TeamTag: Exploring Centralized versus Replicated Controls for Co-located Tabletop Groupware

Meredith Ringel Morris, Andreas Paepcke, Terry Winograd, Jeannie Stamberger
Stanford University
Stanford, CA, USA
{merrie, paepcke, winograd}@cs.stanford.edu, jeans@stanford.edu

ABSTRACT
We explore how the placement of control widgets (such as menus) affects collaboration and usability for co-located tabletop groupware applications. We evaluated two design alternatives: a centralized set of controls shared by all users, and separate per-user controls replicated around the borders of the shared tabletop. We conducted this evaluation in the context of TeamTag, a system for collective annotation of digital photos. Our comparison of the two design alternatives found that users preferred replicated over shared controls. We discuss the cause of this preference, and also present data on the impact of these interface design variants on collaboration, as well as the role that orientation, co-touching, and the use of different regions of the table played in shaping users' behavior and preferences.

Author Keywords
Tabletop interfaces, computer-supported cooperative work, co-located groupware, single display groupware.

ACM Classification Keywords
H5.3. Information interfaces and presentation (e.g., HCI): Group and Organization Interfaces – computer-supported cooperative work.

INTRODUCTION
Single display groupware (SDG) systems [25], such as interactive tabletops, support group work by allowing multiple people to work together with a shared context, thus facilitating communication and productivity. However, designing single display groupware involves many challenges. For instance, there is potential for clutter due to representing information of interest to multiple participants, such as more than one cursor or multiple copies of control widgets.

CHI 2006, April 22–28, 2006, Montréal, Québec, Canada.

Interactive tables are an increasingly popular form of single display groupware that support face-to-face social interaction. There are toolkits available to simplify development of tabletop CSCW applications, such as DiamondSpin [23] and the DiamondTouch Toolkit [6]. These toolkits enable the construction of many interface styles, but provide no guidance as to which design choices are preferable for a particular application or audience.

In this paper we explore an issue that is relevant to designers of tabletop groupware – deciding how many copies of basic interaction widgets to create, and how to position them on the shared display. We compare two endpoints on the spectrum of control placement possibilities: we provided groups with either a single, shared set of control widgets in the center of the tabletop or displayed a separate set of controls in front of each user (still on the shared tabletop display). These controls were menu-like widgets that allowed users to select labels for digital photos. We evaluated the differences between the centralized-controls and replicated-controls designs for TeamTag, a system for collaborative photo annotation.

THE TEAMTAG SYSTEM

Motivation
The increasing popularity of digital photography, which allows users to capture very large numbers of images, has increased the need for photo-labeling applications. These applications, which include commercial systems, such as Adobe Photoshop Album [1], and research systems, such as PhotoFinder [24], allow users to associate custom metadata with their digital photos. This metadata is useful for enabling search of photo collections. TeamTag (Figure 1) is a tabletop photo-labeling application that allows groups of up to four users to collaboratively associate custom metadata with digital images.

Current photo-labeling software is designed for a single user at a traditional PC. However, this is a task that can benefit from a collaborative interface, both for entertainment and efficiency purposes. Labeling a set of vacation photos together as a family could be an enjoyable activity that promotes reminiscence of a shared experience; the group effort could also speed up labeling and result in a more complete set of labels (e.g., Dad forgets the name of a

Figure 1. Four users sit around a DiamondTouch table to label photos using TeamTag.

landmark shown in one photo, but Mom remembers it). Collaborative photo-labeling is also useful beyond the realm of personal collections – it can be an important part of productivity and educational activities. For example, field biologists at our university find it useful to help each other label photos taken on research expeditions. Collaboration allows each biologist to contribute her specific expertise in identifying the species and equipment depicted in the photos. In our exploration of design variants of TeamTag, we used a photo-labeling task inspired by a productivity/educational scenario, rather than personal photo collections, so that we could ask a number of users to interact with the same content.

In the process of developing TeamTag, we faced a design dilemma: should we create a single copy of the labeling widgets for all four users to share, or should we use additional screen space (which is a precious resource in an SDG system) to provide each group member with his own widget set? This question is relevant to a variety of tabletop groupware applications and is a step toward exploring the broader question of how to balance the needs of individuals with the needs of the group as a whole in the design of single display groupware. The issue of balancing individual-oriented versus group-oriented interface needs has been considered in the context of distributed CSCW systems [8], but the design suggestions for distributed systems aren't directly applicable to co-located SDG. This particular design issue is especially relevant for tabletop interfaces, since the underlying hardware (such as the DiamondTouch [7]) allows truly simultaneous interaction by several users, and since direct-touch interaction makes reachability an issue, thus making multiple copies of interaction widgets potentially more appealing despite the additional screen space they occupy.

The remainder of this paper explores this issue in depth. This exploration is situated in the context of the TeamTag application, but collaborative photo-labeling applications themselves are not the focus of this paper.

