing the scheduling). Jobs were replaced with renovation plans, and the different types of contractors became the limited resources. We also placed constraints on how long the renovations could take place, and so forth.

In the whiteboard+calendar trial, participants were given a list of six renovation tasks, and asked to optimally schedule three such tasks.

In the prototype trial, participants were told that four of the six renovation tasks had already been scheduled (and the prototype reflected this scheduling)—their task was to plan the remaining two renovations. We explain this variation in the Measures subsection below.

Review: The experimenter described an ongoing scenario where various individuals (e.g. the carpenter, the electrician) would request information about the schedule from the participant. For example: "The carpenter is calling you, and asking you: what are the days that I am scheduled for, and where am I working?" This type of question mimics the scenario where the user is trying to find and filter information (on his/her) own before communicating it. Similarly: "Your roommate comes in, and says: I love cooking, so how long is the kitchen going to be under construction?" Here, we address the collocated sharing scenario, where a user shows information to another user. Participants would then respond from their work (i.e. they did not need to work from memory), and in the prototype trial, were still permitted to interact with the system to respond to the questions.

Participants

A total of nine participants (three females) participated in the study. These participants were recruited through on-campus advertising and word of mouth. Participants' experience with computers varied, though all were computer literate. None had extensive experience with tablet-based or pen-based computing, though five reported having tried out a SMARTBoard or a TabletPC in the past. Participants' daily job function included financial advisor, book keeper, receptionist, manager, software engineer, and student; of the students, they had backgrounds in various types of engineering (physical, electrical and civil).

Design

Participants completed all tasks, and we alternated presentation order of the tools (whiteboard+calendar and prototype) between participants.

Procedure

Participants were first provided with consent forms, and explained the purpose of the study. The experimenter then read a script that described the scenario. Participants were provided an information sheet that contained all of the renovation tasks along with the constraints, and walked through one of the renovation tasks to ensure that they understood the constraints of the scheduling task. Using the whiteboard+calendar, participants were then given the tools and asked to complete the scheduling task. When using the prototype, participants were also walked through the various features and views of the system to ensure that they understood how each feature behaved. Participants were asked to use a think-aloud procedure. This procedure was used to help the experimenter understand when participants were having difficulty, as well as to gain insight into their immediate reactions as to the technology's utility. As participants completed the scheduling task, the experimenter helped to identify problems in the schedule (i.e. conflicts), and to provide the functionality of the prototype. Once the participant was satisfied with the schedule, the experimenter again read from a script to ask the participant questions about the schedule itself. After responding to these questions in the recall phase, participants were then asked to reflect on their experience with the task.

The entire procedure took about one hour to complete, and participants were remunerated $10 for their participation.

Measures

In observing participants' use of the whiteboard+calendar, we were interested in the types of visualizations that users would construct on the whiteboard to help support their thinking process. In particular, we were interested in how these constructions encoded information (i.e. the extent to which they were symbolic vs. the extent to which spatial semantics were employed).

When participants used the paper prototype, we were interested primarily in when and how the view transitions were used. We had expected users to transition between the use of ink and the use of the calendar throughout the planning task, and the agenda view when completing the recall task.

Note that we are not measuring speed and accuracy as one would expect in a controlled experimental setup—as a consequence, it was not important to have the tasks performed identically in each condition. Instead, we were interested in how users behaved (as a means to enhance the functionality of the paper prototype), and in the choices that participants made with the paper prototype.

Observations

Participants generally completed the tasks without considerable difficulty. Using the whiteboard+calendar, participant behaviour in the planning task could be classified into two groups (unrelated to presentation order). The first group (n=5) would act on the planning task relatively haphazardly, placing items into the calendar until a conflict was encountered and then fixing it when these conflicts arose. A second group (n=4) would construct a timeline-like view on the whiteboard first, which would allow them to identify resource conflicts and "scarce resource" situations. It turned out that members of this second group largely had project planning management training, and while the details of timelines were slightly different from participant to participant, the intention and general outline of the construction was consistent.

With the prototype, users had little to no difficulty navigating between views, and made use of the features to help complete the tasks. Users reported no difficulty in learning or understanding the mechanisms at work, and reported that they appreciated each view for different tasks. In the following subsections, we report on the most salient findings of the study involving the prototype.

Finding 1: Views were desirable & their use was dependent on the task at hand

When using the prototype, all participants made use of the view switching between tasks (9/9). Generally, users used a single view when performing the scheduling task. When the timeline view was not available, all participants used the calendar view to do their planning (6/6). When the timeline view became a part of our prototype, two of three participants made use of it to complete the planning task. When using the calendar view, many would activate the filter function to help identify "free times" for resources. When the timeline view was available, this information was more readily available.

During the recall phase of the task, most users switched views depending on the question. For recall problems that focused on an individual, most users (6/9) toggled to the agenda view, citing the ease in simply "reading off" the information rather than needing to interpret the calendar. The following are some participant comments on the utility of the agenda view:

> *"The great thing [about the agenda view] is that you don't have to think: you just regurgitate what's there. With the calendar, you need to process. I don't like thinking."*
>
> *"I like filter and agenda – it just lets you read it. It gets pretty confusing really fast on a calendar with lot of stuff."*
>
> *"It's good for information retrieval; when information is finite and static – it's just easier to read stuff off."*
>
> *"It's better with more information—easier to go through step by step [compared to the calendar]."*

The calendar view was still useful in this phase for questions that asked about higher-level information, such as for the duration of a renovation project. 8/9 participants toggled to a calendar view when responding to such questions.

What should be clear is that the toggling of views was primarily done when transitioning between different types of tasks. For the planning task, participants preferred (alternately) the calendar and timeline views—such views give a good overview of the state of the problem, and facilitate the planning process; for the recall tasks, participants preferred the agenda view. Because the underlying data for both views was the same, the primary value in the views was in terms of easing visual access: agenda view requires little to no interpretation, while calendar view still requires some interpretation.

Finding 2: Lower level manipulations were infrequently used

In keeping with a desktop-like design philosophy, we also made sure to include the ability to perform drag-and-drop operations on the ink. Yet here, we generally did not see features such as drag-and-drop being used. Only three of nine (3/9) used the drag-and-drop facility. When probed, the remaining participants suggested that it was much clearer (to them) to simply write on the calendar itself.

> *"It's just habit. I'm not used to writing something on the side, and then dragging it onto my calendar."*

We see then that again, work practices with traditional technologies (such as whiteboards) affect how users behave with new technology. That said, there was still utility in being able to drag and drop. Most users (8/9) first wrote the set of tasks (along with the duration) that would need to be completed for each renovation job on the whiteboard. This would act as an on-boarding mechanism for the planning process. While most participants simply used this task list as a reference list, the remaining three used a drag-and-drop process so that they could identify when a task had already been scheduled.

> *"Once I do the drag and drop, I can see that it's done, and I don't have to worry about that [task] any more."*

Finding 3: In views where "work" was being done, users asked for error-notification

An oft-requested feature was the ability for the "work" views (i.e. the calendar view or the timeline view when it was available) to provide simple error checking. Because it was easy to create scheduling problems by simply working in ink, users thought that a computer-based system would be able to detect scheduling conflicts, and notify (likely by highlighting) when a conflict was present.

At a high level, this is interesting, as it suggests that beyond simply *presenting* information and providing a means to interact with the data, the views should also provide meaningful functionality—something that is impossible with traditional technology. Effectively, this is a context where the digital application can provide additional functionality without taking away from basic traditional functionality.

Finding 4: Filtering was desirable

Scheduling data quickly got overwhelming toward the latter stages of the planning process. It became difficult to read the information, and the sheer quantity meant that visually searching the information was difficult. Users quickly took to the use of the filtering tool, as it meant that the system could provide tailored views to answer questions that users might have. For instance, users frequently filtered on individual contractors, or renovation jobs during the recall task. This removed unnecessary clutter, and made the information easier to read. Similarly, during the scheduling task, many users employed the filter to find "next available date" for particular contractors.

Finding 5: Fast accessibility of alternate views is desirable

In addition to menu-based access to the alternate views, users suggested a number of additional ways that these alternate views could be accessed. In general, these suggestions reflected a general desire to be able to toggle or view data in the alternate views more quickly. For instance, many suggested being able to use a radio button at the top of the widget to toggle between views. About half suggested the ability to click on ink items in the calendar and agenda, and to be able to access an associated filtered view, or to bring up an alternate view. Two participants suggested the ability to have both calendar and agenda view visible at all times.

Justification for Paper Prototype

We employed a paper prototype in this study rather than using a functional computer-based system for several reasons:

1. Any system prototype would be comparatively brittle. As with most system prototypes, they are brittle—a consequence of limited development resources and only minimal testing compared to commercial efforts.
2. The purpose of this study was to explore the *concept* of view changing, *not* an actual system. A study involving a system prototype would have produced erroneous results: given poor handwriting recognition, for instance, the system would have completely failed. In such an instance, we would have been unable to evaluate whether the view switching had utility—which was the main point of the study to begin with. Our intention was not to study an *implementation*, but the *concepts* underlying such a system.
3. Designing and exploring multiple views would be time-consuming to implement. Using the paper prototype, we would also be able to explore views that would be time-consuming to implement, though potentially more useful for users in completing the task.
4. Ability to rapidly iterate on interaction mechanisms. Because paper prototypes are easy to develop, we would also be able to rapidly iterate on our designs based on user feedback. This allowed us to test a variety of interaction techniques in a short period of time. Use of a system prototype would limit our ability to examine such a large space of possibilities within the equivalent span of time, because development of system prototypes takes longer.

The paper prototype afforded lessons that would have been missed had we used a system prototype. For example, participants generated a new view (timeline view) of scheduling data that we had not considered in our initial design. Since we were running a study involving a paper prototype, we quickly mocked up this view of the data (Figure 5.1c), and later participants took immediate advantage of the view (and heavily praising it). Had we been using a system prototype for the study instead, it would have been completely infeasible to implement *and* study whether participants would actually find such a view useful.

Phase 2: Implementation of a Scheduling Tool for a Large Display

Given the insights from phase 1, we built an interactive prototype that realized these ideas in the context of an electronic whiteboard application. The main goal of the work was to demonstrate that the development of such a tool for an electronic whiteboard (e.g. SMARTBoard) was feasible given today's ink recognition technologies, and to understand what challenges would arise if the ideas from the first phase were realised in an electronic whiteboard application. The tool currently allows users to create freeform electronic ink (with pen-based input), recognizes dated annotations, allows them to be dynamically grouped, and then presented in different views.

The main challenge we overcame in the implementation of this tool was the ability to support view changes on electronic ink, which is inherently unstructured. Most tools supporting view changing do so in well-defined contexts, or when the data is fully structured (e.g. Excel). However, to maintain the ink affordance, we needed to allow users to retain a level of flexibility. Instead of forcing users to conform to a certain structure (e.g. date first, then annotation afterwards, or by using a form fill-in mechanism), we addressed this challenge by interpreting this electronic ink, and looking for "date signatures" that could appear anywhere in text. This electronic ink would then be recognized as a unit, allowing for view changes, and other higher-level manipulations. As a consequence, we allow for freeform electronic interaction while supporting structured view changes.

The freeform electronic ink interaction was important to us—we wanted to retain the richness of the ink affordance, allowing users to still sketch and draw freeform without restricting them to a highly structured input device or entry method (c.f. Foster & Stefik, 1986). Further, we wanted to explicitly support the different types of interaction with this data, such as editing, sharing, communicating or incidental encounters with the data. Thus, our design requirements were as follows:

1. Support ink strokes as the main method of input (that is, not input from a keyboard or mouse; focusing on interpreting strokes of ink as they were created by users).
2. Enable view changes for this ink input to support different activities or actions.

Our initial design experiments gave rise to additional requirements:

3. Provide a simple mechanism to notify users when ink strokes have been recognized.
4. Provide a simple mechanism for users to correct the system's interpretation of ink strokes.

As we saw in phase 1, it was also important to support view changes on this data. Our thinking was that users would have additional flexibility if we provided this functionality to each individual dated item of text: users would then be able to leave items in an "ink" view when they chose, or to add a number of these items to a consolidated calendar view, or even to create multiple calendars populated with different types of items (e.g. a "work" calendar, a "personal" calendar, and so forth). Because these items would be generated by interpreting native ink strokes, the whiteboard application could still function as normal whiteboards would—we would simply be augmenting the whiteboard with calendar and scheduling functionality.

Figure 5.3 An illustration of how the prototype calendar widget might appear amidst typical whiteboard content in a futuristic large display. Our current prototype (which can be seen inset) is a Windows application, so functions on TabletPCs, and SMARTBoards.

Interacting with the Calendar Widget

Figure 5.3 shows the calendar widget in its current instantiation in an imagined futuristic whiteboard. Here, the calendar widget is already populated with several items. This calendar is a mobile widget that sits atop the rest of the whiteboard application. As expected, it can be moved, and resized as the user sees fit. For instance, we observed that while many users placed reminders along the edges of their whiteboards, some particularly important reminders might be allowed to remain in different regions, or even enlarged to ensure that they would be seen at a later time.

The calendar widget provides several different views of this information, as illustrated by Figure 5.4. There are currently three views, accessed via a menu that is displayed when the user taps on the small green icon on the top right corner of the widget. These views include a standard calendar month view, an ink view (which allows the user to view and potentially edit the underlying ink data), and a combined calendar and ink list view, which provides a context+detail view (clicking on an item, either in the calendar or ink list scrolls the other view into place). We implemented these three views to illustrate the principle of supporting multiple views; however, variations on the views that we constructed for phase 1 (Figure 5.1) would be more appropriate if we had been designing a prototype for deployment.

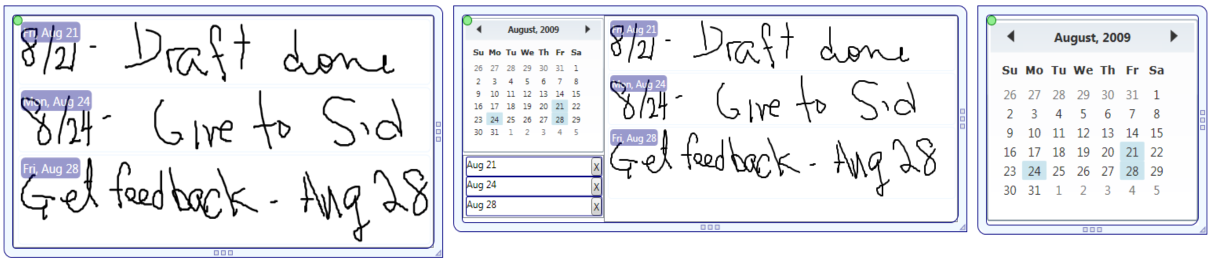

Figure 5.4 The prototype interprets the ink strokes, and allows users to visualize the data from the calendar widget with under three different views: the ink view (left), the calendar+ink view (middle), and the calendar view (right).

What is important here is that the data contained with the widget remains consistent: all that has changed is the view of the information that is contained within. The ink view may be desirable, for instance, if the data contains drawings, or complex amounts of information: as additional dated ink items are added to the widget, it is sorted in dated order, but remains visible. A calendar view may be sufficient if one wishes to only maintain an awareness of the important dates, but requires the use of the whiteboard space for other tasks. Finally, the combined view may be useful for searching and retrieving information that is contained in one of the widgets.

