

# Supporting Transitions in Work: Informing Groupware Design by Understanding Whiteboard Use

Anthony Tang<sup>1</sup>, Joel Lanir<sup>2</sup>, Saul Greenberg<sup>3</sup>, Sidney Fels<sup>1</sup>

<sup>1</sup>Human Communication Technologies Lab, UBC

<sup>2</sup>Department of Computer Science, UBC  
Vancouver, BC CANADA V6T 1Z4

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

<sup>3</sup>Dept of Computer Science, University of Calgary

Calgary, AB, CANADA T2N 1N4

saul@cpsc.ucalgary.ca

## ABSTRACT

Many groupware tools focus on supporting collaborative real-time work; yet in practice, work spans many different modes: from collaborative to independent activity, and from synchronous, real-time activity to asynchronous activity. How can we design tools that allow users to transition between these modes of activity smoothly in their work? We consider how the common office and domestic whiteboard are used for both independent and asynchronous activity, showing how users employ the whiteboard to transition between these and other modes of activity. Our findings suggest that the whiteboard does so by being a contextually located display with visually persistent content, facilitating transitions because it is a flexible, common tool enabling the creation of representations that are useful across modes. We explore the design implications of these findings with respect to interactive whiteboard tools, and discuss how they can be applied more generally to inform the design of groupware tools.

## Author Keywords

Whiteboard, groupware, reflexive cscw.

## ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

## INTRODUCTION

Groupware research, almost by definition, focuses on providing awareness, communication and shared workspace support for collaborative activity. This research typically focuses on real-time activity (e.g. instant messaging [16], shared visual workspaces [5], whiteboard activity [20]), or asynchronous awareness and communication (e.g. email [29], bulletin board displays [7]). In practice, collaborators

	personal	collaborative
synchronous	<b>Worker</b> <ul style="list-style-type: none"> <li>• word processor</li> <li>• spreadsheet</li> <li>• CAD software</li> </ul>	<b>Real-time interaction</b> <ul style="list-style-type: none"> <li>• telephone</li> <li>• video conference</li> <li>• instant messaging</li> </ul>
asynchronous	<b>Personal management</b> <ul style="list-style-type: none"> <li>• PIM: scheduler, agenda, task list</li> <li>• reminders, post-it notes</li> </ul>	<b>Ongoing tasks</b> <ul style="list-style-type: none"> <li>• team rooms</li> <li>• bulletin boards</li> <li>• email</li> </ul>

Figure 1. A modified groupware matrix.

move between asynchronous and synchronous modes of activity—even within the context of the same activity. For instance, immediately before Anne emails a paper to Bob for review (asynchronous), they might engage in a telephone conversation about issues in the paper (synchronous). Notice that email, the word processing tool, or telephone is individually sufficient for meeting Anne and Bob’s needs, forcing them to transition between tools. With some notable exceptions (e.g. [10, 11, 4]), groupware designs have large focused on supporting one mode of activity: *either* synchronous *or* asynchronous.

Yet beyond simply asynchronous and synchronous group activity, many authors observed that groupware tools are also often appropriated and used *reflexively* for *independent activity* (e.g. [29, 26]). As an example of this reflexive-groupware use, Anne may send *herself* an email to remind herself later to edit a figure in the paper she is writing with Bob. While not all groupware tools can be used reflexively, there are some clear benefits to being able to do so [3]: users rely on existing or known practice to smoothly *transition* between independent and collaborative activity, rather than break their flow and use a different tool. For example, after her phone conversation with Bob, Anne might also email Bob to remind him to edit the conclusion of their paper. Notice that she can use the same mechanism (send an email) to remind herself (independent activity) as she does to remind Bob (collaborative activity).

Figure 1 illustrates how we can conceptualize groupware use in a matrix that contrasts independent vs. collaborative work and synchronous vs. asynchronous work. Typically,

*Cite as: Tang, A., Lanir, J., Greenberg, S. and Fels, S. (2008). Supporting Transitions in Work: Informing Groupware Design by Understanding Whiteboard Use. Technical Report TR-2008-04, Department of Computer Science, University of British Columbia, Vancouver, British Columbia, CANADA, V6P 1Z4, April.*

groupware designs concentrate on facilitating activity in one quadrant, yet we have seen that real work often spans several of these quadrants—be it between asynchronous and synchronous activity, independent and collaborative activity, or both. Creating tools that can transition these boundaries smoothly can therefore better support the real-world demands of work (e.g. [3, 4]). Yet, how can we derive requirements for such tools?