Infrastructure

The TeamTag application is designed for use on a DiamondTouch table [7]. The DiamondTouch is a touch-sensitive input device, which is combined with a ceiling-mounted projector so that a display is co-located with the input. Our DiamondTouch table measures 85.6 cm by 64.2 cm, and is top-projected with an SXGA projector (resolution 1280 x 1024 pixels), and controlled by a 3.2 GHz PC. Four users can simultaneously interact with the device. Users sit on pads that are electrically coupled to the table, so the table is able to associate each touch with a particular user. As a result, all four users can interact with TeamTag at the same time, and the table can distinguish the user identity for each touch.

TeamTag is written in Java, using the DiamondSpin tabletop interface toolkit [23]. We chose to use DiamondSpin since it provides primitives useful for tabletop interfaces, such as the ability to display traditional Java widgets at arbitrary orientations. Input to the system includes a set of digital photos and a text file containing categories of metadata that should be associated with the photo collection. Metadata assignments created with TeamTag can be exported automatically to a spreadsheet or other generic format, so that they can be available for search [15] or as input to other applications.

Two Candidate Interface Designs

We have explored two alternative designs for TeamTag. Our two interface designs – *centralized* and *replicated* control placement – are related to the choice between designing groupware interfaces that lend themselves to either more closely-coupled or more loosely-coupled group work. Prior work, such as that by Gutwin and Greenberg [8], has explored this tradeoff as it applies to distributed groupware applications, but this design tradeoff has not been explored in the literature on single display groupware for co-located CSCW.

The "centralized-controls" design (Figure 2) places the metadata in the center of the table. Each metadata category is materialized as one circle in the central region. Each circular widget is subdivided into sectors, with each sector corresponding to one possible value of the respective category. For example, one of the circular controls represents the category of "habitat," with sectors corresponding to "desert," "ocean," "grasslands," "forest," etc. The text label of each sector faces the circle's outer edge. We chose to orient the text in this manner in order not to bias the interface toward favoring any particular side of the table. Users in this design freely distribute photographs around the perimeter of the table by dragging them with their fingers. The orientation of a dragged photo changes dynamically to face the user at the closest table edge. A user creates an association between a photo and metadata by tapping her finger on a photo, thereby selecting it. Any circle sector that this user subsequently touches triggers an association of that metadata with her selected photograph. A user can therefore associate several metadata values with

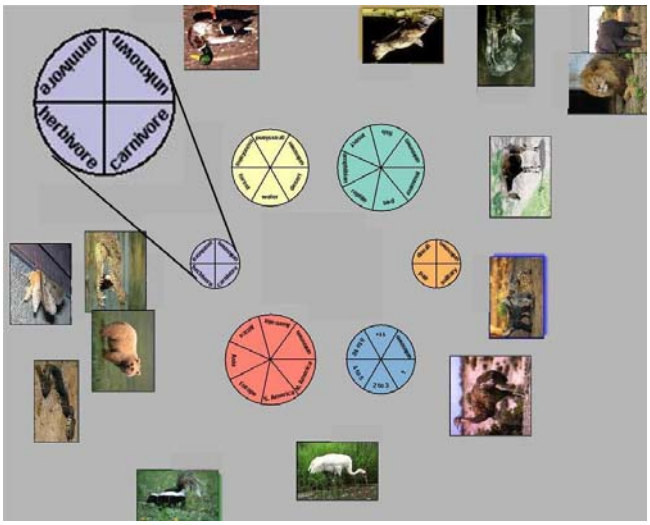

Figure 2. This screenshot shows the "centralized controls" version of the TeamTag interface. Each circle represents a category of metadata, with each sector displaying a label from that category. All four users share this set of annotation controls. The control displaying four labels from the "diet" category has been enlarged in this figure for readability. Photos can be freely moved around the tabletop.

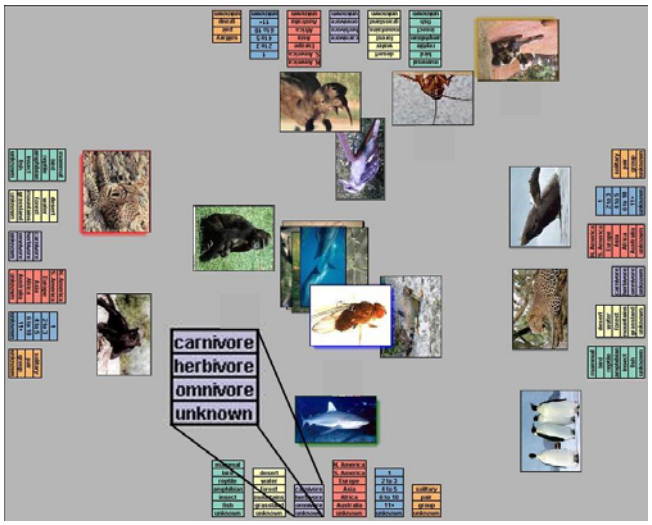

Figure 3. This screenshot shows the "replicated controls" variation of the TeamTag interface. Each rectangle represents a category of labels, with each subdivision displaying a label from that category. The rectangles for each category are replicated for each user. One copy of the control displaying four labels from the "diet" category has been enlarged in this figure for readability. Photos can be freely moved around the tabletop.

the selected photo by tapping rapidly on sectors of each of the circular widgets. For example, a user touches a photo of baboons that is on the table. The photo highlights to indicate it is selected. Then the user sequentially touches multiple attributes that are to be associated with the chosen image. For example, she might touch "omnivore," "Africa," and "grasslands" to indicate the diet, location, and environment of the baboons. Other users can select other photographs and construct associations at the same time; the circle sectors are large enough to fit several fingers at once, and the user-identification features of the DiamondTouch table are used to insure that simultaneous actions are resolved in the appropriate way (e.g., the labels Mary touches are applied to her currently-selected photo and the labels Jack touches are applied to his currently-selected photo).