Creating the widget is done via a simple mixed-initiative mechanism illustrated in Figure 5.5. As ink strokes are added to the underlying ink canvas, the system automatically processes the ink, attempting to interpret the ink strokes as strings of text. Where strings of text have been recognized, the system processes the text to determine whether it contains date information. When a string of ink is recognized as a date, the system immediately underlines the recognized date. The underlined string of ink then becomes a hot-clickable region. Tapping on the region brings up a menu that simultaneously shows the recognized date while allowing users to correct the date through a simple date-picker dialogue. From here, the user can choose to create a calendar widget set to any of the default views. By default, ink strokes that were recognized as being on the same "line of text" are considered as part of the dated item (consequently, the date may appear anywhere

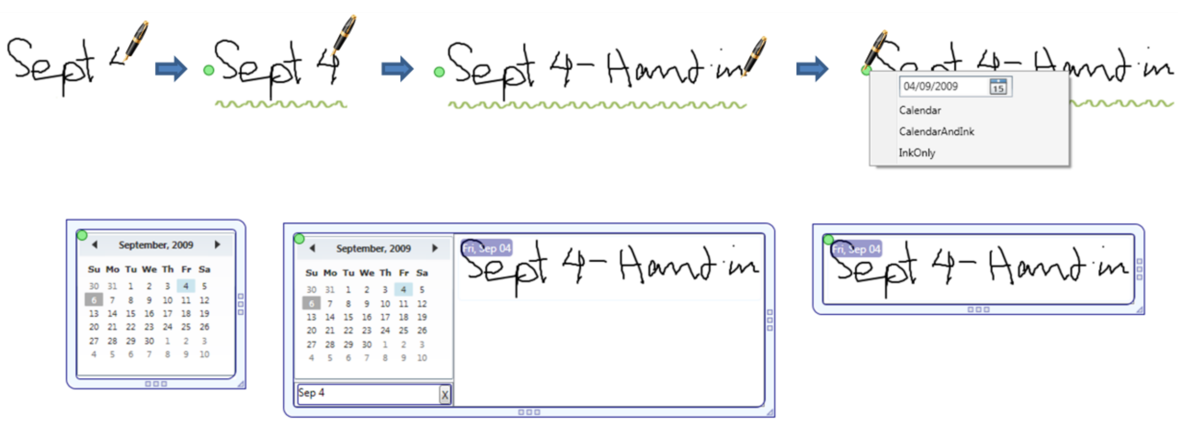

Figure 5.5 The recognition works in the background—once it recognizes a date string, it highlights the line of text containing the date. The user can then create a calendar widget by tapping on the green button. All three representations are immediately available: calendar, calendar+ink, and ink.

along the line of text).

Users may have more than one calendar widget at a time. For example, one may be used for personal events, while another might be for work items. As suggested earlier, it is also possible to add items to existing calendars widgets, and potentially merge widgets. These are effected through a straightforward drag-drop mechanism: dragging one calendar widget atop another combines them. Dragging items out of the calendar widget removes it from the calendar widget, and creates a second calendar widget. Finally, dragging a recognized group of ink onto a calendar adds the item to the calendar widget.

Implementation Details

The prototype was built atop Microsoft's .NET Framework, employing the Windows Presentation Foundation, the Microsoft Tablet PC, and Microsoft Tablet PC Ink Analysis WinFX Libraries. The main application was a simple window containing a number of overlaid transparent canvases: an ink canvas, a widget canvas, and a UI canvas. The ink canvas retains the original ink strokes, allowing the use of the application as a standard whiteboard. We also provided a floating toolbar that allowed users to change the size of the pen, the pen colour, and two methods of erasing ink strokes (by point or by stroke).

The calendar widget itself was a custom-built user control. This control contains a number of control and visual templates that are toggled between when the user changes the view. These calendar widgets are created dynamically through the interaction mechanisms described earlier, and then placed on the widget canvas.

Of particular interest is the recognition mechanism that understands date formats and underlies the entire system. Each ink stroke, as it was drawn, was added to the Ink Analyzer component provided by the Tablet PC Ink Analysis WinFX Libraries. This component provides a recognition tree of the ink on the canvas, along with confidence for recognized strings, as well as basic shapes (e.g. rectangle, circle, etc.).

When a new string is recognized, or an existing node's string is updated, we first check whether a calendar widget has already been created for the node. If not, then our date recognition engine is run on the recognized string. The date recognition engine was derived from a set of open source routines originally written by Sergey Stoyan of CliverSoft[3] for recognizing date and time strings during the parsing of different log file types. We modified this engine so as to relax input requirements, allowing the strings to be poorly formatted, or to be only partial dates, and so forth. This was done based on several pictures of whiteboards from our whiteboard study (Chapter 3), where we found that users rarely used completely well-formed dates, and instead relied on various methods of shorthand (e.g. Aug 3, 8/3, Mon, M, Monday, etc.). Furthermore, these dates might appear anywhere within a string of text. Because our focus was not on the recognition part of this problem, our system is not completely robust to all forms of date entry; however, it does well to interpret most typical forms of lines of text that contain a date. So long as the TabletPC Ink Analyzer can recognize one's writing, the date recognition engine functions reasonably well.

[3] http://cliversoft.com/

Given that a set of ink strokes is recognized as a date a wiggly-widget is created. This wiggly-widget merely places a squiggly line under the strokes belonging to the recognized dates. Tapping on it brings up the dialogue that allows users to instantiate the calendar widget. The calendar widget is populated with a normalized version of the associated ink strokes (translated based on the origin of the ink strokes) and the recognized date. The widget can then be toggled between the views. Information is scaled using a Viewbox in each of these views, so as the calendar frame is scaled, the data is scaled appropriately. When multiple items are contained within a calendar widget, they are displayed in date-ordered format.

Finally, we use the standard Windows drag and drop mechanism to allow users to drag and drop items across and between calendar widgets.

Implementation Challenges

Based on the implementation of this system prototype, we encountered several challenges. Future designers are likely to encounter similar challenges when building applications that enable transitions on whiteboard-like large display applications using view transformations. We describe these issues here, and describe how we addressed them. Some of these issues have been described before (Hong & Landay, 2000; Mynatt et al., 1999): we reiterate these, and add several to account for our interest in supporting transitions with view changes.

Grouping of Ink Strokes (i.e. what should be considered a unit?). The basic input element is an ink stroke. To facilitate interpretation and develop meaning, we group the strokes based on spatial proximity. In principle, temporal proximity should also be accounted for, though it is unclear which type of proximity should take precedence: ink strokes are unlikely to be laid atop one another unless they are related; however a temporally proximate relationship between ink strokes may be superfluous: in the case of Jill (Chapter 3), it may merely indicate a segment of time where the entire whiteboard was being updated.

Interpretation of Ink Strokes. Groups of ink (alternately called *segments* by Mynatt et al. (1999) and *strokes* by Hong & Landay (2000)) can be interpreted by various pluggable interpreters. These interpreters can range in functionality from handwriting recognition to sketch-based recognition (e.g. Igarashi et al., 1997), to interpreters that perform computation on the ink (e.g. arithmetic subsystem from Mynatt et al., 1999). Most importantly, this subsystem should be built as a pluggable architecture (as SATIN) so that interpreters can be added as needed or as new functionality is desired. In the prototype we develop, our main goal was to interpret the ink strokes as a means of generating some structured understanding of the underlying ink strokes.

Correction of Interpretation. Given today's technology, it is still reasonable to expect that the system's interpretation of the ink strokes is faulty: thus, simple mechanisms need to be provided that facilitate correction of the interpretation. For instance, a the grouping algorithm may fail to recognize that a new ink stroke is related to a set, or the interpretation engine may fail to recognize the intended text. In our prototype, we realized this correction mechanism via a date-picker control, allowing users to both see what the system has interpreted, and providing an opportunity to correct it.

Presentation of Alternatives. Interpretation engines should provide a level of certainty in their respective interpretations, and the alternative views that are available. As we articulate in Chapter 6, there are several different general methods of dealing with these alternatives (proactive, reactive, mixed-initiative). For proactive approaches, the most certain interpretation may be simply employed, whereas in reactive ones, no interpretation is enacted until the user decides one should be taken. In the system prototype we develop earlier, we use a mixed-initiative approach. Here, once the system's interpretation reaches a certainty threshold, the system actually shows the user that the text has been recognized and that alternative views are available (Figure 5.5). These alternate views may be presented in the form of call-outs, menu items, and such.

View Transformation. Here, the ink group is transformed into its new view. For instance, a set of words that have been written in ink may be simply transformed into a corresponding view of text. If the text incorporated date/time information, one might be able to drag this into a calendar, or transform the text itself into a calendar format. One might also manipulate it into a timeline format depending on which view was most appropriate. Even after a transformation is made, it is important to maintain the underlying ink strokes, as they may be useful for recovery, and are still a valid view of the information.

Ongoing Ink Interpretation. Further ink strokes on the newly transformed view also needs to be interpreted (within the context of the view). Many of these ink strokes may merely be annotations on the new view, and not require transformation. On the other hand, further ink strokes may be an attempt to augment or change the underlying data. For instance, suppose the underlying data was of a dated task item that had been transformed into a calendar view. If a user then placed additional information into another cell on the calendar, then this information would need to be present in the underlying data structure so that it could be exposed if the view was transformed again (e.g. into a task list view).

General Discussion

The research process that we employed in this chapter provided several lessons about designing tools and applications to support transitions. The paper prototype study revealed the importance of understanding the needs of each set of tasks, and developing useful views for those tasks. Here, there is the question also of granularity—how specialized should each view be, and how much data should be interpreted within each of these views? Our system prototype showed that one of the core challenges with supporting view changes is that the data needs to also have some underlying structure. In the context of our whiteboard application prototype, we *inferred* this structure from the underlying electronic ink based on some loose rules for dated items. Here, we discuss a the possibility of using templates as a more generalizable mechanism to infer structure from unstructured electronic ink.

Understanding the needs of each set of tasks, and what view is most appropriate. When using different views to support sets of tasks, it is important to understand what the needs are, and how the view can address those needs. Our original paper prototype provided calendar and agenda views, yet participants with project management training backgrounds ended up designing their own timeline views in the whiteboard+calendar condition. The timeline view better suits the scheduling process than either of the other two views, because its spatial layout automatically imposes constraints on the planning process (e.g. visually, it was immediately apparent that a resource was scheduled, and it effectively did not allow a resource to be doubly booked). While the other two views technically also provided this information, the timeline view provided this information rather implicitly. We therefore "implemented" this timeline view for later participants in the paper prototype study.

We also saw that when performing the recall task, users preferred setting up and using views where they would not be required to "think" a great deal. Comparing between the calendar and agenda view, the agenda view provided a much easier means to simply "read" information compared to the calendar view, where multiple elements would need to be consulted and cross-referenced before reporting the information.

Granularity of views, and interpreted data. One question that has been unanswered in our work is the granularity of support for each view: to what extent should they be specialized for each individual task that users might accomplish? The two possible extremes are to (1) design lots of views—one optimized for each task that a user may perform, (2) design only one view that is suboptimal for nearly all the tasks, but sufficient. In the paper prototype, we originally designed two views, and later found that the addition of a third was optimal, but this raised several questions: how many views are mentally tractable for users (i.e. how many can they remember to use)? How can tasks be meaningfully grouped together to provide views that are useful for each group?

These questions suggest that there is an additional element that designers need to consider when designing for transitions: granularity of action, and activity support. It also illustrates the importance of finding the right tasks and activities to support.

This lesson also underscores our rationale in using the paper prototyping technique in the study. If we had focused on the use of the system prototype, it is unlikely that this idea would have emerged as an important lesson. Our interest in that case would have focused likely on providing a suitable recognition engine, bugs in our implementation, and how would could iterate on that design. Being able to quickly iterate on the functionality of the system (as a paper prototype) allowed us to explore this idea fully.

Allowing easy toggling between views. As we saw earlier, it is important to provide an easy means to toggle between views. Users saw this view toggling as being useful enough to suggest faster ways to access filtered or highlighted views. Making the toggling between views difficult effectively does away with their utility. Even a feature as useful as filtering is what one might consider cosmetic (given that the information before filtering is still present); however, obscuring access to such a feature effectively makes a user's task more challenging.

Challenges of ink recognition. One of the challenges that we faced in designing and implementing our system prototype was the difficulty of recognizing ink strokes, and being able to interpret them such that alternate views could be generated. This problem is challenging due to the idiosyncratic nature of users' behaviour around whiteboards, and individual writing styles. Even recognizing which strokes belong together is challenging. A common approach is to use spatial proximity (e.g. Ju et al., 2008), but this solution is still subject to errors as the thresholds are again specific to each individual. One possibility is constraining the problem in attempting to find a balance between a form-filling application with the free-form general utility of a whiteboard. Yet the problem with a form-filling application is that the problem space such an application can address is immediately limited (and prescribed).

Templates for semi-structured input. Given what we have seen, a potential solution is the introduction of some generalized "templates" that could be used as a means to support interpretation. These templates could appear as faint outlines on the whiteboard space, helping to guide users into adding information in such a way that facilitates machine interpretation. Many simple templates include: a calendar-like template (as seen in our paper prototype), and perhaps more generally, a grid-like template. The grid template is analogous to spreadsheet applications, which simply provide a grid—the meaning of the grid, how it is interpreted and then used later depends primarily on the user. In such an application, information could remain on the grid, or the user could apply different view transitions on the information contained in the grid.

Even in the context of using templates, it might be possible to introduce templates that facilitate specialized tasks (to reduce the declarative manipulations required later). For instance, one might imagine a specialized "scheduling template" that helps to address exactly the problem that we saw in our study. Such a template might already have functionality built in to support constraint management, but more importantly, alternate views of the data would already be defined.

Summary

In this chapter, we have designed and implemented a tool that supports transitions between calendaring tasks using view changes as a proof-of-concept. We have documented the design process of this tool, which ensured that the functionality could and would be used by users. We then showed that it was possible to implement this tool, addressing in particular the problem with generating views on unstructured ink data. Reflecting on the entire process, we then describe a set of lessons to be learned for designing to support transitions in general. These lessons include: understanding the granularity of the activities that need support; deeply understanding the needs of each activity; supporting easy transition between these activities, and finally understanding the problems of recognition.

In the next chapter, we consider the space of designing to support transitions as a whole. Based on the findings from this study and the previous two chapters, we articulate a design space for supporting these transitions, and then discuss how prior solutions fit into the design space.

Chapter 6

A Design Space for Supporting Transitions in Large Display Applications

In Chapters 3 and 4, we showed that large surfaces should support multiple functional roles while addressing users' changing needs. Chapter 5 showed how view changing could support the transitions between different tasks. These findings provide us with models of different roles that need to be addressed in large display applications, and emphasize the role of transitions in large display application design.

As suggested in Chapter 2, some prior systems have built support for transitions, albeit not explicitly using the term "transition" as we have. Although in Chapter 1 we developed a formal model for transitions, in this chapter, we articulate a focused design space more pertinent to the immediate design of transitions in large display applications. Of particular relevance to this chapter, the operationalization of a user model of transitions can be used directly in the development of proactive approaches that we discuss. To demonstrate the utility of this design space, we show how it can be used to classify existing systems (including the one designed in Chapter 5), as well to discuss design ideas as they relate to transitions. Finally, we articulate challenges in each segment of the design space.

Design Space for Transitions in Large Display Applications

Our design space has two primary axes: the first axis captures the type of transition the system is to support, the second axis captures the form of agency—or the way in which the system supports that transition. We describe each point along each axis in turn.

Types of Transition Support

We classify primarily three types of transitions as illustrated in Table 6.1.

Functional: These transitions are those where the user needs to shift his/her attention to a different, related task during an activity. Supporting a functional transition means enabling users to shift between these tasks in an easy, meaningful fashion while maintaining context. In Chapter 3, for instance, we observed users who used the whiteboard to brainstorm a design, and then used the whiteboard to store the design for later reference. We see here that although the user's needs and activities have changed, they still employ the same underlying data. The scheduling system we designed and described in Chapter 5 supports this type of transition by allowing users to perform both scheduling tasks and recall tasks on the same data.

