

Supporting Transitions in Work: Informing Large Display Application Design by Understanding Whiteboard Use

Anthony Tang¹, Joel Lanir², Saul Greenberg³, Sidney Fels¹

¹Human Communication Technologies Laboratory

²Department of Computer Science

University of British Columbia

2329 West Mall, Vancouver, BC, Canada

hct.ece.ubc.ca

{tonyt@ece | yoel@cs | ssfels@ece}.ubc.ca

³Department of Computer Science

University of Calgary

2500 University Dr NW, Calgary, AB, Canada

grouplab.cpsc.ucalgary.ca

saul@cpsc.ucalgary.ca

ABSTRACT

In this paper, we explore the practice of using a whiteboard for multiple tasks, and specifically how users employ whiteboards to smoothly *transition between related sets of tasks*. Our study underscores several basic, but important affordances of whiteboards that support this practice, including visual persistence, flexibility of interaction primitives, and their situated physicality. We discuss the implications of these findings for the design of large display applications.

Categories and Subject Descriptors

H.5.3. [Information Interfaces and Presentation (e.g., HCI)]: Group and Organization Interfaces—CSCW.

General Terms

Human Factors, Design

Keywords

Whiteboard, large display groupware, reflexive CSCW.

1. INTRODUCTION

Researchers have long been interested in the use of large displays to support work—both collaborative activity (e.g. [1],[3],[5],[7],[10],[11],[12],[21],[23]) and independent activity (e.g. [17],[18],[29]). The inherent size and placement of these displays afford different (and in many cases, new) forms of interaction and use practices from their desktop counterparts [10]. Yet, adoption or large display groupware has been limited: even the most willing users experience difficulties when trying to adopt novel technologies into their everyday practice [11].

Our view is that while users may need to interact with large displays in novel ways, large display applications should support the kinds of *tasks* and *work practices* that users are already accustomed to with analogous traditional display surfaces. Since others have already reported on the task space

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GROUP '09, May 10–13, 2009, Sanibel Island, Florida, USA.

Copyright 2009 ACM 978-1-60558-500-0/09/05...\$5.00.

	<i>Independent</i>	<i>Collaborative</i>
<i>Sync</i>	<i>Worker</i> <ul style="list-style-type: none"> • Word processor • Spreadsheet • CAD software 	<i>Real-time interaction</i> <ul style="list-style-type: none"> • Telephone • Video conferencing • Instant messaging
<i>Async</i>	<i>Personal Management</i> <ul style="list-style-type: none"> • PIM, schedule, agenda, task list • Reminders, post-it notes 	<i>Ongoing tasks</i> <ul style="list-style-type: none"> • Team rooms • Bulletin boards • Email

Figure 1. A modified groupware matrix that emphasizes modes of activity, and the tools that support them.

of traditional displays such as whiteboards (e.g. [26],[17],[32],[2],[19]), our goal was to identify work practices common to traditional display surfaces that would help inform the design of large display applications. In this work, we are particularly interested in studying and understanding the use of the whiteboard artefact—in part due to its ubiquity in the common office environment, and because of its success as a “display device” in the traditional environment. We view whiteboards as an important design resource: the work practices that users have developed around these displays are suggestive of the immediate ways in which users may desire to appropriate novel large display technology.

While there have been some studies of whiteboards, both as a collaborative medium ([19],[26],[32]) and in personal office spaces [17], our study results in a detailed, focused examination of whiteboard use to support *transitions between different tasks and modes of activity*. To begin, our findings will show that most whiteboard tasks can be neatly categorized into a simple 2×2 matrix (Figure 1), where the primary axes are synchronous vs. asynchronous work (i.e. same/different time) and independent vs. collaborative work. In some cases, whiteboard use encompasses related sets of tasks from different quadrants, and here, we will see that the whiteboard facilitates transitions between independent and collaborative activity, and synchronous and asynchronous work. Our analysis will show how the affordances of everyday situated whiteboards support users’ activities across different work modes and tasks.

