The Effects of Interaction Technique on Coordination in Tabletop Groupware

Miguel A. Nacenta*, David Pinelle+, Dane Stuckel* and Carl Gutwin*

*Department of Computer Science
University of Saskatchewan
{nacenta, stuckel, gutwin}@cs.usask.ca

+School of Computer Science
University of Nevada Las Vegas
pinelle@cs.unlv.edu

ABSTRACT

The interaction techniques that are used in tabletop groupware systems (such as pick-and-drop or pantograph) can affect the way that people collaborate. However, little is known about these effects, making it difficult for designers to choose appropriate techniques when building tabletop groupware. We carried out an exploratory study to determine how several different types of interaction techniques (pantograph, telepointers, radar views, drag-and-drop, and laser beam) affected coordination and awareness in two tabletop tasks (a game and a storyboarding activity). We found that the choice of interaction technique significantly affected coordination measures, performance measures, and preference – but that the effects were different for the two different tasks. Our study shows that the choice of tabletop interaction technique does indeed matter, and provides insight into how tabletop systems can better support group work.

CR Categories: H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces—*CSCW*

Keywords

Tabletop groupware, interaction techniques, coordination

1 Introduction

Tabletop groupware systems combine real-world work surfaces with computational interaction techniques, allowing people to collaborate over digital artifacts while still maintaining a face-to-face working style. Tabletop groupware is now becoming more common, and tables have been proposed for a variety of situations, including meetings, design work, and leisure activities.

Users of tabletop groupware manipulate computational artifacts using interaction techniques provided by the application. These techniques can vary along several dimensions—some use virtual cursors [4,21], some use physical reaching [20,31,32,28], and others use three-dimensional pointing [15,5]. This wide variety suggests that there may also be differences in the collaboration that results – for example, some techniques communicate more information about others' actions, and so could result in different levels of group awareness.

Little research has been carried out, however, to investigate how the choice of interaction technique affects collaboration and coordination on tabletops. Ha and colleagues [9] considered the effects of different input devices (mouse and stylus) on group work, but the effects of interaction techniques are still unclear. Without this knowledge, it is difficult for designers to make informed decisions about what techniques to select.

To explore the role that interaction techniques play in tabletop collaboration, we first developed a framework that classifies current techniques according to three dimensions: location of input, location of feedback, and embodiment. We then selected representative techniques from the framework (pantograph, telepointers, radar views, drag-and-drop, and laser beam) and implemented them in two tabletop groupware systems (a cooperative game and a storyboard application). We then compared the techniques in an experiment where ten groups carried out tasks in the game and storyboard systems, using each of the techniques.

Our findings show significant and sometimes unexpected differences across techniques in terms of conflicts, reaching patterns, object transfer, group performance, and user preference. The main result is that no one interaction technique is best for all tasks and all tabletop situations. We found that direct-touch techniques such as drag-and-drop reduce resource conflicts, but at the price of being much less effective when used for reaching distant artifacts. World-in-miniature views such as the radar view were surprisingly effective for a game task, although in a design task the radar did not provide enough awareness information to be successful. There were also strong differences in the ways that users viewed the pantograph, telepointers, and laser-beam techniques, even though these techniques are quite similar on the surface.

This paper makes two contributions: first, the framework provides dimensions on which techniques can be compared and evaluated, and differentiates between a large number of interaction techniques; second, the study provides empirical evidence on how those differences result in changes to group coordination and collaboration in tabletop tasks. The study also provides initial recommendations about where and when the different techniques will be effective, and builds a foundation for the design of novel techniques that can support a wider range of table settings and tasks.

2 Awareness and Coordination on Tables

Group awareness is the understanding of others' presence, locations, and current activities in the shared workspace. A good sense of awareness can help people simplify communication, find opportunities to help one another, and coordinate activities and access to shared resources [14]. Group awareness in tabletop collaboration is maintained through three main mechanisms: territoriality, artifact feedthrough, and consequential communication.

Territoriality is a useful organizing principle for interaction on tabletops. It is based in the physical realities of a person's presence and reach, and the physical extents of the table surface. Scott et al. [25] found that groups divide the space on the tabletop into three territories: personal, group, and storage. These different regions result in different types of shared and individual activity, requiring that awareness information be interpreted in light of the territory in which it was produced.

Feedthrough is the public feedback produced by artifacts when they are manipulated by an interaction technique. Although the feedback is usually intended for the person performing the action, it can also inform others who are watching [6]. In addition, feedthrough can be manipulated to provide better awareness information: for example, delete actions can be drawn out to make them more obvious [8]. One important type of feedthrough on tables is artifact orientation. Collaborators use the orientation of objects in real-world tables to determine use and ownership, and to facilitate communication and coordination with others [12].