Collaborative: These transitions are those that allow users to move between independent and collaborative work or vice-versa—the types of transitions that we described in Chapter 4. Such a transition would allow users to easily partition and merge work or to help interpret each others' work. It should allow users to create and interact with artefacts that allow both meaningful independent activity as well as collaborative activity. One difficulty that has been observed before is that designs that support the activity of a group may not do so well in supporting the activity of individuals, and vice versa (Gutwin & Greenberg, 1998). Of note is that the desire to support this transition has been more enthusiastically explored with tabletop displays compared to upright displays: presumably, this is due in part to how tabletops inherently facilitate independent activity due to territoriality (Rogers & Lindley, 2004; Scott et al., 2004). The use of shared upright displays for simultaneous independent activity, where all activity occurs on the same upright display is considerably rarer (e.g. Brignull et al., 2003).

Temporal: These transitions allow users to shift between synchronous and asynchronous work. For instance, by allowing users to suspend activity, or to create artefacts whose intended use is for a later time. Chapter 3 revealed two examples of the whiteboard being used for such transitions: in John's case, the whiteboard functioned both as a space for brainstorming and design as well as an ambient display for his work that he could use to revisit and revise ideas; similarly, Jill's whiteboard facilitated both asynchronous (periodic review) and synchronous (timetable juggling) work. Similar behaviours have been observed with persistent coordination whiteboards in hospital spaces (Xiao et al., 2001): users engage in both synchronous and asynchronous collaboration. Some

		Transition Type		
		Functional	**Collaborative**	**Temporal**
System Agency	**Proactive**	Vogel & Balakrishnan, 2004 Igarashi et al., 1997	McCarthy et al., 2008	
	Mixed Initiative		Morris et al., in press	Ju et al., 2008
	Reactive	Newman et al., 2003 Mynatt et al., 1999 Tatar et al., 1991	Tobiasz et al., 2009	

Table 6.1 A design space for transitions in large display applications, populated with representative examples from the literature. We have highlighted regions of the design space that have been unexplored.

large display applications designed for these spaces support such transitions via persistent display of information (e.g. Bardram et al., 2006).

System Agency

With this design space, we are also interested in how the support is initiated: is it user initiated, initiated by the computer, or instead from a dialogue between the user and system? This classification scheme is derived from Morris et al. (2004), and builds on the formulation by (Ju et al., 2008). While Morris et al. (2004) used this scheme to describe the design space for how systems resolve interference in tabletop applications, we have found that the dimension works well as a means to classify the kinds of transition support that we have seen in the literature.

Reactive: A system is described as being reactive if it initiates no action on its own, responding (i.e. reacting) only to user input. In the context of our framework, we refer here to explicitly initiated changes on the large display—i.e. the user has taken explicit action to change something about how the system interprets input. An analogous word processor action would be to press the "Bold" button so that further keystrokes are rendered as bold. The paper prototype in Chapter 5 functioned in this way: transitions were effected entirely by actions taken by the user (to toggle views).

Proactive: A system is described as being proactive if actions are initiated on behalf of the user. Typically, these actions are based on some form of implicit input by the user, an inferred model of the user (and the user's goals), or both. For instance, motion detectors in meeting rooms are often used to activate room lights automatically when motion is detected: the premise being that if there is motion (implicit input), it is being caused by users who would prefer the lights to be on (simple model of user's goals).

Mixed-Initiative: A compromise approach articulated by Horvitz (2003) is to employ dialogues with users to enact adaptation. Here, the system's role is to proactively communicate alternatives and possibilities to the user based on inferred intentions based on available input. The user can then choose whether to act on these suggestions. Although in its original instantiation, Horvitz's mechanism to provide dialogue was through the use of a social agent, the contemporary interpretation of "dialogue" in this context is considerably looser. Instead, it simply refers to a reciprocal/iterative interaction between the user and system, wherein the user and system both engage in clarifying acts to iteratively determine some action that is to be taken by the system. The system prototype in Chapter 5 was designed as a mixed-initiative mechanism: the system indicates when a line of text is recognized as being potentially useful for scheduling via a squiggly line, but the user completes the action by clicking on the indicated line of text.

Exploring the Design Space

To demonstrate the utility of this design space as an organizing framework, we classify several existing large display applications. This brief review will show that within this space (Table 6.1), the vast majority of existing large display applications fit under the *reactive* space—responding primarily on user input rather than relying on a model of the user. The few systems that are proactive or rely on a mixed-initiative model are considerably outnumbered. This reflects perhaps the difficulty in articulating a user model that adequately describes the breadth of use for a large display.

Functional/Reactive: DENIM provides the canonical example in this space (Newman et al., 2003). DENIM is a large display application designed to support the rapid iterative design of websites, facilitating both ad hoc and planned activity. The system allows designers to sketch out websites at an architectural level, then transition to page or element level, bringing together each level in a sketch-inspired, zoomable canvas interface. Users can zoom back and forth through different levels, as well as traverse the links. Viewing the website at each level is a different, functional activity, but note that the system is completely reactive: responding to the user-issued commands.

FlatLand allows users to apply an interpretation/view filter on a set of ink strokes in an electronic whiteboard-like application (Mynatt et al., 1999). These view filters are akin to "mini-applications", allowing users to selectively apply computation to specific regions (or segments) of ink. We see then the ability to apply this computation is functionality that is applied on-demand.

A final example in this space is the Cognoter system, which facilitated mediated meeting activity comprised of three tasks: brainstorming, ordering and evaluation (Tatar et al., 1991). Users interacted with the large display using independent computer terminals using a form of input redirection. What is instructive here is that the system facilitated moving from one task to another using the same underlying data, but presented in a different way (thereby facilitating different types of interactions). In the brainstorming phase, participants could post text items to the large display from their individual terminals. In the ordering phase, the data remained the same, but participants could only move, reorder and draw pointers between items. As with the previous two examples, transitions into latter phases of activity with Cognoter was driven by the meeting leader.

Functional/Proactive: Vogel & Balakrishnan (2004) track a user's proximity to the large display, and reacts by smoothly transitioning the displayed content based on proximity. In the specific application explored by Vogel & Balakrishnan (2004), a user's calendar is displayed at varying levels of fidelity: from far away, it is presented in an abstract form, and as the user approaches the display, finer grained details are presented. When the user is immediately next to the display, even information marked as "private" is displayed. Note here that the system is completely proactive: the display transitions between views based exclusively on the user's proximity to the display.

Functional/Mixed-Initiative: At this point, we do not know of a large display application that supports functional transitions using a mixed-initiative approach. The Range system (Ju et al., 2008) uses a mixed-initiative approach, but its features are more correctly classified as supporting temporal transitions. A closer set of examples are the interactive beautification techniques for free-hand sketching (e.g. Igarashi et al., 1997). Here, users' ink strokes are beautified using geometric constraints, and to prevent recognition errors, systems generate multiple candidates from which users can select. This mixed-initiative approach is an "in the small" version of what we might envision for a richer large display application. In such a system, the application infers, based on ink strokes, the intended and underlying structure (i.e. data) that is being constructed. It would then infer users' intended actions (i.e. the role the display might play) and offer these possibilities as views that the user could apply on the underlying structure. Ink (i.e. changes) applied to these views would then need to be re-interpreted back into the underlying structure. While the technical details of such an approach are challenging, the main difficulty remains human-centric: first, what transition states should the system support; second, *how* should the system support these states (i.e. what would actually be useful?). We investigate these challenges further in the next section.

Collaborative/Reactive: Systems in this space need to facilitate the transition between independent and collaborative tasks. Lark is a collaborative tabletop information visualization system designed to support both independent exploration of data sets, as well as collaborative sensemaking and analysis (Tobiasz et al., 2009). The system integrates a visualization of the infovis pipeline from which users can create branches for independent study (i.e. to investigate their independent hypotheses), but because the meta-visualization of the pipeline is present, users can easy re-integrate their independent findings during collaboration phases of activity. For instance, they can determine when their investigations diverged, and where the common points of investigation were.

Collaborative/Proactive: Joe McCarthy and his colleagues have developed an extensive programme of research collectively called the Proactive Display project (McCarthy et al., 2008; McCarthy et al., 2005; McCarthy et al., 2004; McCarthy et al., 2001). Broadly, these displays sense and respond to people or activities taking place nearby. These displays thus take action on behalf of users in attempt to aid their current or expected activities. The most recent generation of these displays provides users with a constantly updating collage of information about nearby users. The information is pulled off associated social networking sites (e.g. Flickr), while proximity is detected using Bluetooth signatures (from the users' mobile phones). Users can interact with the display itself to construct collages of media or to more closely inspect others' media. Thus, the display functions as an ambient shared display (i.e. as a conversational resource) that users can employ for more independent activity (e.g. collage construction) if they so choose.

Collaborative/Mixed-Initiative: The WeSearch system is a collaborative tabletop search and sensemaking system (Morris et al., in press). It provides users the ability to perform independent, parallel work in the form of searching and navigating the web (i.e. loosely coupled work). Beyond this, it allows web-snippets to be created and shared among members of the collaborating group, which can be used for collaborative sensemaking. Further, search terms and titles of webpages that have been clicked on are immediately shared with other collaborators for use in their own queries. The sharing model is presented as a marquee: each term can be touched, which in turn drops the term into one's existing search field. The user model here is fairly straightforward: since users are likely to search for similar topics as their collaborators, making these search terms available for re-use can be of utility. The system thus provides users with opportunities for both independent and collaborative work, as well as opportunities (and suggestions) on how to engage with others' work more closely.

Temporal/Mixed-Initiative: Range is the only large display system that we could find that explicitly supports temporal transitions (Ju et al., 2008). Range does so with two modes of activity: whiteboard mode, and what is effectively a "screen-saver" mode. In whiteboard mode, the display acts as a whiteboard, time-stamping each stroke, and grouping strokes together based on a closeness metric. This mode is intended for real-time, synchronous activity. In screen-saver mode, previous versions of the whiteboard float about the screen in as marquee slideshow. This mode is intended as an asynchronous reminder of prior activity. The transition between these states is based on a mixed-initiative protocol. In the case of the transition between whiteboard→screensaver, the system operates on a combined timeout and proximity mechanism: if people move away from the display for a period of time, the screensaver activates. This is a mixed-initiative protocol because the user can easily step back in and override the transition manually to keep the system in "whiteboard mode" (the transition is a visible one that takes a few seconds to complete). A similar transition occurs from screensaver to whiteboard mode: as the user moves close to the display, images of the prior whiteboard states move to the edges of the whiteboard before moving completely out of view. During this time, a user may grab any of these prior states to use as the basis for continued sketching.

Discussion

The brief review of the design space illustrates several points about large display applications: first, there is an implicit recognition for the need to support a variety of tasks using a range of tools (i.e. functional transitions), a need to support collaborative transitions (i.e. allowing users to transition between collaborative and independent activities), and a need to support temporal transitions (i.e. between synchronous and asynchronous activity); second, the term "transition" as we have described it is not very succinctly defined—consequently, the range of functionality we bring under this umbrella is related mainly at a conceptual level. The common thread among these systems is that they allow users to easily use functionality as their needs change. In many cases, like our calendaring tool, they allow the user to operate on the *same underlying dataset* in these different modes of work. This is fundamentally different from the desktop model, where each application provides a subset of functionality required to complete all the tasks in a given activity, and to accomplish all tasks, one might need to change applications altogether. Here, the systems outlined above operate on the *same* data set across tasks, meaning that the context of the system (i.e. the data) is fixed, regardless of the task being performed.

As Table 6.1illustrates, there is still a dearth of systems that actually support collaborative and temporal transitions, as well as mixed-initiative approaches to supporting transitions. It is likely that the former is due to a poor understanding of what such transitions really mean. The models that we develop in Chapters 3 and 4 of large surface roles and user needs provide two perspectives on the issue of transitions, giving designers a starting point for designs supporting collaborative and temporal transitions.

This design space also bring to light three challenges facing designers of large display applications: first, what types of activities/tasks need to be supported; second, how should those activities/tasks be supported on the large display; finally, how is the support exposed to the user (reactive, proactive, mixed-initiative). The first two of these challenges were already raised in Chapter 5, but it serves to reiterate them here.

As we suggested first in Chapter 5, the first challenge, identifying activity states for transition support is at the heart of the problem for this design space. Without accurately pin-pointing the activity states, and understanding the information and interaction needs of users at each of these states, a design is likely to fail. In the best case, the use of designed mechanisms that support that state may simply go unused. In the case of the original paper prototype we designed in Chapter 5, participants simply made the best of the situation, and used the calendar view for planning. When we finally introduced a timeline view that better supported the scheduling task, users were more apt to use this view to complete their work.

These activity states may be rightfully considered as analogues of "applications" in the desktop-centric application model nomenclature. Providing sufficient functionality in each of these activity states seems like a daunting prospect—mainly because the number of activity states (roles or functionality that large displays should provide) seems intractable. This may, however, be a case where the *context* of a large display's deployment may usefully help limit the number of activity states that are necessary. For instance, the nursing station whiteboard in Figure 1.1 only supports a finite number of tasks, which are defined by its spatial and social context (as were the whiteboards in Chapter 3).

The second challenge, determining how activities and transitions between them should be best supported by the surface, is beyond the scope of this dissertation. Should each activity/task have its own interface, or should there be an independent, separate interface for each activity/task? In the former case, should each interface just be a separate application, as they typically are in most desktop scenarios (i.e. we have separate software for word processing, presentation, and spreadsheet work)? On the other hand, should these activities and tasks all be addressed by a single application with only a limited number of interfaces? If the latter design approach is taken, then the challenge of how to support these transitions becomes highly context-dependent, and the problem space becomes more tractable. In Chapter 5, we explored alternate views of the underlying data as a means to support transitions, and demonstrated it as a viable technique to support transitions.

The third challenge is deciding on how to expose the transition support to the user. Our design space revealed three different design approaches, though it seems the most appropriate approach is the mixed-initiative approach. We advocate this approach for three reasons: first, it explicitly exposes available functionality as an affordance to users who are unaccustomed or unfamiliar with the capabilities of the application (compared to a reactive approach); second, it exposes this functionality in a cautious manner that allows users to accept or decline the functionality (cf. proactive approach); finally, a mixed-initiative approach is considerably more forgiving that a proactive approach: in the event that the system has incorrectly inferred the user's state or desired state, the suggestions provided by the system can be more easily ignored (cf. proactive approach).

Summary

Some generalization of this design space might be possible so as to include single-user desktop systems. The traditional desktop metaphor to support transitions is effectively cut-and-paste—an explicit, reactive mechanism—to enable functional transitions as a user moves between different applications. As suggested in Chapter 2, there is an increasing recognition that application boundaries are far too arbitrary, and difficult to cross. The Microsoft Office suite of productivity tools (which includes a word processor, email client, presentation software, and spreadsheet application) was originally designed to be a set of standalone applications. Over the years, those designers have worked to integrate the applications to enable smooth workflows where users can move between the application seamlessly. The scenario described in the introduction (regarding the sending of a document to a colleague), for example, is now made far more smooth. From the word processing application, a user is now provided with functionality to prepare and send the document from the word processor itself (i.e. without needing to switch to the email client). Similarly, the email client actually makes use of the word processor's document editing component, meaning that from the email client, all of the rich formatting capabilities that one would expect in a word processor are exposed to the user. In this simple way, designers are recognizing the importance of transitions in everyday work, and integrating it also into single-user desktop applications. Our design space therefore also describes many of these scenarios.

In this chapter, we consolidated the ideas and findings from Chapters 3-5 into a coherent design space to support transitions on large display applications. In particular, we focused on the use of view transformations as a means to support these transitions, and introduced a framework to support these transforms. Further, we anticipated and articulated a set of challenges that face designers of large display applications attempting to support these transitions.