While supporting the transitions as we outline in this work may not be the “silver bullet” in successful design and deployment of large display technologies, this study suggests that it is an important work practice in everyday whiteboard use, and that designers should consider some form of support where appropriate. Thus, we make three contributions in this work:

first, we provide a descriptive model of whiteboard activity; second, we provide a detailed exploration of how whiteboards support users' transitions between modes of activity, and third, we identify several design and deployment factors that designers can employ to support this transition practice with situated displays.

In the next section, we review related work that focuses on designing support for multiple modes of work, as well as work that explores traditional and situated digital displays. This review motivates the current study, which we present in the next section. We will see that the whiteboard provides a work space for both temporary and ongoing work (where work continues at a later time or perhaps by a different set of users). Finally, we discuss several design implications of these findings.

2. RELATED WORK

We begin by reviewing how researchers have characterized different modes of work, and particularly how system designs have attempted to support the transition between these modes. We then review literature exploring studies of traditional displays (e.g. whiteboards) and situated digital displays to rearticulate the motivation for our study.

2.1 Modes of Activity and Seamlessness

The standard groupware matrix's primary axes, same/different time and same/different place, define four modes of activity. 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 however, the "place" of activity is fixed and collocated. As a theoretical grounding, we have found that Thimbleby et al.'s articulation of *reflexive-CSCW* [27] (the use of CSCW systems for personal or independent work) to be more pertinent to our interests based the study we report here. Figure 1 illustrates this conceptualization.

Thimbleby et al.'s definition arises from the observation that groupware has sometimes been appropriated for *independent activity* (e.g. [30],[27]): 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 [4]. Thimbleby et al. argue pragmatically for researchers to 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. [4],[6],[8]). Existing email systems are a good example of such a tool since they do not distinguish between independent and collaborative use: people can apply the same mechanism to asynchronously communicate with others or themselves [31].

Some systems have realized the inverse: enabling users to transition from independent to collaborative activity while preserving existing work practices. For instance, TeamWorkStation [14] and ClearBoard [15] overlay a video image of a remote collaborator's drawing workspace on one's own, thereby fusing the two workspaces and allowing users to use the workspace as they would if they were working independently [14]. Both systems realize a form of seamlessness that enables users to smoothly transition from independent work to collaborative work.

TeamRooms [7] 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 [6]). This simple model affords all quadrants of Figure 1, 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 collaborative. If a room is used by an individual who works a bit, then leaves and then comes back to continue where they left off, it covers both modes of individual work.

Finally, many 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 into a collaborative space (e.g. [6]). For instance, MessyBoard [5] and Notification Collage [7] 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 seeing if the same transitions can be afforded by interactions with the same situated large display.

2.2 Traditional Whiteboards and Large Digital Displays

Large display technologies have long been seen as potential vehicles for collocated collaborative interaction (e.g. [20]). This expectation likely stems from their inherent affordance of providing a shared view to a group. While many early designs focused on synchronous group activity (e.g. [28],[23], [20],[1]), recent work has also focused on asynchronous awareness support (e.g. [3],[5],[10]). Parallel to this body of work, other researchers have explored the traditional whiteboard displays to build up design rationale for these digital technologies, both for synchronous activity (e.g. [2], [23]), and asynchronous activity (e.g. [17],[32]).

Perry & O'Hara [19] investigated the motivations for general display-based activity in personal work areas, articulating the role of displays as ready references, and as coordinating resources. Mynatt [17] focused on personal whiteboards in the office context, exploring how spatial management enabled multiple parallel tasks (e.g. reminders, quick capture, and thinking). This work revealed how space (when partitioned into segments) helps to organize work. Teasley et al. [26] provided observations of "war room" whiteboard use, reporting on teams' synchronous and asynchronous use of whiteboards for shared awareness of the team's status or current activity. Xiao et al. [32] presented a case study of an emergency room ward, illustrating how a situated whiteboard became a centralized asynchronous coordinating resource for nurses and doctors in the ward. Whittaker & Schwarz [30] showed how the properties of a material wallboard afforded fundamentally different interactions around a software engineering schedule compared to a digital scheduling tool, where the wallboard's *public physicality* engendered group processes that the digital tool did not. While these studies discuss the roles and contextualized uses of whiteboards in users' activities, we focus specifically on the *transitions* between different types of activities, showing

how whiteboards are central in multiple working activities of our users.