In contrast to the photos, the circular controls are stationary. A user can rotate a circle with his finger (in the same manner as one would use a rotary phone dial), which turns the circle about its center point, thus allowing a user to view different parts of the text right-side-up if he desires. The control rotates back to its original position when the user removes his hand.

The "replicated-controls" design (Figure 3) inverts the location of photos and metadata. Photos are located in the central area of the table, while metadata categories and values are arranged in a series of rectangular controls around the table's edge. Each rectangle stack corresponds to one of the circular widgets from the centralized-controls design, with each constituent element of a stack corresponding to a possible value for the category. For each category, four copies of the rectangle are displayed, one on each side of the table. The text in each rectangle is oriented toward the side of the table that it faces. As in the first design, this "replicated-controls" version allows photos to be dragged, but the metadata widgets remain stationary.

With either design, a user may inspect an image's metadata by touching the photo with two fingers, which opens a popup window rooted to that image. The popup lists the current metadata categories and values associated with that photo. The text of a popup is oriented to face the nearest table edge. Popups may be dragged around the table for easier viewing or sharing with other users; they automatically re-orient their text to align to the nearest table edge.

Design Considerations

When designing the TeamTag software, the choice between centralized or replicated arrangement of the labeling controls was not clear – there were several issues we hoped to clarify through user testing:

Which design best facilitates collaboration? A potential strength of the centralized-controls scheme is a shared focus of attention on the controls, which could result in increased collaboration. This scheme may also promote incidental learning [12], due to the increased visibility of other group members' labeling choices, and may reduce errors by increasing the likelihood of noticing when a teammate mislabels an item. Conversely, the replicated-controls design emphasizes a shared focus on the photos (by locating them centrally), which is also desirable.

Will the oddly-angled text on the centralized controls reduce usability? Unlike in the replicated-controls design,

the display text on the shared, centralized controls is not right-side-up for all users.

Which design uses screen real-estate more effectively? Although the centralized controls take up a smaller proportion of screen real-estate than is required to copy each control four times in the replicated-controls design, the center-of-the-table space that they occupy might be more valuable.

Will users prefer one design over the other? Although each of the designs offers the same functionality, does the variation in the placement of control widgets impact usability and user preferences?

Will the placement of the controls impact comfort? The centralized design requires users to reach farther in order to tag photos, which could impact task speed as well as be a source of discomfort. However, in the replicated variation, users must place the photos farther from themselves, creating similar ergonomic difficulties.

As a result of the variation in placement of the shared and replicated controls, we needed to alter related aspects of these widgets, such as the orientation of the text they displayed and the widget shape. In the replicated-controls condition, it is logical that the text on the widgets be oriented toward the person for whom those controls are intended. However, in the centralized-controls condition, having all of the text oriented in a single direction would have biased the controls to be more useful toward one of the four group members. Having the orientation of the text change automatically to face the currently-interacting user would also not be a viable design, since the DiamondTouch permits all four group members to simultaneously interact with the same widgets. Thus, we chose to orient the text on the centralized controls outward (and thus to design the widgets with a circular shape), to avoid biasing them toward one particular group member. We chose not to use this circular shape in the design of the replicated controls, however, since one drawback of replication is that it uses additional screen real-estate. Therefore, in order to avoid unnecessarily handicapping the replicated-controls design, we chose a rectangular shape that minimized the space occupied by each widget. Previous comparisons between pie-shaped and rectangular menus [4] have found that users are evenly split in their preference for the two widget shapes. Because this prior work implies that it is unlikely that preference would be biased toward either the circular or rectangular shape, we felt the use of different-shaped widgets was justified to avoid the potentially larger confounds of biasing the interface toward a single member of the group by using rectangles for the shared widgets, or of handicapping the replicated widgets by using less space-efficient circles.

EVALUATION

Twenty-four paid subjects (sixteen male, eight female) participated in our evaluation of the centralized and replicated-controls design alternatives of TeamTag. Subjects ranged in age from seventeen to forty-five years old. Subjects were divided into six groups, with four users in each group. The study had a within-subjects design – each group used both interfaces (centralized controls and replicated controls), in a balanced order.

Additionally, subjects experienced a third experimental condition where they used a tablet PC version of TeamTag. This third interface did not use an interactive table; rather, each group member had their own tablet PC. Each tablet contained a copy of all of the labeling controls, and distribution of the photos amongst the four users' tablets was coordinated over a wireless network. This tablet PC interface differed too greatly from our tabletop interface designs to allow for meaningful comparisons. As all conditions were counter-balanced, elision of the third condition should not have an impact on the observed differences between the centralized and replicated conditions. Thus, this paper compares only the two tabletop interfaces.