One place to begin exploring these needs is to study real-world tools that already provide seamless transitions across these boundaries. The ubiquitous whiteboard, found in both the workplace and the home, is one such tool. While we often consider whiteboards as tools for synchronous collaborative activity, there is evidence that suggests personal whiteboards are also used for asynchronous activities [14], such as placing reminders or storing information for later use, and for personal brainstorming activities (e.g. [2]). We have also seen asynchronous communication in “war rooms”, where information is left on the whiteboard for later consumption (e.g. [25]). What then, is the space of activity supported by the whiteboard artefact (or large displays that mimic it)? How does the whiteboard support asynchronous activity, and what enables it to be used across these different modes of activity? Ultimately, what can we learn from how the whiteboard is used that can be applied to groupware design to similarly enable transitions between modes of activity?

This work presents the findings from a survey of 135 regular whiteboard users, and an analysis of 10 in-situ interviews with self-identified “heavy users.” The findings show that users commonly employ whiteboards to communicate with their future selves and others by leaving reminders, storing information or activity states. Whiteboards perform these functions because they are located containers for reliably persistent visual content, facilitating easy perusal and update of information, and because their flexibility allows the use of representations common across modes of activity. These affordances provide groupware designers with an approach to support transitions, both for interactive whiteboard applications and for groupware applications in general.

## RELATED WORK

We begin by reviewing reflexive-CSCW, a concept that provides some motivation for this work, since it suggests the application of groupware tools for independent work as well as collaborative activity. This leads us into a discussion of “seamless” designs that we have seen in groupware. These seamless designs provide different mechanisms to transition between modes of work. Finally, we situate our whiteboard study by exploring work that studies interactive and traditional whiteboards.

### Reflexive-CSCW

Thimbleby et al. [26] first defined *reflexive-CSCW* as the application of CSCW designs, theories and systems for

personal use. Thimbleby et al. argue that individuals often play several roles in their own work, and as a consequence, need to communicate with themselves. Cockburn & Thimbleby [3] elaborate on this conception, identifying potentially two different roles an individual might take during independent activity: the *worker role*, which involves actually executing actions required to complete work tasks, and the *personal management role*, which coordinates the activities of the worker, setting task lists, creating reminders for the worker, and deciding on what to do next. While communication between the roles often takes place internally (i.e. as thought), we often see remnants of this communication in tools or environment as task lists, reminders on post-it notes, and so forth [9]. In effect, an individual often has need to communicate with one’s own *future self*—much in the same way one might communicate with another party altogether [4]. Figure 1 illustrates these ideas in relation to groupware.

This concept suggests that a fruitful way to support asynchronous personal activity is to reapply CSCW tools *reflexively*. As the example in our introduction illustrates, email such a tool that supports smooth transitions between asynchronous collaborative and independent work: users can apply the same mechanism to communicate with both others and themselves.

### Seamless Design in Groupware

Our interest in designing seamless transitions between different styles of work is heavily motivated by earlier groupware researchers. Ishii and colleagues, in their work on the TeamWorkStation [10] and ClearBoard [11], explored seamless workspaces for remote collaborators. TeamWorkstation realizes a seamless connection between personal and collaborative workspaces by overlaying a video image of a remote collaborator’s drawing workspace on one’s own, thereby fusing the two workspaces without changing users’ existing work practice (i.e. the way they might use the workspace if working independently) [10]. ClearBoard takes this notion a step further, and fuses the workspace with interpersonal space, overlaying the video of a collaborator and drawing workspace on one’s own [11]. This fusion allows collaborators to work, and maintain and use eye gaze—much as they would if collocated. These systems both realize a form of seamlessness that enables users to transition from independent work and engage in collaborative work while preserving existing work practice.

TeamRooms [4] employs a “room” metaphor for shared visual workspace groupware, allowing users to use and leave artifacts in different virtual “rooms” in the collaborative environment. The rooms are persistent, and names of the rooms are visible across the users’ clients. These rooms allow users to partition work artifacts, to use the artifacts synchronously with others, and more importantly, to leave and return to artifacts at a later time. In so doing, the TeamRooms design explicitly supported transitions between synchronous and asynchronous activity,

as well as between independent and collaborative activity. We aim to further inform the insight of this metaphor and by exploring the real-life use of whiteboards.

### **Traditional and Interactive Whiteboards**

We are interested in studying whiteboards because although we commonly consider them as tools for synchronous collaboration (e.g. [2, 22]), many authors have also reported their use for both independent activity (e.g. [2, 14]), as well as asynchronous activity (e.g. [14, 30]). The whiteboards is thus a tool that has been appropriated for many varied and idiosyncratic uses, and in many cases, these uses cross many styles of activity (e.g. [14, 30]).