Interactive large display applications can largely be categorized into two classes: those designed for real-time collaboration (e.g. [22],[21],[23],[28],[16],[25]), and those designed as asynchronous awareness applications (e.g. [3],[10],[29]). In the former case, some designs facilitate the ability to transfer data between the main display and auxiliary machines as a means to support the transition between collaborative and independent activity (e.g. [1] [5],[7],[16],[25]), but again, our interest is in how the situated display can support these transitions in-place. Dynamo [1], for instance, affords the ability to post and retain “media parcels” on the display for later retrieval as a means of supporting asynchronous activity on the same display.

Kimura [29] was designed primarily for independent asynchronous use, facilitating deferral and peripheral monitoring of ongoing activities or tasks. It provided ambient display of montages (e.g. thumbnails) representing activity. Similarly, the Semi-Public Display project [10] produced several designs that provided collocated collaborators with awareness of each others’ activities, thus primarily intended for collaborative asynchronous use. In both these cases, the applications afforded means to sketch on the display (as a means to support synchronous activity); however, enabling the transition between asynchronous and synchronous activity was not the focus of the work.

This brief review shows that while some large display applications support multiple modes of activity, their designs frequently do not focus on supporting the *transitions* between the modes; instead, they generally focus on supporting subsets of the modes from Figure 1. In contrast, we have seen that traditional whiteboards enable a broad set of activities. It was this flexibility that motivated our current study: what affordances of whiteboards enable their use in all of these modes and activities? Further, what work practices do users develop to take advantage of these affordances? As we will see, the practice of *using the whiteboard to support transitions across the modes of activity* is remarkably useful, but as yet poorly supported in most large display application designs.

3. STUDY METHODOLOGY

Prior studies provide rich descriptions of contextualized uses of whiteboards (e.g. [17],[30],[19],[26],[32]), but have not conceptualized whiteboard activity in a generalized model. We shall see that the model we adopt (Figure 1) adequately describes the majority of typical whiteboard activity. This model of use immediately raises two questions: How common is activity in each of the quadrants? Further, is the relative prevalence of activity fairly consistent among users? Consequently, 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). We followed this survey with a set of in-situ interviews with the “heavy” whiteboard users in order to understand the practices of users who had deeply appropriated the whiteboard into their work practice.

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. 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. While our sample may not be representative from the entire population of whiteboard users (due to our sampling method), it still provides a data set from a reasonably broad user population.

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 [17]), 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. The survey provided us with a broad basis to understand the scope of whiteboard activity, how these users appropriated whiteboards, and about the whiteboards themselves.

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, we selected two overseas participants for interviews as a check against geographic bias. Ultimately, we conducted interviews until we felt we had exhausted the diversity of uses and were no longer learning anything new.

The one hour interviews were conducted in front of their “most important” whiteboard and audio recorded for transcription (interviews with the two overseas 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 artifact for discussion. We developed a list of questions around theme areas based on the survey, though allowed the flow of the interview to guide the dialogue, referring to the list only to ensure that all themes had been addressed. Participants were paid \$20 remuneration.

Interview Analysis. We conducted an inductive analysis of interview data, iteratively coding the interview transcripts for recurring patterns of behavioural statements (regarding the whiteboard) [24]. We then used an affinity diagramming process to group statements to derive thematic understanding of our participants’ activities.

4. FINDINGS

We begin by describing our survey data, which frames everyday whiteboard environments, and whiteboard use within the four-quadrant model introduced earlier. Drawing from 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.

Table 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.”

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

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

Users. We received 167 survey responses, of which we discarded 32 due to incompleteness: we therefore report on 135 complete responses.