TeamTag was instrumented to record all interactions (e.g., who touched where at what time). Additionally, users completed a questionnaire after the study that contained both free-form and Likert-scale questions. We also collected observational data, both from live observation of the experiments and from analysis of video recordings.

To have a photo-labeling task that we could hold constant across several groups of users, we chose a set of images and metadata that could be labeled based on everyday knowledge, rather than using personal photo collections that required user-specific information. The images were of various common and exotic animals, such as one might see on a trip to the zoo, and the metadata that needed to be applied were items such as the type of animal (mammal, reptile, bird, insect, etc.), the animal's diet (omnivore, herbivore, carnivore), the terrain where the animal was found (grasslands, forests, desert, etc.), and other categories of comparable difficulty. The content was designed to be within the normal trivia and reasoning ability of most people, but challenging enough that assistance from other group members would be helpful.

Subjects were told that their group's goal was to label each photo with the appropriate values for each of the metadata categories presented. Before each condition, groups had a tutorial in which the features of the interface were demonstrated and they were allowed to practice and ask questions about the interface. The tutorial used a different set of images and categories than those used in the study conditions.

For each study condition, groups were given twenty images to label. They were given a different set of images for each condition. Subjects labeled two sets of images with each interface – for one set they had six different categories of metadata to add, and for the other set they had only three metadata categories to add to each image. The order of

image sets, number of categories, and condition was counter-balanced among groups using a Latin Square design.

RESULTS

Preferred Interface

On the post-study questionnaire, subjects' responses to "which session did you prefer" indicated overwhelming preference for the replicated-controls interface. Nineteen of the subjects (79.2%) listed the replicated-controls as their favorite interface (Figure 4). Four subjects preferred the shared, centralized controls.

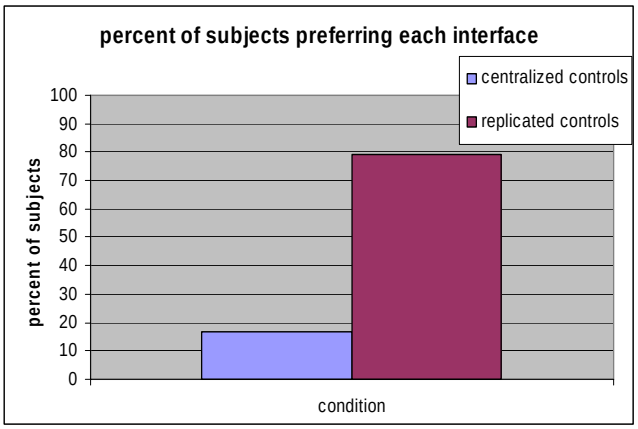

Figure 4. The majority of participants preferred the replicated-controls interface.

Collaboration

Quantifying collaboration when evaluating CSCW systems is challenging [16]. To understand how the centralized and replicated controls designs impacted collaboration with TeamTag, we looked at several factors: how many group members contributed labels to each image, amount of conversation, users' self-evaluations of communication and teamwork, and the quality of the labels assigned.

One indicator of collaboration is how many users added a label to each image (each image needed either 3 or 6 labels, depending on the condition). For example, an extremely parallel work style with no checking of each other's work afterwards might have only one person do all of the labeling for each photo, while a group with a more cooperative strategy or where users checked each others' work might have all four users contribute labels to each image. For the table with replicated controls an average of 1.24 users contributed labels to each image. This number was significantly higher ($t(5)=3.55$, $p<.02$) for the centralized-controls interface, with an average of 1.56 users contributing metadata to each photo.

Automated analysis of the soundtrack from the video recordings of each session found that groups spent 44.7% of each session conversing with each other when using the shared controls, and 42.4% of the time talking to each other when using the replicated controls. There was no statistically significant difference between these two levels of conversational activity.

A free-form question asked subjects to describe the level and quality of their communication with their group in each condition. Subjects distinguished little between the two conditions as to which facilitated more communication. This feeling was reflected in subjects' 7-point Likert scale (7 = strongly agree, 1 = strongly disagree) responses to the statement "I felt that my group worked closely as a team" for each condition. The means for centralized-controls and replicated-controls (5.63 and 6.00, respectively) were not significantly different from each other.

We initially hypothesized that subjects might be uncomfortable with the shared controls because the central location of the controls might subject their labeling behavior to more scrutiny from the team – e.g., they might be more embarrassed if they mislabeled something since it would be more likely that someone would notice the labeling action if it occurred in the center of the table. The questionnaire asked users if they felt hesitant guessing when they were unsure of the answers in each condition. Participants indicated they felt more comfortable guessing with more distributed controls, as reflected by giving a mean score of 4.38 for the centralized controls and 4.17 for the replicated controls, indicating a greater hesitance to guess with the centralized design. However, although these scores trended in the direction we expected, the differences were not statistically significant.