One of these challenges is to deeply understand the states or activities that the system will support transitioning between. As we stated above, a poor understanding of these activities understandably reduces the system's utility. We saw this in our study in Chapter 5, when users had to perform scheduling using a calendar, whereas the timeline view made this task much easier. Furthermore, as suggested in Chapter 3, particular surfaces are typically used for only a fixed set of tasks in a given context.

Chapter 7

Conclusions

This chapter summarizes the contributions and main ideas from this research. It discusses how practitioners should employ the findings in their own designs and applications, as well as outlining several avenues of continued research.

Thesis Contributions

We began this dissertation with a set of research goals. To achieve these goals, our research process produced a number of contributions that we review here.

1. **Development of a new framework for large display interaction focusing on transitions.** The findings of this work suggest that a new framework of application design is needed for large display applications. We inform this framework through several sub-contributions: (a) and (b) address the multiple roles that large surfaces play in supporting users' work (which necessarily implicates the user's model of how large display applications should support their work), while (c) and (d) address the changing information and interaction needs of users.
    a. **Classification scheme describing activity modes around traditional large surfaces.** Based on a study of whiteboards, we outline a classification scheme for the activities that take place around whiteboards. In particular, the work identifies the coordinating role whiteboards play in supporting transitions between synchronous and asynchronous work, and independent and collaborative work. These findings go beyond prior work, which typically identified different whiteboard tasks and activities, but failed to recognize that often, those tasks and activities are related, and that the underlying data to support the tasks is the same.
    b. **Descriptive classification for the role of surfaces in meeting room collaboration.** Our observations of meeting room collaboration inform a descriptive classification scheme for how traditional surfaces are used in collaborative activity. This framework lays the groundwork for the different functional roles large display applications should play in supporting collaboration.
    c. **Model of collaborative coupling around tabletop displays.** By studying collaboration around tabletop displays, we develop a model describing how users

engage and disengage with one another as they complete tasks using a large display application. This model of changing needs illustrates the importance of facilitating mechanisms that enable both independent and collaborative activity around shared displays.

   d. **Design principles to support bystanders around large public displays.** A field deployment of an interactive public display informs a model for bystanders and bystander needs around such displays. These findings illustrate the relevance of supporting bystanders transitions from casual bystander to contributor, as well as the importance of supporting multiple users with different information needs from the same display. These principles aid the design of interactive public displays to encourage public understanding and participation.

2. **An operationalization of transitions.** This conceptual framework gives researchers and practitioners a concrete reference and vocabulary to describe and think about transitions. The framework emphasizes the need to consider the greater context of activity, and we demonstrated through a number of studies how it is a suitable model to describe a wide variety of activities as they relate to large displays and large display use.
3. **A system supporting transitions on large interactive surfaces.** Through the design and implementation of a prototype calendaring tool, we show how transitions can be supported using viewing changes. This proof-of-concept serves as an illustration of the utility of supporting transitions.
4. **A design space for transitions in large display applications**. This design space is both useful for classification, and as a generative mechanism, bringing together existing efforts, as well as suggesting new alternatives for designs. In classifying these systems, and seeing them as a whole, we are able to uncover a number of inherent challenges common to each: first, the challenge of identifying meaningful states to transition between (and understanding the needs of users in these states); second, the challenge of accurately identifying users' desire for transitioning, and finally, the challenge of exposing these transitions to users.

Many of the contributions listed above have been previously published. We outline these publications here:

Conference Publications

Tang, A., Lanir, J., Greenberg, S., and Fels, S. 2009. Supporting transitions in work: informing large display application design by understanding whiteboard use. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work* (Sanibel Island, Florida, USA, May 10 - 13, 2009). GROUP '09. ACM, New York, NY, 149-158.

Tang, A., Finke, M., Blackstock, M., Leung, R., Deutscher, M., and Lea, R. 2008. Designing for bystanders: reflections on building a public digital forum. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 879-882.

Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. 2006. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 1181-1190.

Oral Presentations

Tang, A. 2006. Surface use in meeting room collaboration. In *Conference Companion of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY, 43-44.

Tang, A., Parker, J. K., Lanir, J., Booth, K. S. and Fels, S. 2006. Studying collaborative surface use to guide large display interaction design. In *Conference Companion of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY, 219-220.

Workshop Participation

Tang, A. and Fels, S. 2008. Four lessons from traditional MDEs. ACM CSCW 2008 Workshop on Beyond the Laboratory: Supporting Authentic Collaboration with Multiple Displays.

Tang, A., Finke, M., Blackstock, M., Leung, R., Deutscher, M., Tain, G., and Giesbrecht, C. 2008. Designing for bystanders: reflections on building a public digital forum. *ACM CHI 2008 Workshop on Designing and Evaluating Mobile Phone-based Interaction with Public Displays.*

Implications for the Design of Large Display Applications

The main focus of this work was to develop the foundations for a new framework for large display application design. We initially motivated this work by outlining how large display applications were likely to be used in fundamentally different ways than desktop applications. Our findings strengthen this position, as we articulate the multiple roles of surfaces in supporting the changing needs of users. We have argued that the ability for users to transition between tasks as their needs require should serve as the foundation of this new model for large display application design.

This is an important, and distinguishing factor of large displays compared to our current use of desktop computers: whereas the desktop computer facilitates "transitions" via different applications, large displays are inherently different because they are often used by multiple users. This dissertation provides a starting point for these designers, and we articulate four core lessons that they should take from this work:

1. ***Design not only for a single task, but enable multiple tasks and functions, providing mechanisms to transition between those modes of operation***. Our study of whiteboards (Chapter 3) brought particular attention to collaborative and temporal transitions. And,

while whiteboard users employ specific *representations* of information that can be used across modes of activity, interactive displays can do better by employing different *views* that would better facilitate each of these modes of operation. In Chapter 3, we illustrated the range of activities these surfaces are used for in meeting room contexts. In Chapter 4, we showed that users of large interactive surfaces transition between multiple modes of activity, even when the application does not actively support them. With the tabletop system, users transitioned between independent and group activity—something we called collaborative coupling. With MAGICBoard, we also saw users transition between casual bystanders to being transient users, illustrating that in some cases, the display needs to play several roles simultaneously for each user type. Chapter 5 illustrates the design of a large display application with transitions in mind, using view transformations as the primary mechanism to support these transitions.

2. ***Study and understand the context of use, particularly focusing on the intended role and function of the traditional surface***. In Chapter 3, we develop this understanding with respect to the whiteboard, illustrating that the function of a large surface is determined by a complex interplay between the users of the whiteboard, as well as its location. Later in Chapter 3, we build on this understanding with regard to engineering project-room styles of meetings. We derived a descriptive framework that describes generic classes of activities that occurred around traditional large surfaces. For the design of a large display application, it is important both to deeply understand the role that it is expected to play in people's work, where, and when.
3. ***To better support temporal transitions, consider an always-on display, and what function that display has for asynchronous activity***. As outlined in Chapter 3, we saw that there are many types of asynchronous activity: deferred, periodic, and intermittent. Instead of relying on the user to remember to re-engage in such activities, an always-on display helps remind users of such pending activities by its mere presence. How such information should be presented to the user depends on its role: in some cases, its purpose is awareness, in other cases is as a reminder of pending activity, and so forth. There may be many different types of views that are appropriate depending on the context and purpose.
4. ***To better support collaborative transitions, consider the nature of individual and collaborative work in this context, and how such activities can be best brought together or separated***. Such points of "bringing together" include trying to merge, resolve or explain independent strands of activity. We have seen that these can be visual transitions, or meta visualizations that allow users to see points of divergence.

Future Directions

This dissertation provides us with a bridge between the work practices of users employing large traditional surfaces, and the design of large display applications. The focus of this work was developing a model of users' behaviours and of how large surfaces are used. To do so, we developed a model of transitions (Chapter 1) and a design space (Chapter 6) around our observations of transitions, highlighting their importance to effective large display application design. Our work raises several avenues for continued research, which we discuss next.

Near Term Directions

Contextually Relevant Large Display Applications. Our work has demonstrated the utility of supporting transitions generally, and within the context of a scheduling task. A clear next step is to apply this idea to large display applications with an understanding of the role the display takes within a specific context. These physical contexts, as we outlined in Chapter 3, help define users' needs of the large display. Given that a core challenge in supporting transitions is understanding the needs of each context, it serves to build on knowledge that we have developed in this thesis.

We studied meeting room contexts in Chapter 3, and the results of that study can inform the design of a large display application for such a context. In particular, we identified the basic uses of traditional surfaces in meeting rooms (which we articulated as activities), as well as the roles these surfaces play in supporting such activities. These findings should provide fruitful initial directions for a meeting room large display application.

Similarly, our findings from Chapter 4 regarding interactive public displays provide insight into how to develop systems that engage casual bystanders. We have applied these ideas toward the design of a rich, interactive game for public large displays called Polar Defense, and used this study to further refine our understanding of bystanders. Our methods for supporting their transition from bystander to contributor are reported in Finke et al. (2008).

Designing Tabletop Groupware. With regard to interactive tabletops, our ideas regarding collaborative coupling (Chapter 4), and users' transitions between these states is particularly interesting, because territoriality facilitates independent activity (Scott et al., 2004). In these spaces, it is thus worth exploring how complex activities can be broken into pieces for parallel work before collaborators return to working together. With sufficient understanding of these processes, it would then be possible to design interfaces that support such transitions. We have begun this exploration in the context of visual information analysis. In an early study with co-authors Petra Isenberg and Sheelagh Carpendale, we explored how collaborative exploration of visual information occurs. The study, reported in Isenberg et al., (2008), gave us an in-depth understanding of how users smoothly transition between multiple types of information analysis tasks, some of which involve independent activity, while other involve collaborative activity. In this exploratory study, participant groups (singles, doubles and triples) were provided with a set of paper-based visualizations from which they were asked to solve a number of problems. We analyzed what users were doing with regard to the information items, as well as how they were working with respect to one another, developing a framework that captures the analysis activities of collocated teams and individuals. This type of work provided the basis upon which Petra Isenberg's further work on collaborative information visualization tools (Isenberg & Fisher, 2009). Many such tools aim to facilitate the transition between collaborative and independent activity.

Longer Term Directions

Design Space Directions. The design space shows that proactive and mixed-initiative solutions for supporting transitions are areas of the design space that have not been deeply explored. The earliest approaches have employed proximity as the main mechanism of input to the large display, but more generally, ubicomp has explored a range of sensors that may be of use. For instance, it may be useful to determine the number of other potential users in the room, or to detect their current activity patterns, or level of speech and ongoing dialogue, and use these factors as a means to suggest transitions on the large display. For instance, our prototype in Chapter 5 provides feedback to users when it thinks it has detected a potential dated reminder by underlining the ink. They system may confidently transition into a state useful for independent work if it is confident that no other users are in the room, or that they are not making use of the large display. When more users enter the room, or pay attention to the large display, then it may be useful to transition the system into a state more conducive to collaboration. These ideas could be explored with either proactive or mixed-initiative approaches, though simple proactive approaches are likely not to be as successful (Ju et al., 2008).

It may also be that points in the design space map well to particular contexts (i.e. are some points in the design space better suited for certain contexts than others). Does, for instance, the mixed-initiative approach suit interactive displays that are used by more than one user? What happens when different users have different preferences for how such information is displayed? Would it be appropriate to provide individualized views of displayed content, or is it more important to rely on a consistent representation?

Widgets vs. Applications. This thesis provides evidence that users' actual tasks transcend traditional desktop "application" boundaries, and that their needs and expectations of large displays may be fundamentally different from desktops. In Chapter 3, we suggested designing functional primitives rather than monolithic applications for large display applications—an idea that we realized in Chapter 5 within the context of a scheduling widget. We believe that such an approach, with appropriately designed widgets, can be adopted easily by users to support their activities. The flexibility afforded by this approach also facilitates the fluid and changing needs of users as they appropriate the large display for their idiosyncratic uses.

A Future of Displays. In the near-term future, we envision ubiquitous, cheap displays that are paper-thin, and physically flexible—perhaps not replacing, but certainly augmenting everyday paper. In such an environment, where the displays can be used, where information changes hands, and across tasks, transitions will become even more important. In such a context, what will the role of an "application" as we know in the desktop computing sense play? Information portability, and the way in which it traverses the different displays for different uses will play an increasingly important role. Further, supporting these traversals in the same way that we consider transitions, will help these activities.

Concluding Comments

New technology rarely arrives with a deep understanding of how applications can be designed to exploit it. In the case of large displays, for example, we have continued to employ the traditional desktop model of application design. Yet, our research illustrates that this model may be inadequate in addressing the unique design context of large displays. By studying the use of some rudimentary large display applications, and developing an understanding of how traditional large surfaces are used, we have provided the groundwork for the development of a new framework for large display application design. In particular, we have emphasized the role of transitions, and their importance for large display applications. It is in this understanding that we hope designers and researchers come to see that large display design needs new paradigmatic approaches from the ubiquitous desktop WIMP, and it is our hope that this thesis provides a starting point for exploring some of these deeper questions.

References

[1] Baker, K., Greenberg, S., and Gutwin, C. 2002. Empirical development of a heuristic evaluation methodology for shared workspace groupware. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work* (New Orleans, Louisiana, USA, November 16 - 20, 2002). CSCW '02. ACM, New York, NY, 96-105.

[2] Bardram, J. E. 1997. "I love the system—I just don't use it!". In *Proceedings of the international ACM SIGGROUP Conference on Supporting Group Work: the integration Challenge* (Phoenix, Arizona, United States, November 16 - 19, 1997). GROUP '97. ACM, New York, NY, 251-260.

[3] Bardram, J. E., Hansen, T. R., and Soegaard, M. 2006. AwareMedia: a shared interactive display supporting social, temporal, and spatial awareness in surgery. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY, 109-118.

[4] Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P. Bederson, B., and Zierlinger, A. Drag-and-pop and drag-and-pick: techniques for accessing remote screen content on touch- and pen-operated systems. In *Proceedings of the Ninth IFIP Conference on Human-Computer Interaction* (Zurich, Switzerland, September 1-5, 2003). INTERACT '03. IOS Press, 57-64.

[5] Biehl, J. T. and Bailey, B. P. 2004. ARIS: an interface for application relocation in an interactive space. In *Proceedings of Graphics interface 2004* (London, Ontario, Canada, May 17 - 19, 2004). ACM International Conference Proceeding Series, vol. 62. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, 107-116.

[6] Biehl, J. T., Golovchinsky, G., and Lyons, K. 2008. Beyond the laboratory: supporting authentic collaboration with multiple displays. Workshop at CSCW 2008. San Diego, CA. November 2008.

[7] Blomberg, J., Giacomi, J., Mosher, A., and Swenton-Wall, P. 1993. Ethnographic field methods and their relation to design. In D. Dchuler and A. Namioka (Eds.) *Participatory Design: Principles and Practices*. Erlbaum: New Jersey.

[8] Booth, K. S., Fisher, B. D., Lin, C. J., and Argue, R. 2002. The "mighty mouse" multi-screen collaboration tool. In *Proceedings of the 15th Annual ACM Symposium on User interface Software and Technology* (Paris, France, October 27 - 30, 2002). UIST '02. ACM, New York, NY, 209-212.

[9] Brignull, H. and Rogers, Y. 2003. Enticing People to Interact with Large Public Displays in Public Spaces. In *Proceedings of the Ninth IFIP Conference on Human-Computer Interaction* (Zurich, Switzerland, September 1-5, 2003). INTERACT '03. IOS Press, 17-24.

[10] Brignull, H., Izadi, S., Fitzpatrick, G., Rogers, Y., and Rodden, T. 2004. The introduction of a shared interactive surface into a communal space. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM, New York, NY, 49-58.

[11] Buxton, W. 1995. Integrating the periphery and context: a new model of telematics. In *Proceedings of Graphics Interface 1995*, Morgan-Kaufman, New York, 239-246.

[12] Buxton, W. 2007. *Sketching User Experiences*. Morgan Kaufmann: New York.