Whiteboard use. Table 1 shows 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 compares self-rated frequency of use to a variety of tasks, showing that heavy users appropriate the whiteboard more broadly for independent use than 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 [17], the age of the content on the whiteboards, and so forth. We collected data on 239 such whiteboards. Table 2 shows the location of these boards. Of particular note is that while many of these whiteboards are located in “collaborative” contexts (i.e.

shared workplace area), about half (shaded) are located in personal spaces. Our data on who uses these whiteboards is convergent on this point: users of important whiteboards are primarily limited to those we know fairly well.

4.1 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 4 provides examples of our classification which we elaborate on next. Strikingly, over half of the whiteboards contained deliberately un-erased content for later, asynchronous use.

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%): This type of activity involved a user deliberately putting or leaving information with the intent of using it at a later time for his/her own use. This information was used 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%): 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.

Table 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).

	<i>Independent</i>					<i>Collaborative</i>					
	Brainstorm	Task list	Reminder	Storage	Other	Brainstorm	Conveying Ideas	Task list	Reminder	Storage	Other
Heavy	4	4	4	3	2	4	5	3	2	3	4
Medium	3	3	3	3	0.5	4	4	3	2	2	0
Light	2	1	1	2	0	3	3	1	2	2	0

Table 4. Examples of user reported whiteboard contents, classified in our modified groupware matrix.

	<i>Independent</i>	<i>Collaborative</i>
<i>Sync</i>	<ul style="list-style-type: none"> • “Flow (boxes and arrows) of a presentation I am about to give” • “A mind map of my current largest project” 	<ul style="list-style-type: none"> • “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”
<i>Async</i>	<ul style="list-style-type: none"> • “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” 	<ul style="list-style-type: none"> • “All active projects and their schedules” • “Action items (tasks for team members) from the team meeting”

Collaborative asynchronous (26%): 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.

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 allowed users to transition 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.

4.2 Using Whiteboard Artefacts to Transition across Modes

Our in-situ interviews were thus 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 whiteboards allowed users to easily *transition between* these activities (and quadrants). Of the 11 interviews we conducted, 5 users had created representations (an integral, 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. [17]). Drawing on three vignettes from the former group, we illustrate how users employed the whiteboard to transition across multiple modes of activity.

4.2.1 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 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 being 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 weeklong 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 by both individuals and the group later 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.

This 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.

4.2.2 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, 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 therefore functions as a sort of 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, help him transition into “personal-management” mode to help organize his activity, and then back to resume his thinking activity.

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 2). 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.

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 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 juggled 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).



Figure 2. Part of Jill’s scheduling whiteboard, which is visible from her desk.

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.

Summary: This sample of vignettes illustrates three instances of how users employ the whiteboard 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 them.

4.3 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 2), where they were located (Table 3), 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) 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 *artifact* 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.

Semi-public whiteboards (27%) tended to be in shared location (such as in 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. [26], [30]).

Personal whiteboards (32%) were located primarily in users’ personal workspace, and were therefore primarily used by the user in question (e.g. [17]). 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). 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%) 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 the specific purpose (e.g. a fridge whiteboard for messages or grocery list).

Intuitively, we would expect that with a smaller set of users or a well-known set of users around a whiteboard, a practice would evolve that allows that set of users 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 (e.g. “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




Figure 3. Percentage of whiteboards reported to be used for various activities.

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 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 2 provides a visual example where the use number-letter combinations to represent vacation schedules, and colour to represent task type. Observations of “inside jokes” and differentiated usage between groups with MessyBoard accord with this interpretation [5]. Thus, only on whiteboards with a fairly closed set of users would we expect deep appropriation for the transitioning practice described earlier. 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) shows us that both the *physical context* (where it is located, and what is it nearby [19]) and its *social context* (who uses this whiteboard) shape how the tool is used and perceived by its users.

4.4 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 know this 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 [19]): 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.

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.

5. DISCUSSION & 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 4). Beyond this, however, we have seen that the whiteboard facilitates transitions between different modes of activity. In this section, we 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 [7], 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 [7], we also see the whiteboard as functioning as a *container*. 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 digital displays. For instance, Huang et al. [11] 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. [5],[7],[22]) means that the information is not reliably visually accessible in the same way.