Consequential communication is information produced by arms and bodies as people carry out actions in a workspace [23,24]. This mechanism is clearly important in tabletop collaboration, since the co-present setting allows people to directly observe the position, posture, and movement of heads, arms, eyes, and hands. The way that an interaction technique carries out an action in the workspace will have an effect on the amount of consequential information that gets transmitted. Several researchers have observed the value in making objects and actions visible (e.g., [19,24,26]). However, limiting people to interacting using only their real bodies can restrict their freedom and power on a tabletop (for example, when objects are far away, physical reaching can be difficult). Many other interaction techniques have been proposed to overcome some of the disadvantages of using only physical bodies in large-display settings. The framework described below organizes tabletop interaction techniques and sets out the ways that different techniques could contribute to awareness and coordination.

3 FRAMEWORK OF TABLE INTERACTION TECHNIQUES

To better understand their differences, we organize existing interaction techniques by three dimensions: *location of input*, *location of feedback* and *user embodiment*.

Location of Input

Different interaction techniques allow users to provide input from two possible locations: local space or shared space. Local space is the area close to the user (within about one meter), and shared space is the area between the users that is visible to all of them (but not always easy to access). Local and shared spaces may overlap depending on the position of the users relative to each other.

Many tabletop or large-display systems allow direct input on the workspace: for example, by direct touch [32,10,35,27], through a pen [20,16] or with a tangible object [31]. Direct input forces a group to share the input space in order to manipulate the artifacts on the table. Techniques with shared input space include pick-and-drop [20], mediaBlocks [31], and pen-based drag-and-drop.

Other techniques allow users to interact with objects from their own local space: for example, standard mouse-based drag-and-drop [4], hyperdrag [21], and cursor-extension techniques like the pantograph [10,7,16]. These techniques allow all regions of the display to be accessed without the user having to move around the table.

Laser and pointing techniques [17,15,33] are midway between input space and local input space because they can be used from the user's immediate space (local space) and also directly above the objects that we want to manipulate (shared space). However, these techniques have been proven to be error prone and difficult to control when used from a distance [15].

Location of Feedback

The visual feedback provided by an interaction technique allows the users to adjust and correct their actions. Feedback can also be displayed through shared space, local space, or both. Local feedback is used by miniature-view techniques such as radar views [16,30], which allow users to interact with artifacts through local proxies. Miniature-view techniques provide both direct input and access to the entire workspace; however, the miniature may not provide enough resolution to carry out some tasks.

Tabletops offer the advantage that users can share representations of objects. To provide only private representations of space we don't need a tabletop (multiple miniature views in personal displays would suffice). Consequently, tabletop miniature views are used in conjunction with shared visualization, providing thus both local and shared feedback. Some techniques in this category take advantage of the dual representation by explicitly linking the proxies (local space feedback) with the shared space representations [1,2].

Several further techniques provide feedback only in the shared space. These include techniques using shared input space (pick-and-drop, mediaBlocks), local input space (pantograph, drag-and-drop), or both (laser/tractor beam).

User Embodiment

Embodiment is the real or virtual representation of the user's body in the system, and is a valuable mechanism for conveying presence, location, and activity information. Embodiments can vary widely in both form and size, and different interaction techniques show considerable variation in the representation that they provide. The two main dimensions on which embodiments vary are how literal the representation is, and how large.

Techniques that use shared input space (e.g., pick-and-drop or pen-based drag-and-drop) have full-size, literal representations, because they make use of physical bodies and arms [9]. These techniques have intrinsic advantages for awareness since the system does not have to represent the users' points of action (e.g., with a cursor), and because identity is easy to determine.

Other techniques require virtual embodiments – digital representations that are drawn in the workspace. Techniques like pantograph and telepointers provide each user with a cursor that is controlled by the interaction technique. Although cursors can convey awareness, they are much smaller than physical arms, and it is not obvious which cursor belongs to whom. Colour is traditionally used in groupware to differentiate users' cursors, although richer representations such as photographs are also possible. Techniques such as the pantograph [10] go further by drawing a line between the cursor and the user's pen, providing a visual link between each cursor and its owner.

Last, a few techniques use some information from actual bodies as well as the virtual representation of a cursor – in these techniques, some amount of information is available through the user's interaction with the environment. For example, when using the laser beam [5], others can see where the user is pointing, and so gather extra evidence about which cursor belongs to whom.

The three dimensions of the framework are summarized in Table 1. Note some cells contain no techniques: if input is in the shared space, there is little value in local feedback.

Table 1. Interaction techniques (and their embodiment) according to location of input space and feedback space.

		Location of Feedback/Feedthrough		
		Local	Local & Shared	Shared
Input space	Local	Pure radar (no embodiment)	Radar (cursor) Vacuum (arc of influence, lines)	Telepointers (cursor) Pantograph (line and cursor)
	Local & Shared			TractorBeam (cursor+direction of finger/beam).
	Shared			Pick-and-drop Drag-and-drop MediaBlocks (real body)

4 Exploratory Study

In order to better understand the differences between techniques in collaborative tabletop situations, we carried out an exploratory study that compared five techniques that are representative of the different areas of the framework. The study focused on differences between techniques rather than input devices, and so we used only pen-based versions of the techniques (see [9] for a related study of input devices). The study techniques included drag-and-drop (shared input space, shared feedback space), radar view (local input, local and shared feedback space), pantograph and telepointers (local input space, shared feedback space), and laser beam (local or shared input, shared feedback). Pantograph and telepointers differ in the embodiment dimension.