[13] Cherubini, M., Venolia, G., DeLine, R., and Ko, A. J. 2007. Let's go to the whiteboard: how and why software developers use drawings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA, April 28 - May 03, 2007). CHI '07. ACM, New York, NY, 557-566.

[14] Churchill, E. F., Nelson, L., Denoue, L., Helfman, J., and Murphy, P. 2004. Sharing multimedia content with interactive public displays: a case study. In *Proceedings of the 5th Conference on Designing interactive Systems: Processes, Practices, Methods, and Techniques* (Cambridge, MA, USA, August 01 - 04, 2004). DIS '04. ACM, New York, NY, 7-16.

[15] Cockburn, A. J. and Thimbleby, H. 1991. A reflexive perspective of CSCW. *SIGCHI Bull. 23*, 3 (Jul. 1991), 63-68.

[16] Covi, L., Olson, J. S., Rocco, E., Miller, W. J., and Allie, P. 1998. A Room of Your Own: What Do We Learn about Support of Teamwork from Assessing Teams in Dedicated Project Rooms?. In *Proceedings of the First international Workshop on Cooperative Buildings, integrating information, Organization, and Architecture*. N. A. Streitz, S. Konomi, and H. J. Burkhardt, Eds. Lecture Notes In Computer Science, vol. 1370. Springer-Verlag, London, 53-65.

[17] Crabtree, A., Hemmings, T. and Rodden, T. 2003. Social construction of displays, *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies* (eds. O'Hara, K. et al.), pp. 170-190, Dordrecht: Kluwer Academic Publishers.

[18] Dietz, P. and Leigh, D. 2001. DiamondTouch: a multi-user touch technology. In *Proceedings of the 14th Annual ACM Symposium on User interface Software and Technology* (Orlando, Florida, November 11 - 14, 2001). UIST '01. ACM, New York, NY, 219-226.

[19] Elliot, K., Neustaedter, C., and Greenberg, S. 2005. Time, ownership and awareness: the value of contextual locations in the home. *In Proceedings of the 7th International Conference on Ubiquitous Computing* (Tokyo, Japan, September 11-14, 2005). M. Beigl, S. Intille, and J. Rekimoto, Eds. . Lecture Notes in Computer Science, vol. 3660. Springer-Verlag, Berlin, 251-268.

[20] Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tang, J., and Welch, B. 1992. Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration. In *Proceedings of the SIGCHI*

Conference on Human Factors in Computing Systems (Monterey, California, United States, May 03 - 07, 1992). P. Bauersfeld, J. Bennett, and G. Lynch, Eds. CHI '92. ACM, New York, NY, 599-607.

[21] Erbad, A., Blackstock, M., Friday, A., Lea, R., and Al-Muhtadi, J. 2008. MAGIC broker: a middleware toolkit for interactive public displays. In *Proceedings of the 6th Annual IEEE International Conference on Pervasive Computing and Communications* (Hong Kong, Hong Kong, March 17 - 21, 2008). IEEE PerComm '08. IEEE, 509-514.

[22] Everitt, K., Shen, C., Ryall, K., and Forlines, C. 2006. MultiSpace: Enabling Electronic Document Micro-mobility in Table-Centric, Multi-Device Environments. In *Proceedings of the First IEEE international Workshop on Horizontal interactive Human-Computer Systems* (January 05 - 07, 2006). TABLETOP. IEEE Computer Society, Washington, DC, 27-34.

[23] Fass, A., Forlizzi, J., and Pausch, R. 2002. MessyDesk and MessyBoard: two designs inspired by the goal of improving human memory. In *Proceedings of the 4th Conference on Designing interactive Systems: Processes, Practices, Methods, and Techniques* (London, England, June 25 - 28, 2002). DIS '02. ACM, New York, NY, 303-311.

[24] Finke, M., Tang, A., Leung, R., and Blackstock, M. 2008. Lessons learned: game design for large public displays. In*Proceedings of the 3rd international Conference on Digital interactive Media in Entertainment and Arts* (Athens, Greece, September 10 - 12, 2008). DIMEA '08, vol. 349. ACM, New York, NY, 26-33.

[25] Foster, G. and Stefik, M. 1986. Cognoter: theory and practice of a colab-orative tool. In *Proceedings of the 1986 ACM Conference on Computer-Supported Cooperative Work* (Austin, Texas, December 03 - 05, 1986). CSCW '86. ACM, New York, NY, 7-15.

[26] Geyer, W., Vogel, J., Cheng, L., and Muller, M. 2003. Supporting activity-centric collaboration through peer-to-peer shared objects. In *Proceedings of the 2003 international ACM SIGGROUP Conference on Supporting Group Work* (Sanibel Island, Florida, USA, November 09 - 12, 2003). GROUP '03. ACM, New York, NY, 115-124.

[27] Google Earth. http://earth.google.com/.

[28] Greenberg, S. and Rounding, M. 2001. The notification collage: posting information to public and personal displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, Washington, United States). CHI '01. ACM, New York, NY, 514-521.

[29] Greenberg, S., and Roseman, M. 2002. Using a room metaphor to ease transitions in groupware. In M. Ackerman, et al. (Eds) *Sharing Expertise: Beyond Knowledge Management*, MIT Press, 203-256.

[30] Greenberg, S., Boyle, M. and LaBerge, J. 1999. PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal. *Personal Technologies 3* (1), 54-64, Elsevier.

[31] Grudin, J. 1994. Groupware and social dynamics: eight challenges for developers. *Commun. ACM 37*, 1 (Jan. 1994), 92-105.

[32] Gutwin, C. 2005. Review of "Public and situated displays: Social and interactional aspects of shared display technologies" by Kenton O'Hara, Mark Perry, Elizabeth Churchill and Daniel Russell, the Kluwer international series on computer supported cooperative work, 2003. *Comput. Supported Coop. Work 14*, 3 (Jun. 2005), 287-291.

[33] Gutwin, C. and Greenberg, S. 1998. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (Seattle, Washington, United States, November 14 - 18, 1998). CSCW '98. ACM, New York, NY, 207-216.

[34] Gutwin, C. and Greenberg, S. 1999. The effects of workspace awareness support on the usability of real-time distributed groupware. *ACM Trans. Comput.-Hum. Interact.* 6, 3 (Sep. 1999), 243-281.

[35] Gutwin, C. and Greenberg, S. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Comput. Supported Coop. Work 11*, 3 (Nov. 2002), 411-446.

[36] Ha, V., Inkpen, K. M., Whalen, T., and Mandryk, R. L. 2006. Direct Intentions: The Effects of Input Devices on Collaboration around a Tabletop Display. In *Proceedings of the First IEEE international Workshop on Horizontal interactive Human-Computer Systems* (January 05 - 07, 2006). TABLETOP. IEEE Computer Society, Washington, DC, 177-184.

[37] Han, J. Y. 2005. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th Annual ACM Symposium on User interface Software and Technology* (Seattle, WA, USA, October 23 - 26, 2005). UIST '05. ACM, New York, NY, 115-118.

[38] Hancock, M. and Carpendale, S. 2006. The complexities of computer-supported collaboration. Technical report 2006-812-05, Dept. of Computer Science, University of Calgary, Calgary.

[39] Hawkey, K., Kellar, M., Reilly, D., Whalen, T., and Inkpen, K. M. 2005. The proximity factor: impact of distance on co-located collaboration. In *Proceedings of the 2005 international ACM SIGGROUP Conference on Supporting Group Work* (Sanibel Island, Florida, USA, November 06 - 09, 2005). GROUP '05. ACM, New York, NY, 31-40.

[40] Hepsø, V., Munkvold, G., Olsen, H. H., and Rolland, K. 2008. Evolution of multi-display and distributed collaboration environments in StatiolHydro. In *CSCW 2008 Workshop: Beyond the Laboratory: Supporting Authentic Collaboration with Multiple Displays*. Organized by J. Biehl, G. Golovchinsky, and K. Lyons.

[41] Hong, J. I. and Landay, J. A. 2000. SATIN: a toolkit for informal ink-based applications. In *Proceedings of the 13th Annual ACM Symposium on User interface Software and Technology* (San Diego, California, United States, November 06 - 08, 2000). UIST '00. ACM, New York, NY, 63-72.

[42] Horvitz, E. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: the CHI Is the Limit* (Pittsburgh, Pennsylvania, United States, May 15 - 20, 1999). CHI '99. ACM, New York, NY, 159-166.

[43] Houde, S., Bellamy, R., and Leahy, L. 1998. In search of design principles for tools and practices to support communication within a learning community. *SIGCHI Bull. 30*, 2 (Apr. 1998), 113-118.

[44] Huang, E. M. and Mynatt, E. D. 2003. Semi-public displays for small, co-located groups. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA, April 05 - 10, 2003). CHI '03. ACM, New York, NY, 49-56.

[45] Huang, E. M., Mynatt, E. D., and Trimble, J. P. 2006a. Displays in the Wild: Understanding the Dynamics and Evolution of a Display Ecology. In *Proceedings of the 6th International*

Conference on Pervasive Computing (Dublin, Ireland, May 7 - 10, 2006). K. Fishkin, and B. Schiele, Eds. Lecture Notes in Computer Science, vol. 3978. Springer-Verlag, Berlin, 321-336.

[46] Huang, E. M., Mynatt, E. D., Russell, D. M., and Sue, A. E. 2006b. Secrets to Success and Fatal Flaws: The Design of Large-Display Groupware. *IEEE Comput. Graph. Appl*. 26, 1 (Jan. 2006), 37-45.

[47] Hughes, J., King, V., Rodden, T., and Andersen, H. 1994. Moving out from the control room: ethnography in system design. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, North Carolina, United States, October 22 - 26, 1994). CSCW '94. ACM, New York, NY, 429-439.

[48] Hutchins, E. 1996. *Cognition in the Wild.* MIT Press.

[49] Igarashi, T., Matsuoka, S., Kawachiya, S., and Tanaka, H. 1997. Interactive beautification: a technique for rapid geometric design. In *Proceedings of the 10th Annual ACM Symposium on User interface Software and Technology* (Banff, Alberta, Canada, October 14 - 17, 1997). UIST '97. ACM, New York, NY, 105-114.

[50] Inkpen, K., Hawkey, K., Kellar, M., Mandryk, R., Parker, K.,Reilly, D., Scott, S., and Whalen, T. Exploring Display Factors that Influence Co-Located Collaboration: Angle, Size, Number, and User Arrangement. In *Proceedings of HCI International 2005*, July 22-27, 2005, Las Vegas, NV.

[51] Isenberg, P., and Fisher, D. 2009. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum 28*, 3, 1031-1038.

[52] Isenberg, P., Tang, A., and Carpendale, S. 2008. An exploratory study of visual information analysis. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 1217-1226.

[53] Ishii, H. 1990. TeamWorkStation: towards a seamless shared workspace. In *Proceedings of the 1990 ACM Conference on Computer-Supported Cooperative Work* (Los Angeles, California, United States, October 07 - 10, 1990). CSCW '90. ACM, New York, NY, 13-26.

[54] Ishii, H., Kobayashi, M., and Grudin, J. 1992. Integration of inter-personal space and shared workspace: ClearBoard design and experiments. In *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work* (Toronto, Ontario, Canada, November 01 - 04, 1992). CSCW '92. ACM, New York, NY, 33-42.

[55] Johanson, B., Fox, A., and Winograd, T. 2002a. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *IEEE Pervasive Computing* 1, 2 (Apr. 2002), 67-74.

[56] Johanson, B., Hutchins, G., Winograd, T., and Stone, M. 2002b. PointRight: experience with flexible input redirection in interactive workspaces. In *Proceedings of the 15th Annual ACM Symposium on User interface Software and Technology* (Paris, France, October 27 - 30, 2002). UIST '02. ACM, New York, NY, 227-234.

[57] Johansen, R. 1988 *Groupware: Computer Support for Business Teams*. The Free Press.

[58] Jordan, B., and Henderson, A. 1995. Interaction analysis: foundations and practice. *J of Learning Sciences 4*, 1, 39-103.

[59] Ju, W., Lee, B. A., and Klemmer, S. R. 2008. Range: exploring implicit interaction through electronic whiteboard design. In *Proceedings of the ACM 2008 Conference on Computer*

Supported Cooperative Work (San Diego, CA, USA, November 08 - 12, 2008). CSCW '08. ACM, New York, NY, 17-26.

[60] Karahalios, K. and Donath, J. 2004. Telemurals: linking remote spaces with social catalysts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vienna, Austria, April 24 - 29, 2004). CHI '04. ACM, New York, NY, 615-622.

[61] Kruger, R., Carpendale, S., Scott, S. D., and Greenberg, S. 2003. How people use orientation on tables: comprehension, coordination and communication. In *Proceedings of the 2003 international ACM SIGGROUP Conference on Supporting Group Work* (Sanibel Island, Florida, USA, November 09 - 12, 2003). GROUP '03. ACM, New York, NY, 369-378.

[62] Luff, P. and Heath, C. 1998. Mobility in collaboration. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (Seattle, Washington, United States, November 14 - 18, 1998). CSCW '98. ACM, New York, NY, 305-314.

[63] Luff, P., Heath, C., Kuzuoka, H., Yamazaki, K., and Yamashita, J. 2006. Handling documents and discriminating objects in hybrid spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 561-570.

[64] Newman, M. W., Lin, J., Hong, J. I., and Landay, J. A. 2003. DENIM: an informal web site design tool inspired by observations of practice. *Hum.-Comput. Interact.* 18, 3 (Sep. 2003), 259-324.

[65] McCarthy, J. F. and boyd, d. m. 2005. Digital backchannels in shared physical spaces: experiences at an academic conference. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA, April 02 - 07, 2005). CHI '05. ACM, New York, NY, 1641-1644.

[66] McCarthy, J. F., Congleton, B., and Harper, F. M. 2008. The context, content & community collage: sharing personal digital media in the physical workplace. In *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work* (San Diego, CA, USA, November 08 - 12, 2008). CSCW '08. ACM, New York, NY, 97-106.

[67] McCarthy, J. F., Costa, T. J., and Liongosari, E. S. 2001. UniCast, OutCast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays. In *Proceedings of the 3rd international Conference on Ubiquitous Computing* (Atlanta, Georgia, USA, September 30 - October 02, 2001). G. D. Abowd, B. Brumitt, and S. A. Shafer, Eds. Lecture Notes In Computer Science, vol. 2201. Springer-Verlag, London, 332-345.

[68] McCarthy, J. F., McDonald, D. W., Soroczak, S., Nguyen, D. H., and Rashid, A. M. 2004. Augmenting the social space of an academic conference. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM, New York, NY, 39-48.

[69] Myers, B. A., Stiel, H., and Gargiulo, R. 1998. Collaboration using multiple PDAs connected to a PC. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (Seattle, Washington, United States, November 14 - 18, 1998). CSCW '98. ACM, New York, NY, 285-294.

[70] Miller, T. and Stasko, J. 2001. The InfoCanvas: information conveyance through personalized, expressive art. In *CHI '01 Extended Abstracts on Human Factors in Computing*

Systems (Seattle, Washington, March 31 - April 05, 2001). CHI '01. ACM, New York, NY, 305-306.

[71] Moran, T. P., Chiu, P., Harrison, S., Kurtenbach, G., Minneman, S., and van Melle, W. 1996. Evolutionary engagement in an ongoing collaborative work process: a case study. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work* (Boston, Massachusetts, United States, November 16 - 20, 1996). M. S. Ackerman, Ed. CSCW '96. ACM, New York, NY, 150-159.

[72] Moran, T. P., Saund, E., Van Melle, W., Gujar, A. U., Fishkin, K. P., and Harrison, B. L. 1999. Design and technology for Collaborage: collaborative collages of information on physical walls. In *Proceedings of the 12th Annual ACM Symposium on User interface Software and Technology* (Asheville, North Carolina, United States, November 07 - 10, 1999). UIST '99. ACM, New York, NY, 197-206.