Beyond the virtual metaphor in [7], the whiteboard has physical embodiment, and is *contextually located* near or in a place where action takes place [13]. Thus, the whiteboard *limits*

access to people who would likely be in the context [10], and *aids interpretation* by being in the same context ([17],[19]). Fass et al. provide an instructive example [5]: 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. [30], 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 vs. digital displays. 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. [17], [2], [32]), 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 [17]), and also via the representations that users choose to use. 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 [7]: notifications, reminders, peri-synchronous and synchronous conversation. In contrast to other systems that provided more sophisticated and integrated interaction (e.g. [28],[21]), we see that in many cases, simple, understandable metaphors can be easily appropriated by users for other unintended purposes (e.g. [3] for posting personnel schedules, or [5] for simple game play).

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 [18]. Flatland retains the 'one fixed page' metaphor of a whiteboard, providing "scaling" capabilities rather than relying on a file-based storage metaphor (e.g. [22]) or switching interfaces for different applications altogether (e.g. [12]). An alternative approach is to organize interaction around shared artefacts (e.g. [6], [7]), 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 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*

6. CONCLUSIONS

In designing large digital display applications, we need to be aware of and ideally support users' existing work practices with analogous technology—traditional whiteboards. In this paper, we have explored how users employ whiteboards to accomplish tasks, and have observed that beyond individual tasks, users employ whiteboards for activities and tasks that span the independent/collaborative and synchronous/asynchronous boundaries. From here, we have derived an understanding of whiteboard affordances that support this use: visual persistence, flexibility of the ink primitives, and its situated social and physical context. While supporting the transitions we have described in large display applications may not lead to immediate adoption of large displays, doing so will allow users to apply their existing practices and more easily appropriate novel large display applications.

7. ACKNOWLEDGEMENTS

We thank Leah Findlater and Carman Neustaedter for their thoughts on earlier drafts. We also acknowledge discussions with Kelly Booth, Kirstie Hawkey, Karen Parker, and Garth Shoemaker that helped contribute to some of the ideas in this paper. Finally, we thank NSERC and NECTAR for funding this work.

8. REFERENCES

- [1] Brignull, H., Izadi, S., Fitzpatrick, G., Rogers, Y., and Rodden, T. The introduction of a shared interactive surface into a communal space. In *Proc. CSCW 2004*, 49-58.
- [2] Cherubini, M., Venolia, G., DeLine, R., and Ko, A. J. Let's go to the whiteboard: how and why software developers use drawings. In *Proc. CHI 2007*, 557-566.
- [3] Churchill, E., Nelson, L. and Denoue, L. Digital bulletin boards for social networking. *CSCW 2002 Workshop Paper*, (New Orleans, Louisiana).
- [4] Cockburn, A. J. and Thimbleby, H. A reflexive perspective of CSCW. *SIGCHI Bull* 23, 3 (1991), 63-68.
- [5] Fass, A., Forlizzi, J., and Pausch, R. MessyDesk and MessyBoard: two designs inspired by the goal of improving human memory. In *Proc. DIS 2002*, 303-311.
- [6] Geyer, W., Vogel, J., Cheng, L-T, and Muller, M. Supporting activity-centric collaboration through peer-to-peer shared objects. In *Proc GROUP 2003*, 115-124.
- [7] Greenberg, S. and Rounding, M. The notification collage: posting information to public and personal displays. In *Proc. CHI '01*, ACM Press (2001), 514-521.
- [8] Greenberg, S., and Roseman, M. Using a room metaphor to ease transitions in groupware. In M. Ackerman, et al. (Eds) *Sharing Expertise: Beyond Knowledge Management*, MIT Press (2003), 203-256.
- [9] Gutwin, C., and Greenberg, S. A descriptive framework of workspace awareness for real-time groupware. *CSCW 11*, 3 (2002), 411-446.
- [10] Huang, E. M. and Mynatt, E. D. Semi-public displays for small, co-located groups. In *Proc. CHI 2003*, ACM Press (2003), 49-56.
- [11] Huang, E. M., Mynatt, E. D., Russell, D. M., and Sue, A. E. Secrets to success and fatal flaws: the design of large-display groupware. *IEEE CG&A* 26, 1 (2006), 37-45.
- [12] Huang, E. M., Mynatt, E. D., and Trimble, J. P. Displays in the wild: understanding the dynamics and evolution of a display ecology. In *Proc. PERVASIVE 2006*, Springer-Verlag (2006), 321-336.
- [13] Hutchins, E. *Cognition in the Wild*. MIT Press (1996).
- [14] Ishii, H. TeamWorkStation: towards a seamless shared workspace. In *Proc. CSCW 1990*, 13-26.
- [15] Ishii, H., Kobayashi, M., and Grudin, J. Integration of inter-personal space and shared workspace: clearboard design and experiments. In *Proc. CSCW 1992*, ACM Press (1992), 33-42.
- [16] Moran, T. P., Chiu, P., Harrison, S., Kurtenbach, G., Minneman, S., and van Melle, W. Evolutionary engagement in an ongoing collaborative work process: a case study. In *Proc. CSCW 1996*, 150-159.
- [17] Mynatt, E. D. The writing on the wall. In *Proc. INTERACT 1999*, IOS Press (1999), 196-204.
- [18] Mynatt, E. D., Igarashi, T., Edwards, W. K., and LaMarca, A. Flatland: new dimensions in office whiteboards. In *Proc. CHI 1999*, ACM Press (1999), 346-353.