We tested groups of three participants in two tasks. The first task, a cooperative game, was designed to test coordination, synchronization and possibility of conflict. Our intention was to gain insight into activities with high time demands and with clear performance goals. The second task, a storyboarding activity, was designed to be similar to open-ended design and planning work that occurs in everyday office settings.

Apparatus

Experiments were conducted on a top-projected table system connected to a P4 Windows PC and a Polhemus Liberty magnetic tracker with three pens. The table was 125x160cm, and was not touch sensitive (Figure 1). Custom software was written in Java to implement the interaction techniques and tasks (Figure 2).

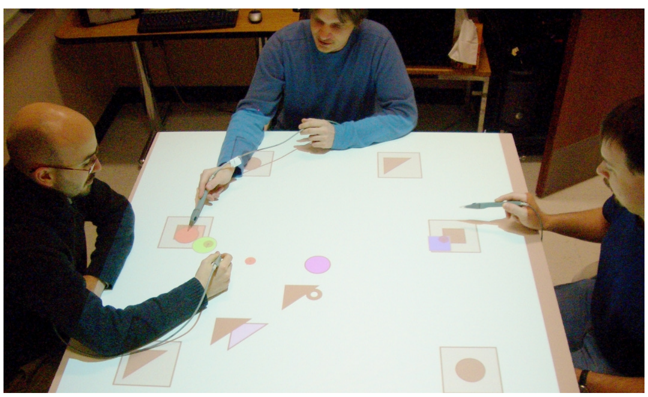

Figure 1. The tabletop setting of our experiments.

Participants

Thirty participants, 14 female and 16 male, were recruited from a local university and organized into groups of three. Participants ranged in age from 17 to 36 years (mean of 24). Each group performed one of two tasks with the five different interaction techniques. Task was a between-subjects factor: groups performed only one of the cooperative game or the storyboard task.

4.1 Procedure

In the *game task*, groups were trained on each interaction technique for two minutes, and then performed a five-minute trial with each interaction technique. In the *storyboard task* there was no specific training, but the trials lasted ten minutes each. The order in which the interaction techniques were presented was distributed among groups using a random Latin square.

During the trials, one participant sat at the head of the table while the other two sat at the adjacent sides (see Figure 1). After each trial they switched places. At the end of the session, participants completed a preference questionnaire. Additional data were collected through system logs and observations.

Interaction Techniques

Drag-and-drop. Users select an object by touching it with the pen, and deselect by lifting the pen from the tabletop. Users can move objects by dragging the tip of their pen across the table (Figure 3).

Radar view. A miniature of the entire workspace (32 x 24cm, 1/5 scale) is displayed in front of each user (Figure 2b). Users manipulate the content of the larger tabletop from within the miniature. Others' cursors, each with a different colour, are visible both in the miniature and the shared space.

Pantograph. The user moves the pen in her local input space as with radar; however, there is no workspace miniature. Instead, the cursor position is proportional to the position of the pen (5X) so that the whole table can be reached with little movement. A colored line is drawn from the cursor to the tip of the pen.

Telepointers. Telepointers are equivalent to the pantograph but without the line that connects the cursor to the pen.

Laser Beam. Users hold the pen above the table and point the tip at the desired location. The cursor is shown at the point where an imaginary line extending from the pen's tip intersects with the tabletop plane. To select an object, the user clicks the stylus button; a selected object can be moved by pointing the pen at a new location on the table.