Mynatt [14] focused on personal whiteboards in an office context, exploring how space on whiteboards was managed by users to facilitate multiple parallel tasks. Some of these parallel tasks included reminders, quick capture, and thinking. This work revealed how space (when partitioned into segments) helps to organize work, especially when allowed to persist for long term (e.g. for reminders). These insights force us to consider the whiteboard also as a tool for asynchronous activity. Teasley et al. [25] provide convergent evidence from observations of “war room” whiteboard use, reporting that whiteboards provide a space for asynchronous communication, acting as a shared awareness display about a team’s status or current activity. In the present work, we aimed to further extend this understanding of asynchronous activity, in particular by studying the content of some of these longer-term segments, and how they were used to resume activity.

Xiao et al. [30] present a case study of the use of whiteboard in an emergency room ward, illustrating how its location and visual persistence afforded many different styles of use. For instance, the whiteboard captured the current state of the ward, providing awareness to a casual passerby (asynchronous collaboration). The location dictated the nature of the content (i.e. by shaping the set of likely viewers), and facilitated central coordination between nurses regarding the schedule. Further, it provided a space for independent activity (e.g. testing out different schedule configurations), and an easy and meaningful location for updating status of the ward in a visible manner for others.

Interactive whiteboard software has primarily aimed to support real-time collaboration. SMART’s software package [21], for example, provides a broad suite of applications for their interactive electronic whiteboard display. The Notebook application supports simultaneous drawing for two users using multiple colours, and a page-based interface for organization. It provides drawing support in the same way as one would expect from a shared drawing tool, thus making it an effective teaching tool.

The BlueBoard [20] and MERBoard [27] projects were designed to support walk-up-and-use collaboration, allowing users to engage with a large display for collaboration. Both provided users with a suite of

applications, such as drawing, web browsing, scheduling, and dedicated project tools. These systems were used for synchronous collaborative activity, and their ability to load information across the network also facilitated their use as shared output displays for work artefacts prepared elsewhere [8]. Dynamo [1] brings focus on the ability to create work artefacts on the display at the system itself, and to visually notes and “media parcels” so they can function as reminders and later retrieved.

In contrast to supporting ad hoc activity, Tivoli aimed to provide structure to meeting room collaboration [13]. Like Cognoter [24] and other GDSSs [18, 22], Tivoli also supported bringing in external resources such as documents over the network, and further allowed voice annotations to be attached to whiteboard activity. These design ideas were to support later (asynchronous) review of meeting room (collaborative, synchronous) activity.

Kimura [28] was designed to facilitate deferral, resumption and peripheral monitoring of ongoing activities or tasks. It provided ambient display of montages (e.g. thumbnails) representing activity. These montages could be manipulated, tagged, and later reactivated. The design allows users to move between synchronous and asynchronous modes of independent activity, and reaffirms the utility of visually persistent information even if no explicit action is taken to manipulate it.

This brief review shows us some gaps in the design of large display technologies for collocated collaboration. Many systems were designed for synchronous activity, and while many support asynchronous activity, their designs clearly do not provide seamless transitions suggested by studies of traditional whiteboard activity (e.g. [30, 14]), nor by the designs of TeamWorkStation [10] or TeamRooms [4]. What aspects of traditional whiteboards allow them to function this seamlessly? This review suggests two factors: persistency and location. The persistent display of information (e.g. reminders) is important [30, 14], but most technological interventions are display/projector-based, so cannot be expected to be as persistent—thus, we have not seen how such a system would be appropriated. We have also seen that a whiteboard’s location has a strong impact on its potential uses due to its expected users/viewers [1, 7, 14, 30], yet electronic whiteboard systems are largely location agnostic. More fundamentally, this review prompted us to revisit our assumptions about interactive whiteboard design: in addition to collaborative synchronous activity, what asynchronous activities are common? How do whiteboards support these activities? Indeed, *why* might they be used for these activities instead of other media?

### **STUDY METHODOLOGY**

We understood at the outset that whiteboards are a fairly ubiquitous artifact both domestically and in the workplace. Furthermore, they have been co-opted in many different ways. Since we were interested in challenging our assumptions about “synchronous only” uses, we focused on



Figure 2: Larry and Mary’s nursery whiteboard is collaboratively authored.

asynchronous use, and on users who would have evolved their work practice with whiteboards. Thus, we targeted “regular whiteboard users”: those who use whiteboards at least once a week. This criteria would exclude “casual” whiteboard users (who would not have appropriated the whiteboard into their everyday use), and give us a sense for the diversity of complex whiteboard practice.