Task outcome is another indicator of collaborative success. Increased collaboration could help groups label more of the images correctly. However, this may not have been true for our task, since subjects had not studied animal facts beforehand and so random variation in groups' a priori knowledge of the chosen subject matter likely overshadows any impact of the interface on accuracy of the labeling. Performance in each condition was similar, with mean scores of 65.6% of metadata assigned correctly with the centralized controls and 66.9% in the replicated-controls condition.

Co-Touching

Even though the DiamondTouch hardware and TeamTag software support concurrent touches by multiple users, and the shared circular widgets were large enough to fit several fingers on a single sector, subjects were hesitant to simultaneously touch the shared circular controls. For instance, we observed one user reaching toward the "diet" category control, but when another user touched it, he immediately withdrew his hand, hovering nearby and waiting for her to complete her action before he began. The hesitance of users to "collide" with each other when using the centralized controls was evidenced by the lack of co-touching on these shared widgets. Across all six groups, there were a total of only thirteen co-touching events (where a co-touch is defined as more than one user touching the same control within .5 seconds of another user). Thus,

co-touching represented only 0.9% of the 1,372 total touches across all groups of the centralized controls.

Subjects had space on the post-study questionnaire for unstructured responses to questions, such as explaining why they liked or disliked a particular interface. One recurring theme in these free-form responses was participants' discomfort with the possibility of accidentally bumping or touching other users' hands when using the shared set of centralized controls. Ten of the twenty-four questionnaires (41.7%) mentioned this in the free-form comments, even though there was no question specifically about this topic. For example, one subject said that he preferred the replicated-controls over the centralized-controls because "our hands didn't go on top of each other."

Although groups mainly used a parallel strategy where each group member would simultaneously label photographs, one group used an assembly-line strategy with the centralized controls, as a means of avoiding the problem of having to collide hands with other users. They took each image and passed it clockwise around the table, and each user was then responsible for adding the metadata from the controls nearest him. A second group discussed using an assembly-line strategy for the shared controls, but ultimately decided on the parallel strategy because some group members felt that they were not knowledgeable enough about the metadata categories that happened to be located nearest them.

Orientation of Information

Despite the fact that users often commented that the ability to rotate the circular widgets with a "rotary phone" interaction in order to view text right-side-up was "cool" or "neat" when it was demonstrated to them in the tutorial, they rarely made use of this capability – each user performed an average of 1.9 rotations during the study.

The unimportance of viewing the labels right-side-up was further supported by the questionnaire responses. The statement "I found it useful to be able to rotate the circles" received a mean score of 3.0, which falls on the "disagreement" end of the 7-point Likert scale. We found this surprising, since we had anticipated that subjects would not want to read text upside-down and might use the rotation to alleviate this difficulty. The one subject who rotated the circles frequently (15 times during the session) was also the only non-fluent English speaker in our study. For someone unfamiliar with the language, the ability to view text right-side-up seemed more desirable. Since the terms were already familiar to subjects (they saw the same categories in each condition, and were familiarized with the potential categories during the initial instructions for the experiment), they probably were able to rely on recognizing the labels at a glance, as opposed to needing to read them each time, thereby lessening the importance of the text's orientation. Users were more likely to rotate the controls during the beginning of the condition (76.1% of rotations occurred during the first half of each session), implying that as subjects memorized the labels, there was less need to re-orient them for reading.

Users also did not seem to mind reading rotated text in the popups that displayed currently-assigned metadata – 57.1% of all popups were opened outside of the initiating user's table quadrant, which meant the text would have been oriented away from that user. However, only 47.4% of popups that were opened outside of a user's quadrant were subsequently dragged (dragging the popup into a user's quadrant would result in the text becoming re-oriented to face that user). This indicates that 27.1% of all popups invoked during the study (395 out of 1,459) were not displayed right-side-up for the invoking user. This suggests that reading text at odd orientations may not be as problematic an issue for tabletop interfaces as was initially thought, at least for relatively short text segments.

Orientation of the images did not seem to be an issue either. While it was very common for a user to drag a photo into the center and then ask other group members to look at it to help her answer a question, this image would still only be right-side-up for one of the users. Only on rare occasions did users then pass the image around so that they could each see it right-side-up – they seemed to do this when the image was particularly tricky to recognize, such as when one group passed a photo of a platypus to each group member in turn, because nobody was able to identify it.

Table Regions

As noted in other studies of tabletops [14, 18], users displayed a strong tendency toward not touching regions of the table that were closest to other group members, perhaps because of informal social rules that suggest that the region in front of each user is for her personal use, while the center of the table is a shared region [21].

The central region of the table was clearly important to users. Although people generally kept an image they were working with on the part of the table nearest them, they used the center as a group area. Groups used the center both as an area for directing group attention, by frequently dragging an unidentified photo into the center and then asking other group members for their opinion on what it might be, and also as a place where groups placed photographs that everyone agreed were properly tagged. Despite the fact that in the centralized-controls condition the center of the table was largely occupied by the control widgets, this shared central area was so important to users that they still used the small amount of space between all of the circular controls as a place to put these focal images and "finished" piles.