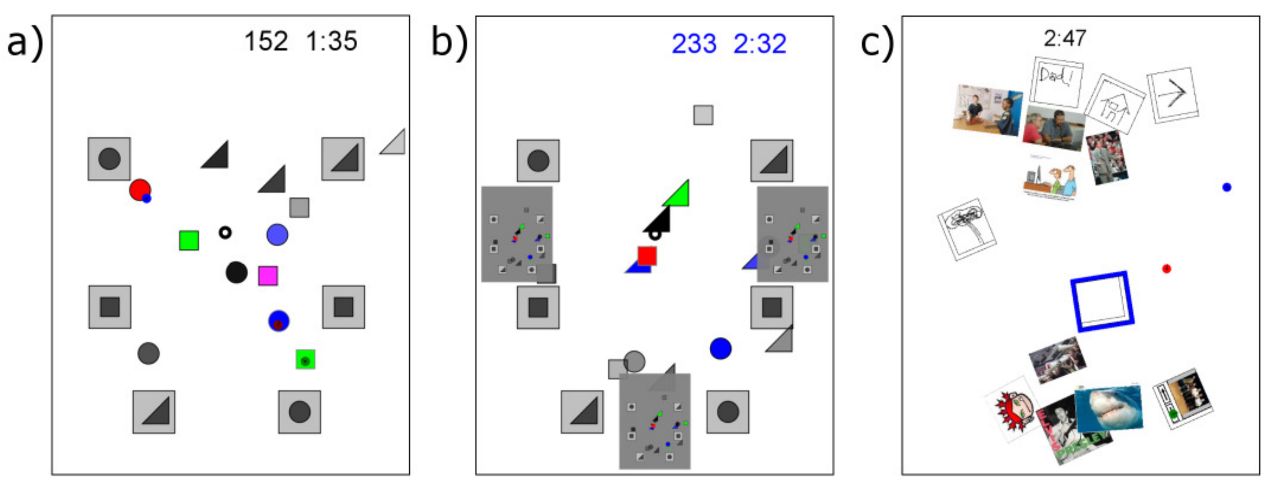

Figure 2. a) The game application, b) the game with radar views, c) Storyboard application.

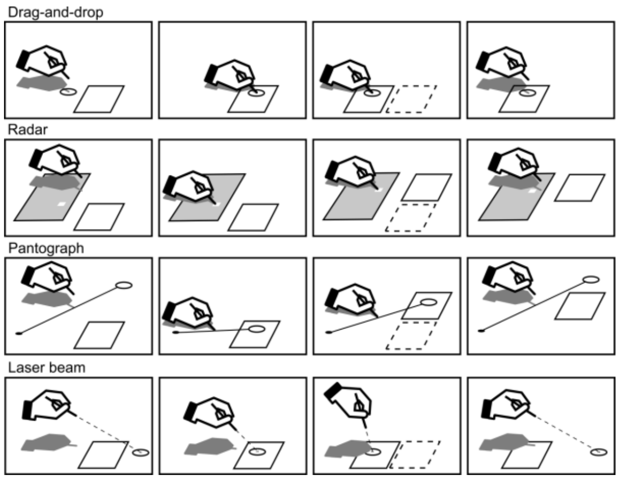

Figure 3. Moving an object with different interaction techniques. The circle in rows 1, 3, and 4 represents the cursor; the grey rectangle in row 2 is the radar's workspace miniature. Telepointers (not shown) is identical to Pantograph but without the line connecting the cursor to the pen tip.

Tasks

Task 1: Cooperative game

The game was a simple puzzle activity where players put shapes into bins as quickly as possible. Triangles, squares, and circles appeared in the center of the screen at random time intervals. Shapes moved away from the center of the screen in a random direction, and each stopped moving after it had traveled a randomly determined distance. Groups received points each time a shape was dropped into a bin that corresponded with the shape, but only if the colour of the shape matched the participant's color (red, green, or blue). There was also a fourth colour (pink) that any of the participants could use – this made conflicts possible, because several users might want to grab a pink shape. While two users interacted with the same shape at the same time, the shape could not be dropped for points; therefore simultaneous grabbing of a shape had no advantage.

Shapes were black when they first appeared in the centre of the table, but revealed their colour when first touched. This ensured that most shapes had to be manipulated by two participants (the participant that revealed the color and the one whose color matches the shape's). To encourage cooperation, only the group score was displayed. In addition, the game awarded extra points if participants were able to place two of the same shape at the same time in different bins. This required two users to coordinate tightly in timing and location of the shape drops, and was the only way to achieve a high score. All the rules and the scoring scheme were explained to the subjects before the start of the trial.

4.1.1 Task 2: Storyboard Activity

Participants were asked to construct a storyline for a short story using images and written notes. Users could write notes with their stylus on virtual pieces of paper taken from a virtual notepad. Images were brought into the table by pulling from a virtual catalog, and participants could navigate through the images by clicking on the catalog's forward or back buttons. The notepad and the catalog could also be passed around the table as well as rotated freely to any orientation (see Figure 3c).

5 Results

We analyzed data from system logs, session videos, and participant questionnaires. Our video analysis focused on general patterns of use and collaboration. We organize the results into five themes related to group coordination: conflicts, transfer patterns, reaching, performance, and preference.

5.1 Differences in conflicts

Conflicts are a useful measure of awareness and coordination. Previous work has clearly established the link between being able to anticipate others' actions, and being able to coordinate activities [8]. The game system recorded the number of times that multiple users simultaneously grabbed the same object. Since it was counterproductive for two users to touch an object at the same time (see Task 1 description), we considered these simultaneous touches to be conflicts.

A one-way repeated-measures ANOVA showed a main effect of technique on conflict rate ($F_{4,16}$=22.43, p<.001). The raw numbers of conflicts were normalized by the number of interactions because different techniques had different overall levels of activity. As can be seen in Figure 4, there were more conflicts when using radar (mean=26.9, SE=3.1) than with the other techniques ($mean_{Drag\&Drop}$=13.6, SE= 2.6; $mean_{Telepointers}$= 8.3, SE=2.28; $mean_{Pantograph}$= 10.5, SE=3.14; $mean_{Laser}$= 4.47, SE= 1.11). Pair-wise post-hoc t-tests showed differences between radar and all other techniques (all t(4)>4.7, p<.01).