[73] Morris, M.R., Lombardo, J., and Wigdor, D. WeSearch: Supporting Collaborative Search and Sensemaking on a Tabletop Display. In *Proceedings of CSCW 2010*, in press.

[74] Morris, M. R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. 2004. Beyond "social protocols": multi-user coordination policies for co-located groupware. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM, New York, NY, 262-265.

[75] Mynatt, E. D. The writing on the wall. 1999. In *Proceedings of the Seventh IFIP Conference on Human-Computer Interaction* (Edinburgh, Scotland, August 30 - September 3, 1999). INTERACT '99. IOS Press, 196-204.

[76] Mynatt, E. D., Igarashi, T., Edwards, W. K., and LaMarca, A. 1999. Flatland: new dimensions in office whiteboards. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: the CHI Is the Limit* (Pittsburgh, Pennsylvania, United States, May 15 - 20, 1999). CHI '99. ACM, New York, NY, 346-353.

[77] Newman, M. W., Lin, J., Hong, J. I., and Landay, J. A. 2003. DENIM: an informal web site design tool inspired by observations of practice. *Hum.-Comput. Interact. 18*, 3 (Sep. 2003), 259-324.

[78] O'Hara, K., Perry, M., and Churchill, E. 2004. *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies* (Cooperative Work, 2). Kluwer Academic Publishers.

[79] O'Hara, K., Perry, M., and Lewis, S. 2003. Social coordination around a situated display appliance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA, April 05 - 10, 2003). CHI '03. ACM, New York, NY, 65-72.

[80] Pedersen, E. R., McCall, K., Moran, T. P., and Halasz, F. G. 1993. Tivoli: an electronic whiteboard for informal workgroup meetings. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands, April 24 - 29, 1993). CHI '93. ACM, New York, NY, 391-398.

[81] Perry, M., and O'Hara, K. 2003. Display-based activity in the workplace. In *Proceedings of the Ninth IFIP Conference on Human-Computer Interaction* (Zurich, Switzerland, September 1-5, 2003). INTERACT '03. IOS Press, 591-598.

[82] Pinelle, D. and Gutwin, C. 2005. A groupware design framework for loosely coupled workgroups. In *Proceedings of the Ninth Conference on European Conference on Computer*

Supported Cooperative Work (Paris, France, September 18 - 22, 2005). H. Gellersen, K. Schmidt, M. Beaudouin-Lafon, and W. Mackay, Eds. ECSCW. Springer-Verlag New York, New York, NY, 65-82.

[83] Pousman, Z. and Stasko, J. 2006. A taxonomy of ambient information systems: four patterns of design. In *Proceedings of the Working Conference on Advanced Visual interfaces* (Venezia, Italy, May 23 - 26, 2006). AVI '06. ACM, New York, NY, 67-74.

[84] Rekimoto, J. 1997. Pick-and-drop: a direct manipulation technique for multiple computer environments. In *Proceedings of the 10th Annual ACM Symposium on User interface Software and Technology* (Banff, Alberta, Canada, October 14 - 17, 1997). UIST '97. ACM, New York, NY, 31-39.

[85] Rekimoto, J. and Saitoh, M. 1999. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: the CHI Is the Limit* (Pittsburgh, Pennsylvania, United States, May 15 - 20, 1999). CHI '99. ACM, New York, NY, 378-385.

[86] Ringel, M., Ryall, K., Shen, C., Forlines, C., and Vernier, F. 2004. Release, relocate, reorient, resize: fluid techniques for document sharing on multi-user interactive tables. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems*(Vienna, Austria, April 24 - 29, 2004). CHI '04. ACM, New York, NY, 1441-1444.

[87] Rogers, Y. & Lindley, S.E. 2004. Collaborating around vertical and horizontal large displays: Which way is best? *Interacting with Computers 16*, 1133-1152.

[88] Russell, D. M. and Gossweiler, R. 2001. On the Design of Personal & Communal Large Information Scale Appliances. In *Proceedings of the 3rd international Conference on Ubiquitous Computing* (Atlanta, Georgia, USA, September 30 - October 02, 2001). G. D. Abowd, B. Brumitt, and S. A. Shafer, Eds. Lecture Notes In Computer Science, vol. 2201. Springer-Verlag, London, 354-361.

[89] Russell, D. M., Drews, C., and Sue, A. 2002. Social Aspects of Using Large Public Interactive Displays for Collaboration. In *Proceedings of the 4th international Conference on Ubiquitous Computing* (Göteborg, Sweden, September 29 - October 01, 2002). G. Borriello and L. E. Holmquist, Eds. Lecture Notes In Computer Science, vol. 2498. Springer-Verlag, London, 229-236.

[90] Ryall, K., Forlines, C., Shen, C., and Morris, M. R. 2004. Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM, New York, NY, 284-293.

[91] Salvador, T., Scholtz, J., and Larson, J. 1996. The Denver model for groupware design. *SIGCHI Bull.* 28, 1 (Jan. 1996), 52-58.

[92] Schmidt, K. and Simone, C. 1996. Coordination mechanisms: towards a conceptual foundation of CSCW systems design.*Comput. Supported Coop. Work* 5, 2-3 (Dec. 1996), 155-200.

[93] Scott, S. D., Carpendale, M. S. T., and Habelski, S. 2005. Storage Bins: Mobile Storage for Collaborative Tabletop Displays.*IEEE Comput. Graph. Appl.* 25, 4 (Jul. 2005), 58-65.

[94] Scott, S. D., Carpendale, M. S. T., and Inkpen, K. M. 2004. Territoriality in collaborative tabletop workspaces. In*Proceedings of the 2004 ACM Conference on Computer Supported*

Cooperative Work (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM, New York, NY, 294-303.

[95] Scott, S. D., Grant, K. D., and Mandryk, R. L. 2003. System guidelines for co-located, collaborative work on a tabletop display. In *Proceedings of the Eighth Conference on European Conference on Computer Supported Cooperative Work*(Helsinki, Finland, September 14 - 18, 2003). K. Kuutti, E. H. Karsten, G. Fitzpatrick, P. Dourish, and K. Schmidt, Eds. ECSCW. Kluwer Academic Publishers, Norwell, MA, 159-178.

[96] Shen, C., Everitt, K.M., and Ryall, K. 2003. UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces. In *Proceedings of the 5th International Conference on Ubiquitous Computing* (Seattle, USA, October 12-15, 2003). A. Dey and A. Schmidt, Eds. Lecture Notes in Computer Science vol. 2864. Springer-Verlag, London, 281-288.

[97] Shoemaker, G., Tang, A., and Booth, K. S. 2007. Shadow reaching: a new perspective on interaction for large displays. In*Proceedings of the 20th Annual ACM Symposium on User interface Software and Technology* (Newport, Rhode Island, USA, October 07 - 10, 2007). UIST '07. ACM, New York, NY, 53-56.

[98] Smart Technologies, http://www.smarttech.com/.

[99] Snowdon, D. and Grasso, A. 2002. Diffusing information in organizational settings: learning from experience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Changing Our World, Changing Ourselves* (Minneapolis, Minnesota, USA, April 20 - 25, 2002). CHI '02. ACM, New York, NY, 331-338.

[100] Snyder, Carolyn. 2003. *Paper Prototyping: the fast and easy way to design and refine user interfaces*. San Francisco, CA: Morgan Kaufmann.

[101] Star S. L., and Griesemer, J. R. 1989. Institutional ecology, translations and boundary objects: amateurs and professionals in Berkeley's museum of vertebrate zoology, 1907-39. *Social Studies of Science 19*, 4, 387-420.

[102] Stefik, M., Foster, G., Bobrow, D. G., Kahn, K., Lanning, S., and Suchman, L. 1987. Beyond the chalkboard: computer support for collaboration and problem solving in meetings. *Commun. ACM 30*, 1 (Jan. 1987), 32-47.

[103] Stone, M. C., Fishkin, K., and Bier, E. A. 1994. The movable filter as a user interface tool. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Celebrating interdependence* (Boston, Massachusetts, United States, April 24 - 28, 1994). B. Adelson, S. Dumais, and J. Olson, Eds. CHI '94. ACM, New York, NY, 306-312.

[104] Strauss, A. L., and Corbin, J. 1999. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications.

[105] Streitz, N. A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. 1999. i-LAND: an interactive landscape for creativity and innovation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: the CHI Is the Limit* (Pittsburgh, Pennsylvania, United States, May 15 - 20, 1999). CHI '99. ACM, New York, NY, 120-127.

[106] Tan, D. S., Gergle, D., Mandryk, R., Inkpen, K., Kellar, M., Hawkey, K., and Czerwinski, M. 2008. Using job-shop scheduling tasks for evaluating collocated collaboration. *Personal Ubiquitous Comput. 12*, 3 (Jan. 2008), 255-267.

[107] Tang, A., Lanir, J., Greenberg, S., and Fels, S. 2009. Supporting transitions in work: informing large display application design by understanding whiteboard use. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work* (Sanibel Island, Florida, USA, May 10 - 13, 2009). GROUP '09. ACM, New York, NY, 149-158.

[108] Tang, A. and Fels, S. 2008. Four lessons from traditional MDEs. *ACM CSCW 2008 Workshop on Beyond the Laboratory: Supporting Authentic Collaboration with Multiple Displays*.

[109] Tang, A., Finke, M., Blackstock, M., Leung, R., Deutscher, M., and Lea, R. 2008. Designing for bystanders: reflections on building a public digital forum. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 879-882.

[110] Tang, A., Finke, M., Blackstock, M., Leung, R., Deutscher, M., Tain, G., and Giesbrecht, C. 2008. Designing for bystanders: reflections on building a public digital forum. *ACM CHI 2008 Workshop on Designing and Evaluating Mobile Phone-based Interaction with Public Displays.*

[111] Tang, A. 2006. Surface use in meeting room collaboration. In *Conference Companion of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY, 43-44.

[112] Tang, A., Parker, J. K., Lanir, J., Booth, K. S. and Fels, S. 2006. Studying collaborative surface use to guide large display interaction design. In *Conference Companion of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (Banff, Alberta, Canada, November 04 - 08, 2006). CSCW '06. ACM, New York, NY, 219-220.

[113] Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. 2006. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 1181-1190.

[114] Tang, J. C. 1991. Findings from observational studies of collaborative work. *Int. J. Man-Mach. Stud. 34*, 2 (Feb. 1991), 143-160.

[115] Tatar, D. G., Foster, G., and Bobrow, D. G. 1991. Design for conversation: lessons from Cognoter. In *Computer-Supported Cooperative Work and Groupware*, S. Greenberg, Ed. Academic Press Computers And People Series. Academic Press Ltd., London, UK, 55-80.

[116] Teasley, S., Covi, L., Krishnan, M. S., and Olson, J. S. 2000. How does radical collocation help a team succeed?. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work* (Philadelphia, Pennsylvania, United States). CSCW '00. ACM, New York, NY, 339-346.

[117] Thimbleby, H., Anderson, S., and Witten, I. H. 1990. Reflexive CSCW: supporting long-term personal work. Interact. Comput. 2, 3 (Nov. 1990), 330-336.

[118] Tobiasz, M., Isenberg, P., Carpendale, M. S. T. Lark: Coordinating co-located collaboration with information visualization. *IEEE Transactions on Visualization and Computer Graphics* (Proceedings Visualization / Information Visualization 2009), 15(6), November-December, 2009. To appear.

[119] Tolmie, P. and Crabtree, A. 2008. Deploying research technology in the home. In *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work* (San Diego, CA, USA, November 08 - 12, 2008). CSCW '08. ACM, New York, NY, 639-648.

[120] Trans2D. http://mail.rochester.edu/~mabernet/trans2d/.

[121] Trimble, J., Wales, R., and Gosswiler, R. 2003. NASA's MERBoard: an interactive collaborative workspace platform. *Public and situated displays: social and interactional aspects of shared display technologies*, K. O'Hara, et al. (Eds.), Springer, 18-44.

[122] Tse, E., Histon, J., Scott, S. D., and Greenberg, S. 2004. Avoiding interference: how people use spatial separation and partitioning in SDG workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM, New York, NY, 252-261.

[123] Tyler, J. R., and Tang, J. C. 2003. When can I expect an email response? A study of rhythms in email usage. In *Proceedings of the Eighth European Conference on Computer Supported Cooperative Work* (Helsinki, Finland, September 14 - 18, 2003). K. Kuutti, E. H. Karsten, G. Fitzpatrick, P. Dourish, and K. Schmidt, Eds. ECSCW. Springer, 239-258.

[124] Vogel, D. and Balakrishnan, R. 2005. Distant freehand pointing and clicking on very large, high resolution displays. In *Proceedings of the 18th Annual ACM Symposium on User interface Software and Technology* (Seattle, WA, USA, October 23 - 26, 2005). UIST '05. ACM, New York, NY, 33-42.

[125] Voida, S., Mynatt, E. D., MacIntyre, B., and Corso, G. M. 2002. Integrating Virtual and Physical Context to Support Knowledge Workers*. IEEE Pervasive Computing 1*, 3 (Jul. 2002), 73-79.

[126] Wang, J., Zhai, S., and Canny, J. 2006. Camera phone based motion sensing: interaction techniques, applications and performance study. In *Proceedings of the 19th Annual ACM Symposium on User interface Software and Technology*(Montreux, Switzerland, October 15 - 18, 2006). UIST '06. ACM, New York, NY, 101-110.

[127] Ware, C. and Lewis, M. 1995. The DragMag image magnifier. In *Conference Companion on Human Factors in Computing Systems* (Denver, Colorado, United States, May 07 - 11, 1995). I. Katz, R. Mack, and L. Marks, Eds. CHI '95. ACM, New York, NY, 407-408.

[128] Weiser, M. 1991. The computer for the 21st century. Scientific American 265, 3, 94-104.

[129] Whittaker, S. and Schwarz, H. 1999. Meetings of the Board: The Impact of Scheduling Medium on Long Term Group Coordination in Software Development. *Comput. Supported Coop. Work 8*, 3 (Jun. 1999), 175-205.

[130] Whittaker, S. and Sidner, C. 1996. Email overload: exploring personal information management of email. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Common Ground* (Vancouver, British Columbia, Canada, April 13 - 18, 1996). M. J. Tauber, Ed. CHI '96. ACM, New York, NY, 276-283.

[131] Wigdor, D., Jiang, H., Forlines, C., Borkin, M., and Shen, C. 2009. WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system. In *Proceedings of the 27th international Conference on Human Factors in Computing Systems* (Boston, MA, USA, April 04 - 09, 2009). CHI '09. ACM, New York, NY, 1237-1246.

[132] Wilson, A. D. 2004. TouchLight: an imaging touch screen and display for gesture-based interaction. In *Proceedings of the 6th international Conference on Multimodal interfaces* (State College, PA, USA, October 13 - 15, 2004). ICMI '04. ACM, New York, NY, 69-76.

[133] Xiao, Y., Lasome, C., Moss, J., Mackenzie, C. F., and Faraj, S. 2001. Cognitive properties of a whiteboard: a case study in a trauma centre. In *Proceedings of the Seventh Conference on European Conference on Computer Supported Cooperative Work* (Bonn, Germany, September

16 - 20, 2001). W. Prinz, M. Jarke, Y. Rogers, K. Schmidt, and V. Wulf, Eds. ECSCW. Kluwer Academic Publishers, Norwell, MA, 259-278.

Appendix A: Whiteboard Study Materials

Contains:

- Screenshot
- Questionnaire
- Interview Questions

Screenshot

A screenshot of the first page of the questionnaire as it was seen online by participants.