The questionnaire asked users to indicate on a 7-point Likert scale whether they felt either of the tabletop interfaces seemed cluttered. We hypothesized that subjects might find the centralized-controls interface more cluttered, since it took up valuable center-of-the-table real-estate. Although the responses trended in this direction, there was

no statistically significant difference, with a mean score of 4.21 for the centralized controls, indicating more agreement that the table seemed cluttered than the mean of 3.42 for the replicated controls.

General Usability

We were pleased that subjects found both interfaces easy to learn, indicated by their agreement with the statement "I found it easy to learn how to label photos" for each of the interfaces (mean of 5.75 and 6.33 for centralized controls and replicated controls, respectively, on a 7-point Likert scale). It was also encouraging to note that the groups' discussions focused on the task and not on the interface – the most common type of speech was asking other group members if they knew the appropriate metadata for a particular image.

DISCUSSION

Based on the quantitative and qualitative results of our user study, we can revisit the design questions that we initially posed:

Which design best facilitates collaboration? Although we expected that the centralized controls would facilitate more collaboration than the replicated set, there turned out to be little difference between the two interfaces in this regard. The two interfaces were statistically indistinguishable on all but one of our measures of collaboration (the centralized controls resulted in more contributors per image than the replicated set).

In this study we did not assess the impact of the interface on incidental learning, although we hypothesize that the centralized controls, by increasing the visibility of others' actions, might better facilitate peripheral awareness. Incidental learning and awareness may also be facilitated with the replicated-controls interface by using controls that display feedthrough based on other group members' interactions with their copies of each widget, such as those provided by the MAUI groupware toolkit [10].

Will the oddly-angled text on the centralized controls reduce usability? Although we initially thought that readability of oddly-angled text might be a large factor in making the centralized controls unappealing, from our study we learned that this was not a major consideration for applications where the angled text consists of relatively short strings that remain constant throughout a session, although other work [11] has shown that proper orientation can be important for longer documents. Recent work [26] suggests that deviations from upright orientation impact reading performance only slightly on tabletop interfaces, a finding confirmed by our results.

Which design uses screen real-estate more effectively? Even though the centralized-controls scheme devotes fewer pixels to the control widgets than replicating them for each user, their occupation of the valuable center-of-the-table space was problematic, perhaps because it violated users' tendencies to establish personal "territories" [21] around the table, and because it reduced the availability of the central region for accomplishing shared tasks, such as examining difficult-to-identify photos.

Will users prefer one design over the other? Users had a clear preference for the replicated-controls interface. This preference seems to result from users' desire to use the center of the table for other tasks and their aversion to the physical proximity with teammates' hands created by the shared, centralized control design.

Will the placement of the controls impact comfort? Reach distance did not play a major role in the usability of either interface. Although the tabletop was large (85.6 cm by 64.2 cm), the center of the table was easily within arm's length of most adults. Although none of the users complained that the centralized controls were difficult to reach, many users felt socially uncomfortable about the forced physical proximity of using the shared set of widgets. This negative response to co-touching echoes findings in the area of proxemics (the study of personal space) [9].

Design Guidelines

Based on our study of the centralized- and replicated-controls variants of TeamTag, we offer suggestions relating to two aspects of tabletop interface design: control replication (i.e., how many copies of control widgets should appear on the table) and control location (i.e., where on the table control widgets should be placed).

Control Replication

Our experience with TeamTag suggests that creating multiple copies of frequently-used controls (one copy per user) is a good design strategy. Even though replicating controls uses additional screen real-estate, it alleviates users' proxemic and hygienic concerns that can result from control sharing. However, the generalizability of this advice depends on both task type and group composition.

The photo-labeling task supported by TeamTag enabled a high degree of parallel activity. Tasks that engender a less parallel work style might encounter less difficulty with a shared-controls interface, since the incidence of simultaneous control access by multiple individuals would be far lass likely, thus reducing the potential for problematic co-touching incidents. TeamTag's content also modeled a productivity/education-oriented task (i.e., the photos labeled were from a collection relating to classifying fauna rather than a collection of personal photos). We hypothesize that the intimate proxemics of shared controls may be considered more acceptable in the context of entertainment/socially-oriented tasks, such as tabletop games.

In addition to the formality of a task and the work style it engenders, the composition of a group can also impact the acceptability of a tabletop interface that uses shared, rather than replicated, controls. In our evaluation of TeamTag,

group composition was random – group members did not know each other prior to the study. In many real-world applications, however, group members working together at a tabletop would have a higher degree of familiarity. As a result, the aversion to co-touching that made shared controls unpopular with our test population may be less of an issue for applications whose target user population is more closely-knit groups.

Control Location

Based on our comparison of variants of the TeamTag interface, it seems preferable to design tabletop UIs that locate controls near users' seats, thus leaving the center of the table open. This central space is then available to users for a variety of communicative purposes (e.g., as a focal area for items currently being discussed, or as a storage area for organizing sets of related objects). However, we qualify this recommendation by noting that certain aspects of task type or physical configuration may make controls located along a table's edges less optimal.