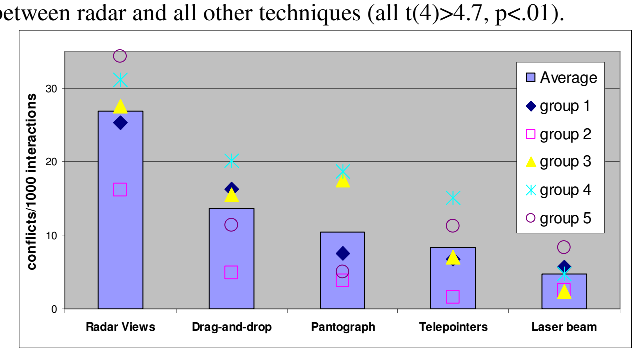

Figure 4. Conflicts per 1000 interactions in the game task. The symbols show the conflicts for each participant group.

In post-experiment questionnaires, participants ranked the techniques according to how well they were able to avoid conflicts. A Friedman test found significant differences between these rankings ($\chi^2(4)$=14.0, p<.01). We combined data from the two tasks, since there was no significant difference between the tasks in terms of rankings.

Participants reported that drag-and-drop (mean rank 2.2, where lower rank is better) and pantograph (mean 2.6) were best at helping to avoid conflicts. Laser beam and telepointers both had an average rank of 3.3, and radar was seen as the worst (mean 3.5). Post-hoc Wilcoxon Signed Rank tests separate the techniques into two groups: drag-and-drop and pantograph were seen as low-conflict while the rest were seen as high-conflict (the difference between pantograph and telepointers is significant, z=-2.58, p<.01, with a Bonferroni correction).

Several participants stated that drag-and-drop provided more information about others' activities since they could directly observe their actions, making it easier to avoid unwanted conflicts (confirming the earlier work on social protocols for coordination [14]). According to one participant, with drag-and-drop "you are able to see an arm reach across the table… which would deter you from grabbing the same shape". Several participants also mentioned the benefits of the pantograph, and stated that the line between the cursor and the pen helped avoid conflicts. The radar was clearly seen as giving the least information for avoiding conflicts: as one participant stated, "it was horribly difficult for this because you couldn't see the physical movements of the others" (even though telepointers were shown both in the miniature-views and in group space).

5.2 Differences in transfer patterns

Transferring objects to another user is a second indicator of whether people can successfully coordinate their activities. Participants ranked the techniques according to ease of transferring objects to others. A Friedman test showed significant differences in these rankings ($\chi^2(4)$= 15.7, p>.01). Again, tasks were considered together as there was no significant variation in rankings between the tasks. Participants felt that drag-and-drop and pantograph were best at supporting transfer (mean ranks of 2.3 and 2.5); the other techniques were markedly lower (laser, 3.7; telepointers, 3.2; radar, 3.3). Post-hoc Wilcoxon Signed Rank tests show that the difference between pantograph and telepointers is not significant when using a Bonferroni correction (z=-2.38, p=.017).

In the questionnaires, the most commonly cited reason for thinking drag-and-drop was useful for transfers was because it was clear where people were located in the workspace. One person wrote that "we knew where people were physically located, so if we put an item in front of them, they can notice it easier". The stated reasons were similar for pantograph, but were based on the use of a virtual embodiment: "it is easier when you've got the line pointing to the place". Participants gave similar reasons for feeling that the other techniques were not as useful in facilitating transfers—that it was more difficult to know where others were working, and where objects should be left. Even though others' cursors were represented in the miniature views in real time, one participant wrote about radar: "I just dragged the objects that were for others wherever because I didn't know very well where the others' cursors were". Clearly, the colored cursors provided by the miniature views were not as easy to track as the embodiments provided by other techniques.

The video records also showed marked differences in the ways that people transferred objects. In the game task, users could only deposit a shape if it had the same color as their cursor. Although transfer occurred with all techniques, drag-and-drop allowed people to rapidly exchange shapes and to drop shapes in front of others based on their colors.

5.3 Differences in reaching patterns

The applications recorded activity maps during the sessions to summarize the amount of activity by different users on different parts of the table. Every time that a user clicked on an object, released an object, dragged the cursor or moved the cursor, it left a trace on the maps (Figure 5). These maps allow us to tell how often different parts of the table were used with the different techniques.

The maps show obvious differences in the movement patterns of both the tasks and the interaction techniques. The structure of each task shows up in the activity maps. In the game, the six bright spots of interaction correspond to the bins where shapes could be deposited; in the storyboard the activity traces are distributed more homogeneously. More importantly, there are also consistent differences across techniques. Drag-and-drop and laser beam show highly regionalized activity (strong colors in each region). Pantograph shows more homogeneous use of the space, but not as mixed as telepointers and radar views (for which it is difficult to distinguish red, green, or blue regions. Similar differences appear for the storyboard maps (second row), even though this task was far less constrained.