### Survey

We deployed a web-based survey using a snowball recruitment sampling technique. We received over 167 responses, of which we used the 135 responses that were complete. Our sample included a wide demographic, including a broad spectrum of occupations (engineers, consultants, business analysts, academics, designers), and a range of frequency of use (1-2 uses/week, to 9+ uses/week). The survey itself consisted of 53 items, asking users about their whiteboard behaviour: what activities did they engage in, and how frequently? How frequently were these activities engaged in with others, and so forth. We also asked users about two whiteboards important to them, collecting detailed information about 250 such whiteboards, including where the whiteboards were located, what they were used for, what was currently on them, how long that content had been there, 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 also conducted in-situ interviews with users selected from our survey pool. These users were selected on the basis of being self-identified “heavy” whiteboard users who frequently used whiteboards by themselves. We conducted interviews until we felt we had exhausted the diversity of uses and were essentially hearing the same ideas over and over again. Ultimately, 10 interviews were completed (2 females), and these users came from a broad variety of occupations such as academics, managers, and engineers.

The interviews lasted an hour, and all were conducted in the user’s place of work (save for two which were conducted over the phone with overseas participants with the aid of digital photos of their whiteboards). We collected photographs of users’ whiteboards and their physical context, and used the whiteboard as a grounding artifact for discussion. While we developed a list of questions around theme areas, we allowed the flow of the interview to guide the dialogue, referring to the list only to ensure that all themes had been addressed.

### Analysis Method

We conducted an inductive analysis of interview data, iteratively coding the interview transcripts [23]. The coding process was aimed at deriving a thematic understanding of our participants’ activities—focusing on independent use of whiteboards, asynchronous activities, and how persistency was used on the whiteboard.

### FINDINGS

In this section, we combine the findings from both the survey and interviews to provide an in-depth picture of some whiteboard activities, how persistency supports these activities, and the role of location in facilitating asynchronous communication.

#### Asynchronous Activities

For many users, their whiteboards function as a type of coordination centre, similar to the function of an agenda or computer/PDA-based personal information management tool. They allow users to leave reminders, messages, and tasks lists for their future selves. These notes become relevant to one or more contextual triggers: *temporal* (e.g. deadlines), *location* (e.g. shopping list for the store), *activity* (e.g. a prior brainstorm), or *person* (e.g. messages left for someone else). As alluded to earlier, what makes the whiteboard interesting is that the same whiteboard and content can and often is used for multiple styles of work, sometimes crossing the independent *vs.* collaborative threshold, sometimes the synchronous *vs.* asynchronous threshold, and sometimes both. The following three vignettes drawn from our interview data provide examples of how the whiteboard bridged gaps across several styles of work, illustrating decision making, deferral and resumption of activity, and asynchronous reference.

#### *Asynchronous, located, at-a-glance decision making*

Larry and his wife Mary have just had a newborn child, and converted the old home office into the nursery. One of the last things Larry intended to do was to move the whiteboard out of the nursery (Figure 2); however, now, six months later, Larry reports that the whiteboard is an indispensable part of how they take care of their newborn.

The first hectic days of bringing their newborn home were routine-less—Larry and Mary had a seemingly endless set of tasks that needed to be done, feeding the newborn, taking medication, cleaning various body parts, purchasing last minute or forgotten items—all of which were new and

unusual. On a whim, they both began authoring this information with timestamps on the whiteboard, allowing Larry and his wife to track what had been done to the baby and when. The whiteboard served thus served three functions: as a *memory aid*—communicating with oneself (“When was the last time the baby was fed?”, “Have I already taken my baby-related medication?”), as a *communication device* (“I did this to/with/for the baby when you were not around.”), and as a *task list* (“These need to be purchased for the baby.”). The chaotic few days of the newborn’s home life led to the unplanned adoption of the nursery whiteboard as a coordination centre.

Now, six months later, Larry says that the content on whiteboard has evolved and normalized because while the baby is a large undertaking with occasional unexpected tasks, their daily life has also stabilized somewhat. The whiteboard now primarily tracks the baby’s feedings, the relative duration, and which breast was used for the feeding. Although this information is only briefly salient in the moment, it is important to Larry and Mary for making decisions. At each feeding, Larry or Mary can quickly refer to the whiteboard to determine which breast the newborn should feed from (since this needs to alternate). This visible record is crucial since Mary can easily forget, and at other times, Mary’s hands are full with the newborn, so she can ask Larry to look. At night time, when the newborn awakens, the board also plays a decisive role in determining who must wake up. If, for example, the last feeding was over four hours ago, then Mary must wake up to feed the baby; otherwise, it is Larry’s responsibility to put the baby back to sleep. Of course, it is still always Larry’s job to go and check the whiteboard, though.

The whiteboard’s information, as a jointly constructed artefact, functions as a coordinating mechanism for Larry and Mary’s childcare activities, both independent and collaborative. They asynchronously employ this information as a shared memory device (deciding which breast to use), and for decision making (determining whether the newborn’s feeding patterns have regressed). In this case, we see the whiteboard being used for both asynchronous independent and collaborative activity.