The photo-labeling task studied was one in which all data on the table was public (i.e., there were no objects on the table that were restricted for use by only a single member of the group). However, some tabletop tasks involve the presence not only of public data, but also of individually-owned materials. Observations of traditional table use [21] show that people prefer to locate personal and private materials along the edges of the table nearest their seats. Thus, a tabletop UI that locates controls in front of users' seats might reduce the available table-edge screen real-estate for storage of personal materials, and applications that involve a large amount of data of this type may find centrally-located controls to be preferable.

The physical configuration of a tabletop workspace also impacts the preferability of different control-location schemes. In particular, very large table sizes (e.g., interfaces for conference-room sized tables) further reduce the desirability of centrally-located controls since they may be physically unreachable by users. Single, shared copies of controls that have been located near the table's edge in order to preserve the availability of the central region would be similarly problematic for users located at a distant end of a very large table. Re-locatable controls, rather than controls that remain fixed in place, might be preferable in such an environment.

Future Work

Exploring additional variations of the widget-layout design space (e.g., repositionable and/or collapsible widgets) merits further attention. In this study, we intentionally chose the fixed-location widgets in order to facilitate understanding the endpoints of the control-placement spectrum. Allowing users to choose the widget placement themselves would have inhibited our ability to contrast these divergent design choices. However, further work to separate out the effects related to the number of copies of the controls (shared versus replicated) from the effects related to the placement of the controls (center versus borders) would provide additional insight into this issue.

Our study presents an initial exploration of the tradeoffs in deciding whether to use replicated or shared widgets for co-located tabletop applications. Because tabletop technology has only recently been introduced, there is currently a dearth of published advice available to guide the design of multi-user tabletop interfaces. The aim of this study is to increase awareness about the subtle design decisions that can impact the usability and acceptability of tabletop interfaces. We hope that our initial experiences and findings in this area provide a jumping-off point for additional exploration of basic tabletop interface design issues.

RELATED WORK

The introduction of new touch-sensing technologies such as DiamondTouch [7] has facilitated research on tabletop interfaces for supporting collaborative work, such as [14, 18, 23]. However, none of these interfaces deal with issues of centralized versus replicated placement of controls. The Personal Digital Historian (PDH) [22] is a tabletop system that allows a co-located group of users to search through and discuss a collection of digital photos. PDH allowed users to replicate menus around the edge of the table, but did not compare this choice in menu placement to other alternatives. Also, since simultaneous multi-user interaction was not possible with that system (the PDH project used a single stylus for interaction, which had to be shared among group members), the issue of shared versus replicated controls was not as relevant (i.e., the impact of this design choice on parallel work styles was not applicable). Other CSCW systems for photo management, such as [5], focus on distributed interfaces, while our work focuses on co-located systems. The focus of the TeamTag work is not on interfaces for digital photo management; rather, we use a digital photo management application as a motivating example and vehicle for studying the more general issue of control placement on tabletop displays.

Although some systems such as Pebbles [17] and STARS [13] use distributed controls via PDAs in combination with single-display groupware [25], these projects do not explore how a shared-controls design would affect their systems. In contrast, we have presented two designs – shared, centralized controls and replicated controls on a single display – and empirically evaluated their impact on group dynamics and productivity. Sharing a single display is important for a number of applications. For example, research on educational game software for children [19] has found that children playing in pairs found it easier to work and form shared understandings when using a single monitor rather than two separate monitors. Our work focuses specifically on the issue of widget placement for small-group interaction, and does not replicate the entire application on separate displays for each user.

The topic of widget placement for single display groupware is raised by Bederson et al. [2] who propose "local tools" (repositionable widgets). MMM [3] is a multi-device, multi-user editor. The MMM system allowed user-positionable menus to deal with the problem of wasting space by replicating menus on each screen versus only allowing them to appear in one fixed position. In contrast to these systems, our work presents an empirical study to address the issue of widget placement, and compares fixed-place, rather than relocatable, widgets.

Zanella and Greenberg [27] identify design considerations for single display groupware widgets, pointing out that widgets should be placed to avoid interference between users (e.g., one user's view or reach of a widget is blocked by another user), but also to provide awareness of others' activities. We explore these tensions in our comparison of widget placement. In their list of guidelines for tabletop displays, Scott et al. [20] identify "support of simultaneous user actions" and "providing shared access to digital objects" as key areas for future research. Our exploration of impacts of centralized versus replicated control placement on tabletop usability and collaboration is a step toward addressing these issues.

CONCLUSION

We have presented a usability study comparing alternative designs of the TeamTag collective photo annotation software, which differed in whether they provided a centralized set of shared controls or replicated controls for each user. Users strongly preferred the replicated controls for two main reasons: (1) the desire to use the center of the table for other semantically important tasks, and (2) aversion to accidentally touching a teammate's hand when using the shared controls. The unusual orientation of the text on the shared controls, however, did not appear to be a factor in the unpopularity of that design. Our findings relating to co-touching, orientation of information, and the importance of the central region of the table are applicable to the design of a variety of collaborative tabletop interfaces.