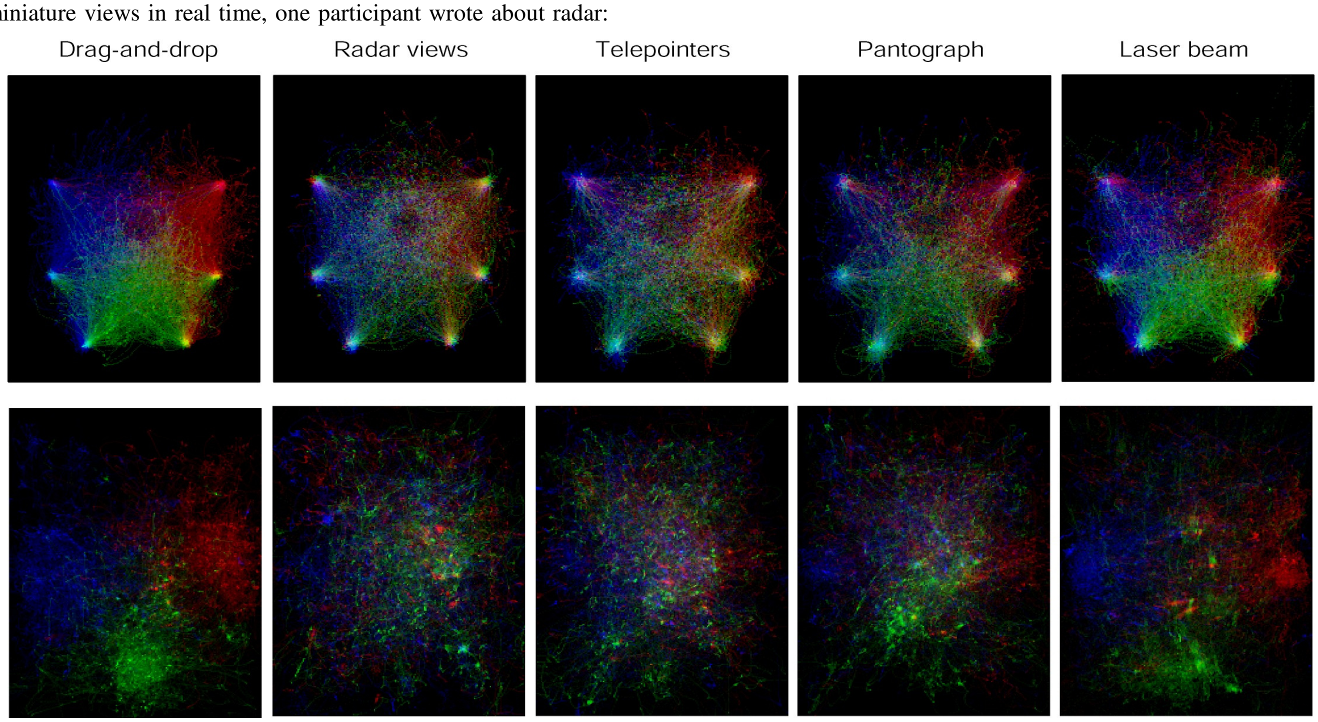

Figure 5. Movement maps for all techniques in the game task (first row) and the storyboard task (second row). Each map aggregates data across groups. Each of the primary colors (blue, green, red) represents the users sitting to the left, bottom, and right of the image respectively. Higher color intensity represents a higher number of cursor movements in that pixel. White pixels are points accessed equally by all users.

Although only two of the eight sets of maps are reproduced here, all other groups of maps (clicks, releases and drags for each task) show similar regionalization patterns.

In the questionnaires, participants ranked the techniques according to how well they support reaching outside their personal territory. A Friedman test showed significant differences in these rankings ($\chi^2(4)=58.0$, $p<.001$), and data from both tasks were considered together. Figure 6 shows the mean rankings given to each technique. Radar views was ranked first on average (mean 1.7), followed by pantograph (mean 2.4) and telepointers (mean 3.0)—all techniques explicitly designed to allow interaction at a distance. The techniques that require direct or indirect physical access to the table surface were ranked next (laser beam, 3.3; drag-and-drop, 4.6). Post-hoc Wilcoxon Signed Rank tests show that the difference between pantograph and telepointers is significant ($z=-2.87$, $p<.005$), as is that between laser and drag-and-drop ($z=-3.36$, $p=.001$).

People reported that the radar was easiest to use since it gave access to all items on the table: one participant wrote that the "radar view and pantograph allow easy collection of distant objects". People felt that drag-and-drop was difficult since "you cannot reach the objects which are far away".

The videos also showed differences in reaching. For example, in a group carrying out the storyboard task, one individual took a leadership role, directing the layout of images and notes through most of the activity. However, when the group used drag-and-drop, he was not sitting near the storyboard layout, so he had difficulties contributing to the task, and the group struggled. In the game task, we also observed differences in the way groups managed distant artifacts with the techniques. When people used drag-and-drop, they placed shapes primarily into nearby bins. However, when people used radar, pantograph, or telepointers, they also used distant bins since they were all equally accessible.