#### *Deferral, storage and resumption of personal activity*

John is a researcher for a small telecommunications start-up, and is responsible for delivering architectural designs that link together hardware and software components with customers’ systems. Creating these designs is an iterative process for John, requiring him to continually re-evaluate his designs, communicating with customers about their needs and interfaces, and with his developers about the details of their own components. John uses his office whiteboard (visible from his desk) to sketch through these ideas, capturing and storing his design ideas. This process often reveals uncertainties about different components or requires making assumptions about them—all of which are

captured in the sketches. John then uses these sketches as a task list for issues that need to be clarified.

At the time John was interviewed, 30% of the space contained a small task list, remnants of brief whiteboard meetings with others, and a joke; the remaining 70% of the whiteboard was filled with three design sketches. These design sketches, reports John, are captured “brain states”—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). Each sketch is like informal documentation about decisions he has made or his ideas about each problem. The sketches structure his ongoing work, focusing his energies when he comes into the office: persistently remind him of unresolved issues or uncertainties in designs, 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 continually updates the design sketches so that the brain states are always up-to-date. John also sometimes uses the sketches as a means of communicating with his team members.

John uses the whiteboard as a *storage device*, and its constant visual availability means that ideas can “percolate” when he needs to work on another task. More importantly, John says that he is often interrupted in his work, so his use of the whiteboard means that when he returns, he can use sketches to immediately bring himself up to date: What was he thinking about when he was interrupted? What issues need to be dealt with right now? The brain states allow John to temporarily suspend his work and thought, and allow him to later pick up and resume activity seamlessly. Often, as mentioned earlier, these activities entail gathering information from others about questions or issues, and the sketch therefore functions as a sort of task list for John.

As time goes on, and the ideas become more stable, John says that the representation on the whiteboard becomes less important because the ideas are already either “ingrained” in his mind, or toward the end of the process, captured by some other form of formal documentation (e.g. a design specification). It is around these times that the sketches are removed and the whiteboard space reclaimed.

John’s whiteboard is primarily used for independent activity, for both synchronous and asynchronous activity. The whiteboard is an up-to-date representation of John’s work, providing him a space for synchronous activity, such as brainstorming and design work. The brain states allow John to smoothly suspend activity on problems, and to later use them asynchronously to recontextualize himself when he resumes activity, or to derive tasks he must engage in. Even without re-engaging with the problems deeply, John employs the visual manifestation of these brain states as an *ambient display* to remind himself of the current, or “present state” of every one of the problems.





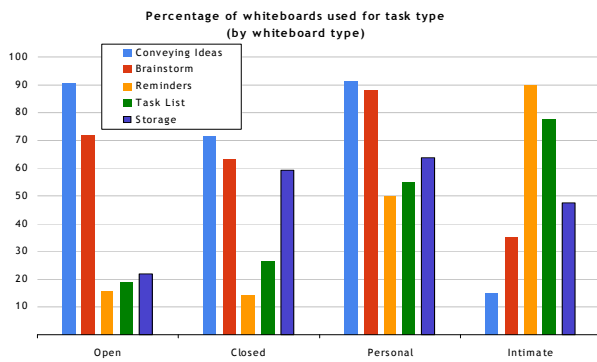


Figure 4. How different whiteboard “types” are used.

We performed an exploratory cluster analysis of the 250 whiteboards reported on by our respondents. As a word of caution, these whiteboards were survey respondents’ “most important” or “second most important” whiteboards. Thus, many whiteboards may not fit this classification, such as the in/out whiteboards commonly found in reception areas. This cluster analysis was based on descriptors such as the whiteboard’s location, who used the whiteboard, what activities the whiteboard was used for, the content currently on the whiteboard and the authors of that content, and so forth. The analysis revealed four different clusters (Figure 4), showing that while the whiteboard *artifact* may be similar in across contexts, they are different tools *depending* on the context.

*Open shared 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 are 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 fell into this category. *Closed shared whiteboards* (27%) also tend to be in shared location (such as in a lab), but the users and viewers of the board are typically known. They are used for similar tasks as public shared whiteboards, but in addition are occasionally be used storage of information or shared knowledge. These activities are made possible because the user pool is known and fairly fixed. Whiteboards in office common areas, “war rooms”, or labs are a good example of this type of whiteboard. *“Personal” whiteboards* (32%) are located primarily in users’ personal workspace (in the office), and are therefore primarily used by the user in question. A small set of trusted, or close co-workers are sometimes invited to use these whiteboards, too. It is on these whiteboards that, in addition to the usual synchronous activities, we begin to see asynchronous activities: reminders, task lists, and long-term storage of information are common on these whiteboards. *Intimate whiteboards* (22%) are similarly often located in personal workspaces, but also at home and in home offices. The users of these boards are almost exclusively the owner, and are *primarily* used for asynchronous activities (e.g. reminders, task lists).