ACKNOWLEDGMENTS

We would like to acknowledge Mitsubishi Electric Research Laboratories for donating a DiamondTouch table and providing a license to use the DiamondSpin toolkit. This work was funded by an NSF Graduate Research Fellowship and by the AT&T Labs Fellowship Program.

REFERENCES

1. Adobe Photoshop Album. http://www.adobe.com/products/photoshopalbum/main.html
2. Bederson, B., Hollan, J., Druin, A., Stewart, J., Rogers, D., and Proft, D. Local Tools: An Alternative to Tool Palettes. *Proceedings of UIST 1996*, 169-170.
3. Bier, E., and Freeman, S. MMM: A User Interface Architecture for Shared Editors on a Single Screen. *Proceedings of UIST 1991*, 79-86.
4. Callahan, J., Hopkins, D., Weiser, M., and Shneiderman, B. An Empirical Comparison of Pie vs Linear Menus. *Proceedings of CHI 1988*, 95-100.
5. Crabtree, A., Rodden, T., and Mariani, J. Collaborating Around Collections: Informing the Continued Development of Photoware. *Proceedings of CSCW 2004*, 396 – 405.
6. Diaz-Marino, R., Tse, E., and Greenberg, S. Programming for Multiple Touches and Multiple Users: A Toolkit for the DiamondTouch Hardware. *UIST 2003 Conference Companion*.
7. Dietz, P. and Leigh, D. DiamondTouch: A Multi-User Touch Technology. *Proceedings of UIST 2001*, 219-226.
8. Gutwin, C. and Greenberg, S. Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness. *Proceedings of CSCW 1998*, 207 – 216.
9. Hall, E.T. The Hidden Dimension. New York: Anchor Books, 1966.
10. Hill, J. and Gutwin, C. Awareness Support in a Groupware Widget Toolkit. *Proceedings of GROUP 2003*, 258-267.
11. Kruger, R., Carpendale, M.S.T., Scott, S.D., and Greenberg, S. How People Use Orientation on Tables: Comprehension, Coordination and Communication. *Proceedings of Group 2003*, 369-378.
12. Lave, J. and Wenger, E. Situated Learning: Legitimate Peripheral Participation. Cambridge University Press, Cambridge, UK, 1991.
13. Magerkurth, C., Memisoglu, M., Stenzel, R., and Streitz, N. Towards the Next Generation of Tabletop Gaming Experiences. *Proceedings of Graphics Interface 2004*, 73-80.
14. Morris, M.R., Morris, D., and Winograd, T. Individual Audio Channels with Single Display Groupware: Effects on Communication and Task Strategy. *Proceedings of CSCW 2004*, 242-251.
15. Morris, M.R., Paepcke, A., and Winograd, T. TeamSearch: Comparing Techniques for Co-Present Collaborative Search of Digital Media. *The First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, January 2006, 97-104.
16. Morris, M.R. and Winograd, T. Quantifying Collaboration on Computationally-Enhanced Tables. *CSCW 2004 Workshop on Methodologies for Evaluating Collaboration Behaviour in Co-located Environments*.
17. Myers, B., Stiel, H., and Gargiulo, R. Collaboration Using Multiple PDAs Connected to a PC. *Proceedings of CSCW 1998*, 285-294.

18. Ryall, K., Forlines, C., Shen, C., and Morris, M.R. Exploring the Effects of Group Size and Table Size on Interactions with Tabletop Shared-Display Groupware. *Proceedings of CSCW 2004*, 284-293.

19. Scott, S.D., Mandryk, R., and Inkpen, K. (2003) Understanding Children's Collaborative Interactions in Shared Environments. *Journal of Computer-Aided Learning*, 19(2), 220-228.

20. Scott, S.D., Grant, K., and Mandryk, R. System Guidelines for Co-located Collaborative Work on a Tabletop Display. *Proceedings of ECSCW 2003*, 159-178.

21. Scott, S.D., Carpendale, M.S.T, and Inkpen, K.M. Territoriality in Collaborative Tabletop Workspaces. *Proceedings of CSCW 2004*, 294-303.

22. Shen, C., Lesh, N., Vernier, F., Forlines, C., and Frost, J. Sharing and Building Digital Group Histories. *Proceedings of CSCW 2002*, 324-333.

23. Shen, C., Vernier, F., Forlines, C., and Ringel, M. DiamondSpin: An Extensible Toolkit for Around-the-Table Interaction. *Proceedings of CHI 2004*, 167-174.

24. Shneiderman, B., and Kang, H. Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos. *International Conference on Information Visualization 2000*, 88-98.

25. Stewart, J., Bederson, B, and Druin, A. Single Display Groupware: A Model for Co-present Collaboration. *Proceedings of CHI 1999*, 286-293.

26. Wigdor, D. and Balakrishnan, R. Empirical Investigation into the Effect of Orientation on Text Readability in Tabletop Displays. *Proceedings of ECSCW 2005*, 205-224.

27. Zanella, A. and Greenberg, S. A Single Display Groupware Widget Set. *Western Computer Graphics Symposium 2000*.