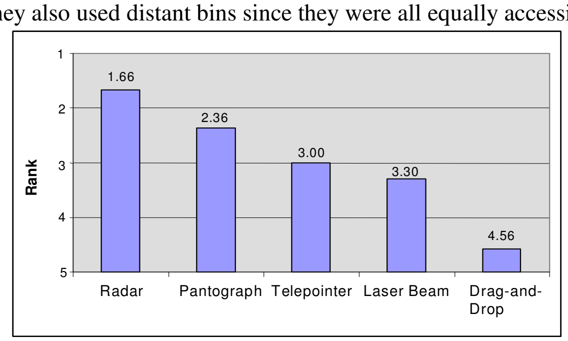

Figure 6. Preferred techniques for access to distant areas. Taller bars mean better ranks.

5.4 Differences in performance

In the game task, participants were given a group score based on their performance in the game. While the scoring was somewhat arbitrary, we attempted to encourage group cooperation by providing more points for synchronized actions than for individual actions (5X). Participants were instructed in the scoring system, and were encouraged to try to maximize their score with all techniques, so the score provides a limited indicator of how groups performed using each technique during the game task.

One-way repeated-measures ANOVA showed a significant main effect of technique on score ($F_{4,16}=8.45, p<.001$). Scores were higher with drag-and-drop and radar (means 1133 and 1099) and lower with the other techniques (see Figure 7). Post-hoc t-tests showed a significant difference between radar and pantograph ($t(4)>4.6$, $p=.01$), but failed to find a difference between drag-and-drop and pantograph ($t(4)>2.925$, $p=.043$, using a Bonferroni correction).

Several participants stated that they were able to perform better using drag-and-drop, since it closely mirrors the way that they manipulate objects on real tables. According to one participant, "it is the most effective one because it's similar to our daily life—when we need to get something or transfer something." While people felt that radar view led to more conflicts, many also indicated that they were able to achieve good performance using that technique. One participant stated: "although there were many conflicts with the miniature view, the smaller distance to drag helped speed the process up. One could throw an object to a further bin faster" (note however that the index of difficulty of the targeting task was the same, so this comment relates to the difficulty of reaching across the table).

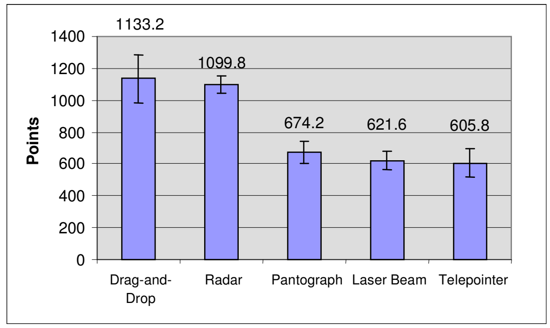

Figure 7. Mean scores (error bars = 95% confidence intervals)

5.5 Preference

Participants were asked to rank the techniques according to their personal preference. We carried out independent-sample Mann-Whitney tests to compare the rankings for each technique in the two tasks and found that the rankings were different. The differences are striking—radar view was highly ranked for the game task, but was the lowest by far for the storyboard task (see Figure 8). The change in ranking between tasks for the radar technique was significant ($z=-3.86$, $p<.001$).

Due to the differences between tasks, the data were analyzed separately for each task. Two Friedman tests showed significant differences between techniques in both the game task ($\chi^2(4)=33.4$, $p<.001$), and the storyboard ($\chi^2(4)=15.6$, $p<.003$).

Drag-and-drop was the favorite for both tasks. In the game task, radar and drag-and-drop were preferred over the other three techniques, as shown by a significant post-hoc Wilcoxon Sign Rank test between radar and pantograph ($z=-2.8$, $p=.005$). In the storyboard task, radar was the least preferred among all techniques, but this difference is not significant with the Bonferroni correction ($z<-2.3$, $p<.021$).

Questionnaire results suggest that the temporal constraints of the game may have played a role in people's preference for certain techniques. In the game task, we observed that people often focused exclusively on the radar view and rarely looked at the main table display. Many people indicated that they liked the radar since it allowed them to move objects efficiently to anywhere on the table—for example, one person wrote that "with the radar view you can access the whole field with very little movement". In the storyboard task, however, many people stated that the radar view was too small for the task. One person wrote that, "in the radar view objects were too small, so I had to switch my view from the mini view to the table to see what objects are". We observed this regularly in the videos—people repeatedly switched back and forth between the miniature view and the main table view so that they could see the details on the artifacts and details of other's actions. Additionally, in the storyboard task people had to collaborate more closely, and several people indicated that the radar view shifted their focus away from what others in the group were doing, making it more challenging to work together on the activity.