Figure 5. Whiteboards with asynchronous information.

A whiteboard’s *location* is therefore highly associated with whether it is used for asynchronous activities. Intuitively, this makes sense since we should expect that many of these reminders or lists would only be useful if they could be expected to be seen regularly (inadvertently or otherwise). Accordingly, of the 108 whiteboards that had reminders or lists, 70% of were visible from the user’s workspace; of the 98 users who reported putting reminders on their whiteboards, 77% reported seeing them regularly at least once a day. Thus, they are readily visually accessible to users. In contrast, of the whiteboards used for synchronous activities such as conveying ideas (175) or brainstorming (164), only about half were visible from the user’s workspace (47%, 49%).

Asynchronous activity is interesting because it is often one-way (especially in the case of leaving messages for oneself)—unlike a dialogue, where one can ask for clarification, the “receiver” of the message must be able to interpret the message (e.g. know what the reminder or list is referring to). During interviews (and all three of the vignettes presented above), whiteboards contained many instances of short-hand or abbreviations—many of which were incomprehensible to the interviewer, though were readily interpreted by the interviewee. Accordingly, we should expect asynchronous activity to occur primarily on whiteboards whose users know each other well. Figure 4 illustrates this relationship: open shared whiteboards are rarely used for any asynchronous activities (reminders, task lists, storage). This finding accords with the presence of common asynchronous content, such as reminders and lists on these whiteboards. Figure 5 illustrates that this content is more common with personal and intimate whiteboards.

This data suggests that whiteboard location largely determines whether asynchronous activity takes place. A likely reason for this is that some locations afford access to individuals capable of interpreting the content.

### Asynchronous Information Practice

The whiteboard does not operate in isolation to support asynchronous activity: information is often copied to or from the whiteboard onto other media. In some cases, the purpose of copying information off the whiteboard is to retain the information in a more permanent record. For instance, 79% of survey participants reported having taken notes of whiteboard content (be it via paper or computer), 65% having taken photos of that content, and of these, 80%

reported reviewing notes or photos again later at least half the time. Participants also reported transcribing from the whiteboard to more mobile media, such as Post-It notes. These Post-It notes could then be reviewed in another, more contextually appropriate location or time.

Many users (49%) also reported transcribing prepared content onto the whiteboard, too. This means the content was prepared elsewhere, and then recorded here so that it could be seen later, or on an ongoing basis. John of the brain state board is one such user: during business trips, he takes notes on paper since his whiteboard is unavailable. When he returns home, he transcribes the relevant information to his whiteboard.

The whiteboard can also perform a reminding function beyond most personal information management tools, because whiteboard content has the property of being visually persistent, which changes the meaning of being “reminded.” One participant reported, “[*Outlook*] doesn’t do the task thing right because I have to set a date, but lots of stuff doesn’t have dates. I also don’t like the priorities, because it only has three, and on top of that the categories thing doesn’t work for me either... [With the whiteboard] I can just organize it the way I want.”

Finally, we should be clear that persistent content on personal whiteboards is not always left intentionally for later review, as in these examples. Psychologists have demonstrated that often, the *act* of writing facilitates the encoding or retrieval process (i.e. for memory) [16], and that for some users, reviewing the content is secondary. For instance, we saw some task lists that were severely out of date (e.g. old or inaccurate). This suggests that users have difficulty keeping these lists up to date if completing the task elsewhere, and that for some, the *act of constructing* the task list was sufficient for their needs. For these users, after constructing a task list on the whiteboard, it may never even be referred to again. In spite of this, the fact that the task list is persistent on the whiteboard is useful, as one participant stated: “It’s nice to know that it’s there, and I can see if I’ve forgotten to do anything.”

## DISCUSSION

Beyond the common conception of the whiteboard as a tool for shared synchronous use, we have seen that the whiteboard can and does perform roles in many other task contexts. It facilitates transitions between different modes of activity because of several properties: (a) it is a *container* for task and coordinating information [4], (b) information is easily *revisitable*, (c) information is readily *updatable*, and (d) the flexibility of the device *allows users to build representations of information suitable for many modes of activity*. Many of these properties can be reapplied in many groupware systems—particularly interactive whiteboard systems—to facilitate transitions between modes of activity, whether it be between independent and collaborative activity, or synchronous and asynchronous activity. In this section, we review these properties in turn.