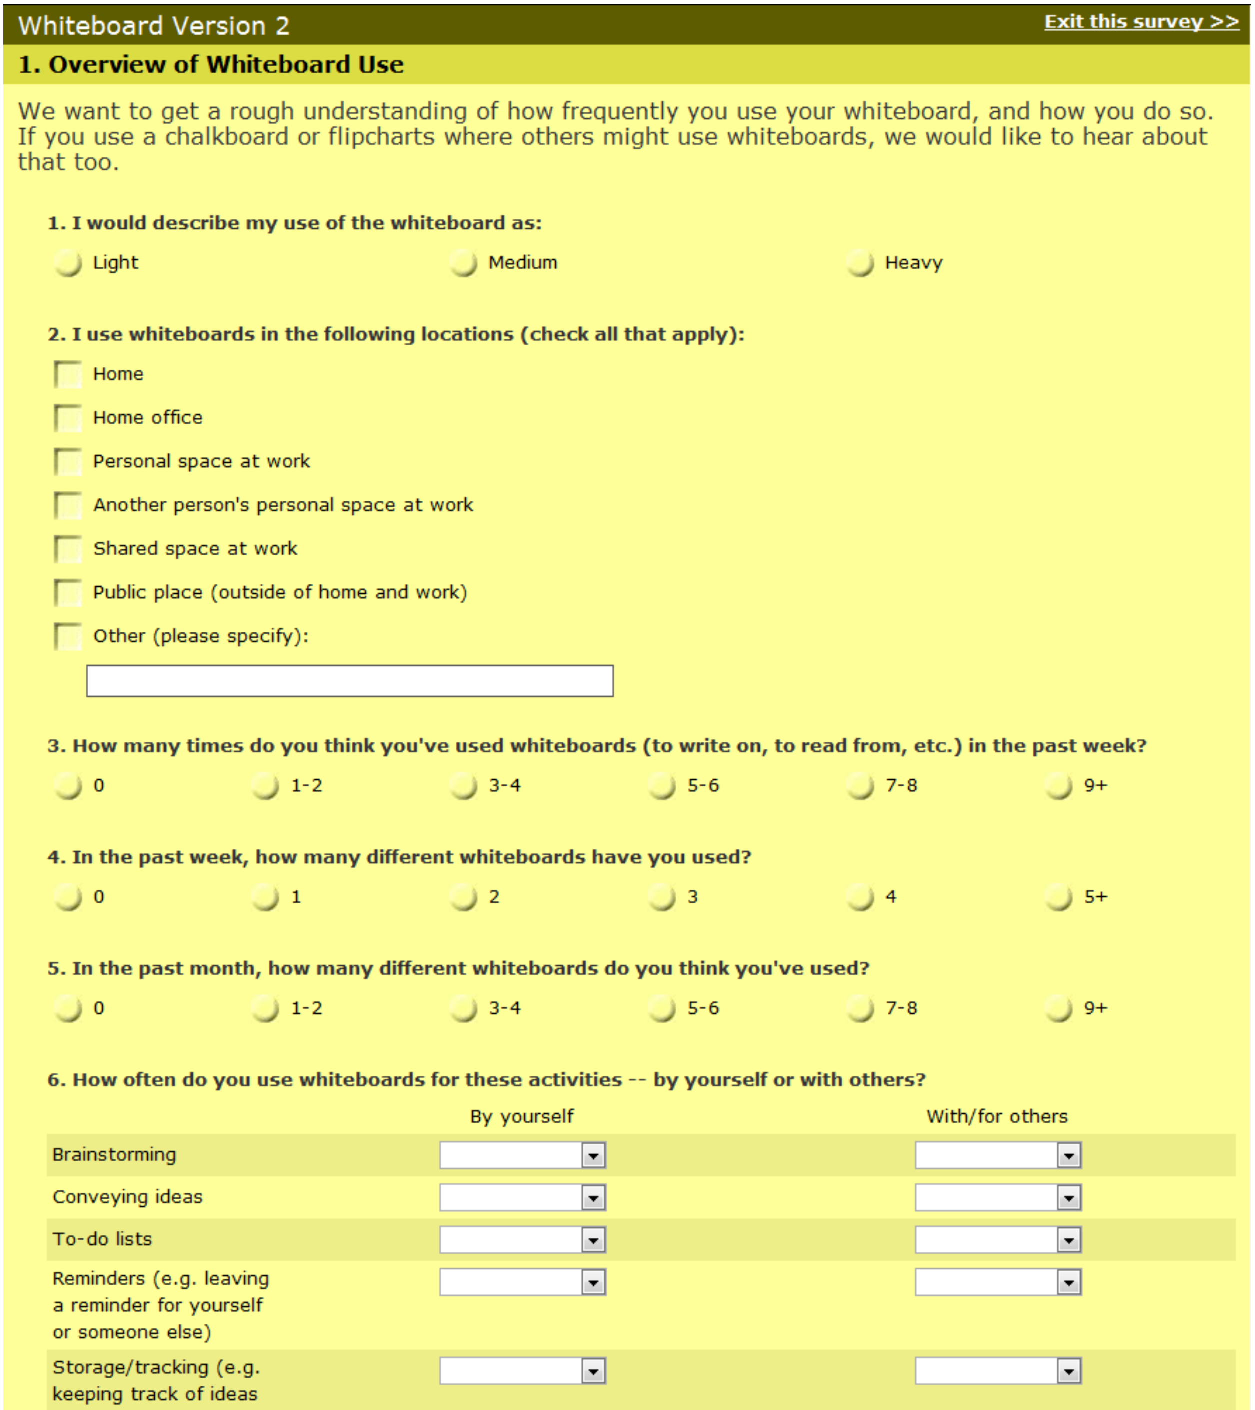

Whiteboard Version 2 Exit this survey >>

1. Overview of Whiteboard Use

We want to get a rough understanding of how frequently you use your whiteboard, and how you do so. If you use a chalkboard or flipcharts where others might use whiteboards, we would like to hear about that too.

1. I would describe my use of the whiteboard as:

Light Medium Heavy

2. I use whiteboards in the following locations (check all that apply):

- Home
- Home office
- Personal space at work
- Another person's personal space at work
- Shared space at work
- Public place (outside of home and work)
- Other (please specify):

3. How many times do you think you've used whiteboards (to write on, to read from, etc.) in the past week?

0 1-2 3-4 5-6 7-8 9+

4. In the past week, how many different whiteboards have you used?

0 1 2 3 4 5+

5. In the past month, how many different whiteboards do you think you've used?

0 1-2 3-4 5-6 7-8 9+

6. How often do you use whiteboards for these activities -- by yourself or with others?

	By yourself	With/for others
Brainstorming		
Conveying ideas		
To-do lists		
Reminders (e.g. leaving a reminder for yourself or someone else)		
Storage/tracking (e.g. keeping track of ideas or information)		

Whiteboard Questionnaire

Note that this questionnaire was conducted online. The materials here reflect the questions that were used.

1. I would describe my use of the whiteboard as:
    a. (radio) light, medium, heavy
2. I use whiteboards in the following locations:
    a. (checkbox) home office, home, my personal space at work, another's personal space at work, shared space at work, public place outside of work, elsewhere: _____
3. How many times do you think you've used whiteboards (to write on, to read from, etc.) in the past week?
    a. (radio) 1-2, 3-4, 5-6, 7-8, 9+
4. In the past week, how many different whiteboards have you used?
    a. (radio) 1, 2, 3, 4, 5+
5. In the past month, how many different whiteboards have you used?
    a. (radio) 1-2, 3-4, 5-6, 7-8, 9+
6. How often do you use whiteboards for these activities with others or by yourself? [grid matrix with radio: very often, often, sometimes, rarely, never X 2 (by myself, with others)]
    a. Brainstorming
    b. Conveying ideas
    c. To do lists
    d. Reminders (e.g. leaving a reminder for yourself or someone else)
    e. Storage/tracking
    f. Other activity: ____
    g. Other activity: ____
7. I use whiteboards with others:
    a. (radio) almost all the time, usually, half-and-half, rarely, almost never
    b. If you do use whiteboards with others…
        i. Do you copy information from notes or your laptop to the whiteboard?
            1. (checkbox) I have once—but not usually, I bring notes—though don't copy them to the whiteboard, it is too cumbersome to do that, I do this frequently, I usually just wing it on the spot, other: _____
        ii. Do you make notes, take photos, or somehow record these group activities with the whiteboard?
            1. (radio) almost all the time, usually, half-and-half, rarely, almost never
        iii. If you do record these group activities, what techniques have you used?
            1. (checkbox) taking notes in a notepad, taking notes on a computer, taking a photo of the whiteboard, other: ____
        iv. If you do record these group activities, when do you do it?
            1. (radio) during the activity, immediately after the activity, long after the activity, other: _____
        v. Do you return to these notes or recordings at a later time?
            1. (radio) almost all the time, usually, half-and-half, rarely, almost never
        vi. Have you ever put a note on the whiteboard telling others not to erase the whiteboard (after a group whiteboard activity):
            1. (checkbox) Yes—to remind myself of what I'd done earlier, Yes—to remind others of what we'd done earlier, Yes—so we could later pick up on the activity where we left off, Yes—so I could take notes from it later, Yes, reason: _____, No—I've never left content on the whiteboard purposefully
8. I use whiteboards for myself:
    a. (radio) almost all the time, usually, half-and-half, rarely, almost never

b. If you use whiteboards in a individual context,
    i. If you leave lists of things or reminders on your whiteboard, how frequently in a day do you see these items?
        1. (radio) I don't leave lists or reminders on my whiteboards, I look at the items about once a day, I look at the items a few times a day, I see the items many times a day
    ii. Are you responsible for this whiteboard (e.g. control who writes on it)?
        1. (radio) This whiteboard is primarily under my control, Yes—but I am not solely responsible (others also control the board's contents), No—someone else owns this whiteboard, Nobody really owns or controls this whiteboard (it is shared), Other: ____

9. How many whiteboards would you say you use frequently?
    a. [textbox]
10. Please consider the two whiteboards that are most important to you to answer the next questions: [Matrix x2]
    a. Where is it located?
        i. (radio) home office, home, my personal space at work, another's personal space at work, shared space at work, public place outside of work, elsewhere: ____
    b. Is this whiteboard primarily for your own activities or for group activity? || (another way of phrasing:)Do you use this whiteboard by yourself or with others?
        i. (radio) primarily for myself, usually myself—rarely involving others, about half-and-half, usually group activity—sometimes involving myself, almost always group activity, other: ____
    c. What is this whiteboard used for?
        i. (checkbox) brainstorm, todo list, reminder, conveying ideas, storage/tracking, other: ____
    d. In as much detail as possible, what is on the whiteboard right now?
        i. [textbox: open]
    e. Who wrote what is on the whiteboard right now (if there are multiple people, please check all that apply)?
        i. (checkbox) me, my family members, my close co-workers, co-workers, other: ____
    f. Roughly when was it written or put there (if there are multiple times, please check all that apply)?
        i. (checkbox) today, within seven days, within 14 days, within the last month, over a month ago, other: _____
    g. Are there identifiable "regions" on this whiteboard?
        i. (radio) yes—4+ regions, yes—2-3 regions, no—it's one big blob
    h. If there are "regions" on the whiteboard, are these regions separated in any meaningful way?
        i. (checkbox) no regions, they're separated because they look different (e.g. drawings vs. writing), they're separated by whitespace, they're separated by lines, other: ____
    i. Are any of these regions meant for specific purposes (e.g. for phone numbers):
        i. (radio) yes, no
    j. Are there regions on this whiteboard that almost never change (e.g. phone numbers):

            i. (radio) yes there are several of these regions (3+), yes there are a few of these regions (1-2), no—there is no such region
   k. Are some of these regions more important than others?
            i. (radio) yes, no
   l. Are there any reminders on this whiteboard?
            i. (radio) yes, no
   m. Are there any lists of items (e.g. a todo list) on this whiteboard?
            i. (radio) yes, no
   n. Is there anything on this whiteboard that usually does not get erased (e.g. phone number list)
            i. (radio) yes, no
   o. As you were answering those questions, were you looking at the board, or working from memory?
            i. (radio) looking at it, working from memory
   p. Is this whiteboard generally easily visible to you?
            i. (radio) It is out of the way (e.g. behind a door), I need to be very close to it to see it, I can see it from most places in the room, I can basically see it everywhere in the room
   q. Do you ever attach stuff to this whiteboard? (e.g. sticky notes, or using tape)
            i. (radio) yes, no
   r. How big is the whiteboard?
            i. (radio) size of a page of paper, size of a big computer screen, size of a regular chalkboard or whiteboard, an extra big whiteboard, other: ____

11. What is your occupation? [textbox: open]
12. Are you a student?
    a. (radio) yes, no
    b. If so, what is your year and major [textbox: open]
13. Do you attend meetings where there is a whiteboard present? (radio) yes, no
14. Do you have your own personal workspace? (radio) yes, no
15. Does your personal workspace have a whiteboard? (radio) yes, no
16. Do you live with others? (radio) yes, no
    a. Do you have a whiteboard at home? (radio) yes, no
    b. Is this whiteboard used by multiple members of your home or by yourself? (radio) only myself, mostly myself, shared with others, mostly by someone else
17. Age:
    a. (radio) 19-21, 22-25, 26-30, 30-35, 36-40, 41+
18. Can we contact you for an interview?
    a. (radio) yes, no

Interview Script

We have two primary goals in each of these interviews: understand the content and form of their whiteboard, and understand how it is used in their everyday work practices. We will focus on the "all-purpose" users since these are individuals that will have found novel ways of using the whiteboard, and/or will have *meaningfully integrated the whiteboard into their everyday practice*. We have three areas of focus with respect to interventions that we will investigate: (1) the structure and organization of persistent information; (2) the potential utility of linking whiteboard spaces/content, and (3) the organization of content reuse.

General Questions

- What do you use this whiteboard for?
- When do you look at this whiteboard?
- How frequently would you say you look at this whiteboard?
- How is this whiteboard positioned with respect to your workspace?
- Who else uses this whiteboard? What do they use it for?

Artefact Questions

- What is on the whiteboard right now?
- How long has that content been there?
- When was that generated, and who was around when it was generated?
- Was that copied <u>from</u> somewhere else?
- Was it copied <u>to</u> somewhere else afterwards?
- How many different past activities are documented on this whiteboard right now?

Whiteboard Practice

- What other whiteboards do you use in your everyday life? Include at home, at work.
- What do you use these whiteboards for, mostly?
- Is there anything about how you use whiteboards that you think is unusual?

Structure and Organization of Persistent Information

- What on this whiteboard has been here for a very long time? Why is it here? When can you remove it? Will you remove it then?
- How do you know what you can erase?

- When do you erase it as oppose to leave it or stroke it out?
- Is each piece of persistent information as important as the next? How do you identify what is important, how do you know?

Utility of Linking Whiteboard Spaces

- Was the information here copied from somewhere else? Has it been copied somewhere else?
- For each item here, how many other things is it related to or associated with (real and virtual people, things, etc.)? If you were to tackle each of these items, what will you need to do first? (*draw*)
- Have you ever encountered a situation where you wanted to see what was on this whiteboard or a different whiteboard, but were somewhere else? What did you do? What is it that you wanted to see?
- For each of the things here, to what extent is it important to be able to see all of the content? Is it the reminder that there is something there that is important to you? When you look at the whiteboard, do you register it as a bunch of regions with things inside, or do you actively read what's inside?
- Have you ever wanted to "save" whiteboard content?
- Have you ever wanted to look at a whiteboard, but not been in the right place?