- [19] Perry, M., and O'Hara, K. Display-based activity in the workplace. In *INTERACT '03*, IFIP (2003), 591-598.
- [20] Rønby-Pedersen, E. R., McCall, K., Moran, T. P., and Halasz, F. G. Tivoli: an electronic whiteboard for informal workgroup meetings. In *Proc INTERACT '93 and CHI '93*, ACM Press (1993), 391-398.
- [21] Russell, D. M., Drews, C., and Sue, A. Social aspects of using large public interactive displays for collaboration. In *Proc. UBICOMP 2002*, Springer (2002), 229-236.
- [22] Smart Technologies, <http://www.smarttech.com/>.
- [23] Stefik, M., Foster, G., Bobrow, D. G., Kahn, K., Lanning, S., and Suchman, L. Beyond the chalkboard: computer support for collaboration and problem solving in meetings. *Commun. ACM* 30, 1 (1987), 32-47.
- [24] Strauss, A. L., and Corbin, J. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications (1999).
- [25] Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. i-LAND: an interactive landscape for creativity and innovation. In *Proc CHI 1999*, ACM Press (1999), 120-127.
- [26] Teasley, S., Covi, L., Krishnan, M. S., and Olsen, J. S. How does radical collocation help a team succeed? In *Proc. CSCW 2000*, ACM Press (2000), 339-346.
- [27] Thimbleby, H., Anderson, S., and Witten, I. H. Reflexive CSCW: supporting long-term personal work. *Interacting with Computers* 2, 3 (1990), 330-336.
- [28] Trimble, J., Wales, R., and Gosswiler, R. 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 (2003), 18-44.
- [29] Volda, S., Mynatt, E. D., MacIntyre, B., and Corso, G. M. Integrating virtual and physical context to support knowledge workers. *IEEE Pervasive Computing* 1, 3 (2002), 73-79.
- [30] Whittaker, S., and Schwarz, H. Meetings of the board: the impact of scheduling medium on long term group coordination in software development. *CSCW* 8, 3 (Jun. 1999), 175-205.
- [31] Whittaker, S., and Sidner, C. Email overload: exploring personal information management of email. In *Proc. CHI 1996*, ACM Press (1996), 276-283.
- [32] Xiao, Y., Lasome, C., Moss, J., Mackenzie, C. F., and Faraj, S. Cognitive properties of a whiteboard: a case study in a trauma centre. In *Proc. ECSCW 2001*, Kluwer Academics Publishers (2001), 259-278.