Building on Greenberg & Roseman’s articulation of the “room metaphor” in supporting transitions [4], we also see the whiteboard as functioning as a *container*. Information placed on many whiteboards (e.g. closed shared, personal, intimate) can be expected to be *persistent*: without explicit action, the information will not be removed. Within the whiteboard, information is freely organized, and reliable: it is always available without needing to be loaded (as say on a computer). Similarly, the container is *permeable*, and readily provides access to that information. Whiteboards are typically constantly *visually available*—unless information has been deliberately obscured. Further, the information is easily *updatable*—in terms of one’s ability to rapidly add, erase, or change content on the whiteboard. Beyond the virtual metaphor in [4], the whiteboard has physical embodiment, and is *contextually located* near or in a place where action takes place [9]. Thus, the whiteboard *limits access* to people who would likely be in the context [7] (e.g. consider the nursery whiteboard, accessible primarily only to Larry and Mary), *aids interpretation* by being in the same context [14] (e.g. the nursery whiteboard’s information only makes sense in that context). *Thus, the whiteboard is a container for contextually located information that users are able to constantly revisit*. The ability to revisit this information thus allows the information to be used across a variety of task contexts, especially between synchronous and asynchronous work.

The whiteboard, of course, is a highly flexible and malleable surface. Meaning, as has been alluded to by several authors (e.g. [14, 2, 30]), 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 [14]), 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 and Mary’s 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). *With whiteboards, users can therefore evolve representations that are consistent and meaningful across different modes of activity*. These representations can be used across modes with little loss of generality or meaning (e.g. Jill’s whiteboard can be used by herself or by others both asynchronously and synchronously).

The whiteboard is therefore a realization of the “room” concepts from [4], and we have seen in this study that they do, in this instance, support transitions between modes of activity. Furthermore, we have seen the *ways* in which they support those transitions: by enabling revisitation (through persistency and location), and by providing a means to create representations meaningful and interpretable (via flexibility) across activity modes.



## OPPORTUNITIES FOR DESIGN

Our study of the whiteboard tool provides an understanding of how groupware tools can be built to support not only asynchronous independent activity, but also the transitions to synchronous and collaborative activity. Based on how whiteboards facilitate these transitions, we see several clear opportunities for design:

*Facilitating revisitation.* The whiteboard supports revisitation through its fixed physical location and persistent content. Technology-based solutions can go further, and allow messages to be tagged with time, time spans, locations or activities. These tags can then be used as triggers to promote revisitation—even if the intent is not to trigger action on the messages themselves. PlaceMail, for instance, facilitates sending messages to users when they are *near* a location tagged to the message [12]. One might imagine augmenting displays so that these messages are just *present* when the user is in the location, rather than *notifying* the user as PlaceMail does. Most PIM tools typically trigger reminders based on time, but most are triggered based on a “due date” rather than facilitating a time span. Displaying information over a time-span facilitates incidental, unbidden encounters with the information. Beyond this “tag and display” notion, we even consider showing information across unused displays as an ambient form of revisitation a useful notion. This idea has been applied to displaying hand-written notes successfully [6]. We might also imagine that displays might be linked together, facilitating revisitation of content from different displays to further promote this reminder functionality.

*Facilitating updatability.* A simple idea is to facilitate updating or modifying information on fixed displays remotely. While this is not a new idea, it further enables fixed displays to function asynchronously because it removes a potential inhibitor of use (not being physical co-present with the display when it should be updated).

*Facilitating interpretation.* One of the problems with leaving messages is that they are often ambiguous, and the meaning can be forgotten. The whiteboard may overcome this problem in part because the limited whitespace forces users to constantly re-evaluate what can be erased, and this evaluation forces revisiting and reinterpreting content (so the user is constantly “reminded” of what the meaning is). Further, while the contextual location may aid in interpreting such meaning, technology can be more explicit. To help disambiguate these messages, it may be possible to link content, such as documents or other digital information to the messages themselves (e.g. [3]). Similarly, because the final representation of a message may have been modified, obscuring its meaning, it should be possible to explore the “change history” of a message to determine its meaning.

*Enabling flexible representations vs. complex operations on those representations.* More difficult is enabling flexible representations of information, while allowing both flexible

and powerful operations on those representations. One might imagine, for instance, that updating Jill’s schedule board might update her team’s electronic schedulers. The design tension is that enabling this kind of functionality necessarily constrains users’ ability to create meaningful representations, because this flexibility is necessary to allow users to evolve representations as they encounter new needs. Flatland provides an approach where powerful functionality is enabled only to specified segments on an on-demand basis [15].