Content Re-use

- Was the content used again after it was generated? Roughly how many times? Who used it, when, and who was around when it was used again?
- Will you ever copy some of this information and bring it with you somewhere else?
- If you could package up each of the items on the board, how would you package them up?
- Now imagine you could put each of these packages somewhere. When and where would you look at them? Why would you look at them there?
- Have you ever erased something on this whiteboard by accident? Roughly, what was it? How long did it take for you to discover the error? What was the context that you discovered the error? (Where were you, who was around, etc.)
- Try to imagine yourself _now_ adding some content to the whiteboard. (a TODO list, conveying an idea to a friend). Where would you add it? Would you erase something to make place? (maybe ask to do it)

Probing

- What is the hottest area on your whiteboard?
- Which piece of your whiteboard would get voted off the whiteboard, and why?
- If you were to shoot your whiteboard, where would you shoot?
- Imagine you could put the whiteboard in your hand. Where would you put it?
- If you could put your computer on one place on the whiteboard, where would you put it? What would you make it do?
- What would be the first thing you would do if your whiteboard disappeared?
- *Without looking at the whiteboard, can you write out what's on the whiteboard?*

Other Questions

- Do you ever leave anything on the whiteboard to look at later (for yourself and/or for others)? What is it?
- Why do you leave it on the whiteboard? (What makes it special?)
- How do you know you will leave it on the whiteboard?
- Are there areas of the whiteboard that you leave for this?
- How much later would you say you use it? Give me some examples.

Appendix B: Tabletop Study Materials

Contains:

- Instructions & Protocol
- Questionnaire

Route Planning Task Instructions & Protocol

Part 1: Route Planning with 4 Different Displays

You will repeat this task 4 times, each time with a different display configuration.

With each display, you will work together to plan 2 bus routes on a city map:

- One bus route goes between the two black circles marked with an "X" and the other bus route goes between the two black circles marked with an "O". These start and end points will change each time you repeat the task, but otherwise the map will remain the same.
- Your bus routes must follow the roads (black lines on the background map).
- To create your routes, tap along consecutive squares on the map. They will be coloured in to indicate the routes you have chosen.

Your routes will be assigned a penalty score that will be displayed & updated as you work. Your objective is to minimize this penalty score. To achieve this goal, you should:

- Keep the route as short as possible
- Go through commercial (blue) zones. Industrial zones (yellow) are next best.
- Go through areas with high foot traffic (indicated by red lines)
- Avoid areas with high vehicle traffic (indicated by blue lines)

You may share or divide the work any way you like.

Part 2: Design Your Space

Now that you have experienced the route planning task with several different interfaces, work together to design an interface that you think would work best for this task. You may use the display configurations you experienced and paper and pens to discuss your ideas. We are very interested to hear the reasons for your choices, so please discuss the ideas together and think out loud. The facilitators will ask questions as you work.

Instructions for the Facilitators

This is to remind us what to say & do. Add whatever you can think of.

Before part 1 begins we need to:

- Sign the consent form.
- Ask them about their experience with computers/tables/large screens and collect demographic info (i.e. age/sex/handedness)
- Hand out the written part 1 instructions.
- Demonstrate the interface for the first display, including:

    - How to draw a line to create a route
    - Introduce each map layer, so they know what each colour means
    - Explain the icons
    - Show them how to turn layers on and off or use lenses (whichever will be used first)
    - For lenses, need to generate, move, resize, and hide offscreen
    - Show them how to create, move, and resize personal spaces (if used in first session)
- Allow them 5 minutes to practice with the first display configuration

Before each of the remaining display configurations, we need to

- Describe the available features in the current interface
- Demonstrate any new features they haven't used yet
- Allow them 5 minutes to practice with the current display configuration

After Part 1, give instructions for part 2 (we may want to do this verbally).

Ask questions while the participants work.

Interview Questions

After part 2 (and/or during part 2 if the questions fit in), participants will be interviewed together.

- Of the 4 display conditions you tried, overall which did you like best? Why?
- Were there times when there was a different display you liked best? Why?
- Overall, which display condition did you like least? Why?
- How did you divide up the work? Did you do this differently in the different conditions?
- Which condition made it easiest for you to coordinate your work? Why?
- What did you find most difficult about this task?

At the end, pay the $10. Provide a general debriefing of the study and its methods!

Questionnaire

Instructions
Please respond to all of the items listed below.
Questions

- How much time do you spend working with a computer every week? ____________ hours per week

- List the three most frequent activities you perform with a computer. How many hours per week do you spend doing each of these activities?

  1. ____ Time Spent: ______ hours per week
  2. ____ Time Spent: ______ hours per week
  3. ____ Time Spent: ______ hours per week

- Do you have any experience with large-screen computer displays (i.e. CAVEs, tiled wall displays) or special-format displays (i.e. tabletop displays)? If so, what kind of experience?

Check One:	**Kind of Experience:**
[ ] Yes [ ] No	

- Do you have any experience with computerized spatial mapping tools (e.g. Mapquest, Google Maps, or GIS tools) or have you taken a university-level course in geography or geographic information systems (GIS)? If so, what kind of experience?

Check One:	**Kind of Experience:**
[ ] Yes [ ] No	

- How much time do you spend playing computer and/or video games (i.e. Xbox, PS2) every week?

Less than 2 hours	2 to 4 hours	4 to 6 hours	6 to 8 hours	More than 8 hours

- Are you left or right handed? ____________________
- What is your age? __________________
- What is your gender? [ ] Male [ ] Female

Appendix C: Scheduling Tool Study Materials

Contains:

- Paper Prototype Materials
    - Sample agenda view
    - Sample calendar view
    - Sample timeline view
- Study Protocol

Sample Agenda View

Monday, Sept 7
Electrician - Bedroom
Carpenter - Master Bedroom
Plumber - New Washroom

Tuesday, Sept 8
Electrician - Kitchen
Carpenter - Master Bedroom
Plumber - New Washroom

Wednesday, Sept 9
Carpenter - Bedroom
Electrician - Master Bedroom

Thursday, Sept 10
Contractor - Master Bedroom

Friday, Sept 11
Contractor - Kitchen
Electrician - New Washroom

Monday, Sept 14
Electrician - Kitchen
Carpenter - New Washroom

Tuesday, Sept 15
Contractor - Bedroom
Carpenter - New Washroom

Wednesday, Sept 16
Plumber - Kitchen

Thursday, Sept 17
Contractor - New Washroom
Plumber - Kitchen

Friday, Sept 18
Carpenter - Kitchen

Monday, Sept 21
Contractor - Kitchen

Tuesday, Sept 22
Contractor – Kitchen

Sample Calendar View

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
		1	2	3	4	5
6	7 Electrician - Bedroom Carpenter - Master Bedroom Plumber – New Washroom	8 Electrician - Kitchen	9 Carpenter - Bedroom Electrician – Master Bedroom	10 Contractor – Master Bedroom	11 Contractor – Kitchen Electrician – New Washroom	12
13	14 Electrician - Kitchen Carpenter – New Washroom	15 Contractor – Bedroom	16 Plumber - Kitchen	17 Contractor – New Washroom	18 Carpenter - Kitchen	19
20	21 Contractor - Kitchen	22	23	24 House Party	25	26
27	28	29	30			

Sample Timeline View

<table>
<tr><th>Electrician</th><th>Carpenter</th><th>Plumber</th><th>Contractor</th></tr>
<tr><td>Bedroom</td><td rowspan="2">Master Bedroom</td><td rowspan="2">New Washroom</td><td></td></tr>
<tr><td>Kitchen</td><td></td></tr>
<tr><td>Master Bedroom</td><td>Bedroom</td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td>Master Bedroom</td></tr>
<tr><td>New Washroom</td><td></td><td></td><td>Kitchen</td></tr>
<tr><td>Kitchen</td><td rowspan="2">New Washroom</td><td></td><td></td></tr>
<tr><td></td><td></td><td>Bedroom</td></tr>
<tr><td></td><td></td><td rowspan="2">Kitchen</td><td></td></tr>
<tr><td></td><td></td><td>New Washroom</td></tr>
<tr><td></td><td>Kitchen</td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td rowspan="2">Kitchen</td></tr>
<tr><td></td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td></td></tr>
</table>

Study Protocol

1. Introduce self, and that the study is about designing electronic whiteboard systems—particularly, interaction techniques for such systems, and that their feedback is used to better inform the design of such techniques.
2. Introduce scheduling task:

Imagine that you just bought a brand new two floor house. You really like the house: it's got everything you ever wanted, the location is great, and the colour is perfect, and your neighbours are friendly. You did, however have some renovations in mind when you saw the layout of the house. Your job is to schedule these renovations. Each renovation plan requires a different set of individuals working on the job, and each task may require different amounts of time. What is important is to maintain the *order* that the workers come in to work on the task. *[Go through an example]*

3. Depending on condition:

Paper+Calendar: Your task is to select three of these renovations, and to schedule them into the calendar. The workers can start on Sept 7, but do not work weekends. The work must be complete by Sept 24, as this is the day you are having a house party. Try to each of the renovations complete done as quickly as possible: both overall, and individually.

Paper Prototype: You will be using a paper "system" to do the scheduling. Here, you have already completed scheduling four of the six renovations—you only need to add the last two. This prototype mimics a computer-based system, and so there are some interaction techniques that I will describe to you.

- Adding an item to the schedule ← Plumber, yada →
- View change & Filtering
- Copying

4. During task, help them out to complete the task—the purpose is not to have them solve the problem, but to have them experience what it would be like to use the system.
5. Once the task is complete, have them answer several questions regarding the data that is in the calendar.

- The contractor would like to know what his schedule is—he has come up to your whiteboard, and is asking you to show him so he can add it to his own day planner.
- The electrician has called, and would like to know what his schedule is.
- How many days does it take to complete the [kitchen]?
- What is the scheduled order of workers for the [main washroom]

6. Ask them experience questions.

Paper+Calendar:

> What did you think of this task? What made it difficult? What made answering the questions difficult? If this was on a whiteboard, how could the system have helped you to solve your problem? How could it have helped to answer the questions?

Paper Prototype:

> What did you think of this task? What made it difficult? What made answering the questions difficult? Were the additional views useful in helping you answer questions? Were the additional views useful in helping you to complete the task? How could the views have been changed to help you complete the task or to answer the questions? What would you do differently? Were the interaction techniques easy to learn? How could they have been different?

Appendix D: UBC Research Ethics Board Certificates

Contains:

- Approval for: (H05-81061) B05-1061 - Understanding Tabletop Work Practices
- Approval for: H07-01773 - Exploring Adoption Patterns of Large Displays in Public Settings
- Approval for: H08-02511 - User-Centered Design of an Electronic Whiteboard

The University of British Columbia
Office of Research Services
Behavioural Research Ethics Board
Suite 102, 6190 Agronomy Road,
Vancouver, B.C. V6T 1Z3

CERTIFICATE OF APPROVAL - FULL BOARD

PRINCIPAL INVESTIGATOR:	INSTITUTION / DEPARTMENT:	UBC BREB NUMBER:
Sidney S. Fels	UBC/Applied Science/Electrical and Computer Engineering	H07-01773

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Vancouver (excludes UBC Hospital)

Other locations where the research will be conducted:
On campus, indoor locations where people (students, staff, faculty, public) congregate: e.g. near a coffee shop, in a public concourse in the student union building, etc.

CO-INVESTIGATOR(S):
Rock Leung
Matthias Finke
David Vogt
Mike Blackstock
Rodger J. Lea
Kellogg S. Booth
Sidney S. Fels

SPONSORING AGENCIES:
Canadian Heritage

PROJECT TITLE:
Exploring Adoption Patterns of Large Displays in Public Settings

REB MEETING DATE:	CERTIFICATE EXPIRY DATE:
August 9, 2007	**August 9, 2008**

DOCUMENTS INCLUDED IN THIS APPROVAL:	DATE APPROVED:
	August 31, 2007

Document Name	Version	Date
Consent Forms:		
Consent to be interviewed, photographed and videotaped	2007-08-24	August 24, 2007
Advertisements:		
Recruitment Flyer	2007-08-24	August 24, 2007
Participant Information Sheet	2008-08-24	August 24, 2007
Questionnaire, Questionnaire Cover Letter, Tests:		
Questionnaire Items	2007-08-24	August 24, 2007
Interview Script	2007-08-24	August 24, 2007
Other Documents:		
Entire Appendix (All documents)	2007-08-24	August 24, 2007

The application for ethical review and the document(s) listed above have been reviewed and the

procedures were found to be acceptable on ethical grounds for research involving human subjects.
Approval is issued on behalf of the Behavioural Research Ethics Board and signed electronically by one of the following: Dr. Peter Suedfeld, Chair Dr. Jim Rupert, Associate Chair Dr. M. Judith Lynam, Associate Chair Dr. Laurie Ford, Associate Chair

The University of British Columbia
Office of Research Services
Behavioural Research Ethics Board
Suite 102, 6190 Agronomy Road,
Vancouver, B.C. V6T 1Z3

CERTIFICATE OF APPROVAL- MINIMAL RISK RENEWAL

PRINCIPAL INVESTIGATOR:	DEPARTMENT:	UBC BREB NUMBER:
Sidney S. Fels	UBC/Applied Science/Electrical and Computer Engineering	H05-81061

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Vancouver (excludes UBC Hospital)

Other locations where the research will be conducted:
N/A

CO-INVESTIGATOR(S):
Joel Lanir
Kellogg S. Booth
Anthony Tang
Karen Parker

SPONSORING AGENCIES:
Unfunded Research - "Understanding Tabletop Work Practices"

PROJECT TITLE:
Understanding Tabletop Work Practices

EXPIRY DATE OF THIS APPROVAL: December 1, 2009

APPROVAL DATE: December 1, 2008

The Annual Renewal for Study have been reviewed and the procedures were found to be acceptable on ethical grounds for research involving human subjects.

Approval is issued on behalf of the Behavioural Research Ethics Board

Dr. M. Judith Lynam, Chair
Dr. Ken Craig, Chair
Dr. Jim Rupert, Associate Chair
Dr. Laurie Ford, Associate Chair
Dr. Daniel Salhani, Associate Chair
Dr. Anita Ho, Associate Chair

The University of British Columbia
Office of Research Services
Behavioural Research Ethics Board
Suite 102, 6190 Agronomy Road,
Vancouver, B.C. V6T 1Z3

CERTIFICATE OF APPROVAL - FULL BOARD

PRINCIPAL INVESTIGATOR:	**INSTITUTION / DEPARTMENT:**	**UBC BREB NUMBER:**
Sidney S. Fels	UBC/Applied Science/Electrical and Computer Engineering	H08-02511

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Vancouver (excludes UBC Hospital)

Other locations where the research will be conducted:
N/A

CO-INVESTIGATOR(S):
Kellogg S. Booth

SPONSORING AGENCIES:
N/A

PROJECT TITLE:
User-Centered Design of an Electronic Whiteboard

REB MEETING DATE:	**CERTIFICATE EXPIRY DATE:**
November 27, 2008	**November 27, 2009**

DOCUMENTS INCLUDED IN THIS APPROVAL:

DATE APPROVED:
November 27, 2008

Document Name	Version	Date
Protocol:		
Design Session Protocol	N/A	November 5, 2008
Consent Forms:		
Consent Forms	N/A	November 5, 2008
Advertisements:		
Recruitment Poster	N/A	November 5, 2008
Questionnaire, Questionnaire Cover Letter, Tests:		
Sample Interview Questions	N/A	November 5, 2008
Sample Design Tasks	N/A	November 5, 2008
Other Documents:		
All Appendices	N/A	November 5, 2008

The application for ethical review and the document(s) listed above have been reviewed and the procedures were found to be acceptable on ethical grounds for research involving human subjects.

Approval is issued on behalf of the Behavioural Research Ethics Board
and signed electronically by one of the following:

Dr. M. Judith Lynam, Chair
Dr. Ken Craig, Chair
Dr. Jim Rupert, Associate Chair
Dr. Laurie Ford, Associate Chair
Dr. Daniel Salhani, Associate Chair
Dr. Anita Ho, Associate Chair