## FUTURE WORK

This study was conducted to provide an understanding of how to develop interactive whiteboard applications, and groupware. While it has revealed several dimensions that can be realized and extended with technology, we are still actively engaging in several avenues of work. First, we need to understand the space of groupware applications that can meaningfully be extended to support transitions in work. This relates largely to the kind of tasks a groupware application is to support. For instance, the whiteboard facilitates informal, transient, or pre-production activity—content not generally considered as formal documentation [15]. Other groupware systems support different kinds of activity: TeamSpace, for instance, provides meeting facilities [4]; TeamWorkstation and ClearBoard facilitate sketching activity [10, 11]. Still others facilitate transitions between formal and informal representations of work [13]. Second, we intend to investigate these implications in actual groupware, and to see how users appropriate them over the long-term (e.g. [8])—in particular, to understand how the affordances actually engender transitions in work.

## CONCLUSIONS

In real life, work takes place on artifacts across the personal and collaborative boundary, as well as across the synchronous and asynchronous boundary. To build an understanding of how to design groupware that can provide seamless transitions across these boundaries, we studied how whiteboards are used, because this tool enables such transitions. This study shows that the whiteboard’s fixed contextual location, combined with the content’s reliable visual persistence promotes revisitation. The flexibility of the whiteboard allows evolved representations that can be used across many task contexts. The design opportunities arising from this study are directly applicable to interactive whiteboard applications, and also provide avenues for groupware design in general to support transitions. Microsoft Word, for instance, is a tool that has readily incorporated several collaborative features (e.g. revision system) that enables it to be used collaboratively while maintaining the single-user experience. Such integration is valuable to the user: s/he no longer needs to switch tools between different modes of (independent vs. collaborative) activity. We believe that in general, groupware designers should similarly aim to support smooth transitions between different modes of activity.

## 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*, ACM Press (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*, ACM Press (2007), 557-566.
3. Cockburn, A. J. and Thimbleby, H. A reflexive perspective of CSCW. *SIGCHI Bull* 23, 3 (1991), 63-68.
4. 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.
5. Gutwin, C., and Greenberg, S. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work* 11, 3 (2002), 411-446.
6. Hsieh, G., Wood, K., and Sellen, A. Peripheral display of digital handwritten notes. In *Proc. CHI 2006*, ACM Press (2006), 285-288.
7. Huang, E. M. and Mynatt, E. D. Semi-public displays for small, co-located groups. In *Proc. CHI 2003*, ACM Press (2003), 49-56.
8. 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.
9. Hutchins, E. *Cognition in the Wild*. MIT Press (1996).
10. Ishii, H. TeamWorkStation: towards a seamless shared workspace. In *Proc. CSCW 1990*, ACM Press (1990), 13-26.
11. 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.
12. Ludford, P. J., Frankowsky, D., Reily, K., Wilms, K., and Terveen, L. Because I carry my cell phone anyway: functional location-based reminder applications. In *Proc. CHI 2006*, ACM Press (2006), 889-898.
13. 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*, ACM Press (1996), 150-159.
14. Mynatt, E. D. The writing on the wall. In *Proc. INTERACT 1999*, IOS Press (1999), 196-204.
15. 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.
16. Naka, M., and Naoi, H. Effect of repeated writing on memory. *Memory and Cognition* 23, 2, (1995), 201-212.
17. Nardi, B. A., Whittaker, S., and Bradner, E. Interaction and outeraction: instant messaging in action. In *Proc. CSCW 2000*, ACM Press (2000), 79-88.
18. Nunamaker, J. F., Dennis, A. R., Valacich, J. S., Vogel, D., and George, J. F. Electronic meeting systems to support group work. *Comm of ACM* 34, 7 (1991), 40-61.
19. 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.
20. 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.
21. Smart Technologies, <http://www.smarttech.com/>.
22. 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.
23. Strauss, A. L., and Corbin, J. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications (1999).
24. Tatar, D. G., Foster, G., and Bobrow, D. G. Design for conversation: lessons from cognoter. *Computer-supported cooperative work and groupware*, S. Greenberg (Ed), Academic Press Ltd. (1991), 55-80.
25. 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.
26. Thimbleby, H., Anderson, S., and Witten, I. H. Reflexive CSCW: supporting long-term personal work. *Interacting with Computers* 2, 3 (1990), 330-336.
27. Trimble, J., Wales, R., and Gossweiler, R. NASA's MERBoard: an interactive collaborative workspace platform. *Public and situated displays: social and interactional aspects of shared display technologies*, K. O'Hara, M. Perry, E. Churchill, D. Russell (Eds.), Springer (2003), 18-44.
28. 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.
29. Whittaker, S., and Sidner, C. Email overload: exploring personal information management of email. In *Proc. CHI 1996*, ACM Press (1996), 276-283.
30. 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.