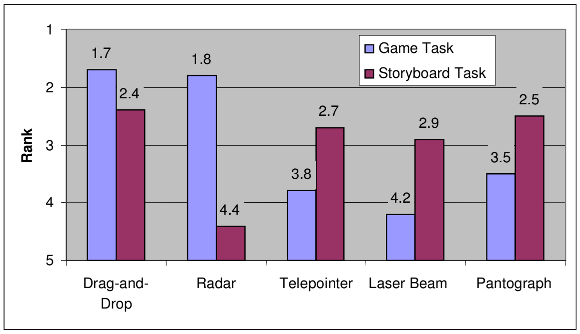

Figure 8. Preferred techniques (taller bars are better).

In the storyboard task, the higher opinions that were given to telepointers, laser beam, and pantograph seem to be linked to the need to access distant regions of the table. In the game task, people could ignore shapes that were out of their reach with only a minimal penalty. However, in the storyboard task, task dependencies forced users to access specific objects and table regions that might be out of their reach, making techniques that allow distant reaching potentially more useful. It is indicative that, for the game task, many people stated that the laser beam technique was frustrating and made their arms tired, while in the storyboard some participants stated that laser beam was "the most functional", "the simplest and the best", and "the easiest to use". Nevertheless, this was not enough to make laser beam the preferred technique in the storyboard task.

6 Discussion

The two tasks of this study were designed to cover a wide range of hypothetical tabletop tasks: from time-constrained tightly-coordinated tasks to informal group tasks similar to those studied previously [28]. Unfortunately, tabletop displays are not yet widespread, and there exist no canonical tabletop activities or applications. Therefore the results of the study should be interpreted in the light of the particular focus of this study and compared to the desired application domain. Nevertheless, we believe that having empirical evidence over the effect of interaction techniques on low level issues such as coordination, performance and reaching patterns will help researchers and designers to anticipate problems in table groupware interaction technique choices.

The single most important lesson learned is that the choice of interaction techniques in tabletop groupware applications significantly affects the way that users carry out group activities. Our study suggests that each technique has strengths and weaknesses, and that the overall utility of each technique is strongly related to the task that it supports. In the next paragraphs, we synthesize our findings and make design recommendations based on the strengths and limitations of each technique.

Drag-and-drop. Drag-and-drop had the best all-around performance and rankings of the techniques included in this study. The benefits of the technique are related to its parallels with the way that people interact with objects in the real world—it is easy for people to learn the technique, and it is easy for people to interpret others' actions. The user's body acts as the embodiment, allowing people to collect more information about others' actions and intentions than with most virtual embodiments. The main limitation of the technique is that people have difficulties with accessing resources that are outside of their immediate reach. This limitation could potentially become more of a problem as the size of the table increases, and more areas are out of reach.

Recommendation: Drag-and-drop enables people to track others' actions, and is the best technique for tabletop collaboration when reaching demands are low.

Radar. Radar had mixed results. It had good overall performance, and was highly ranked on preference scores by users of the game application. Furthermore, people felt it was the best technique at helping them to interact with items that were out of their reach. However, preference scores were much lower on the storyboard task, where radar had the lowest ranking, and it had some obvious coordination problems. More conflicts were seen with radar than with other techniques, and people felt that they were less aware of others' actions. These results suggest that radar is a good technique choice when a group task can be carried out in parallel, such as the game task, but it is not a good choice when there are strong task dependencies and when people must interact with each other to accomplish a task (e.g. the storyboard task). Furthermore, the value of radar may be increased when the table size is large, and when users have to work in distant areas.

Recommendation: Radar is a good choice when tasks can be carried out in parallel, and where awareness requirements are low. It is one of the best techniques for working with objects that are out of peoples' reach.

Pantograph. Pantograph had high rankings on reach and transfer and it had a high preference ranking on the storyboard task, where awareness requirements were high. Most participants' comments indicate that they felt that they were able to track other's pointers successfully using the technique (but there were a few negative comments as well). Overall, pantograph provides the best combination of embodiment and reaching, although it did not clearly win in any performance measures or rankings.

Recommendation: Pantograph is a valuable technique when both reaching and awareness are important.

Telepointers. Telepointer and laser beam had comparable measures in both tasks, and they uniformly had the lowest rankings. Telepointer is an indirect technique (local input, shared feedback), and there is no spatial mapping between the pointer and the person. This made it difficult for people to stay aware of others' actions, and people also had difficulties at times with keeping track of their own pointer.

Recommendation: While telepointers allow people to reach distant objects, its minimal embodiment limits its usefulness in group tasks. Spatial embodiments need to be added to improve the overall usability of the technique.

Laser beam. Intuition suggested that laser beam would be useful in group situations since other users can interpret body and hand position to gain awareness of each other's activities (shared input space, shared feedback). However, people had difficulties with targeting distant and small objects using the technique, and these individual problems seemed to outweigh any possible collaborative benefits. Furthermore, we observed that users prefer to use the individual input space and point without moving very far from their positions rather than reaching into the shared space (individual input space rather than shared), thus reducing the amount of information available to others about their intended target.

Recommendation: Poor accuracy limits the usefulness of laser beam. Consequential communication is low due to limited virtual embodiments and limited physical pointing in the shared space.

We believe that many of the limitations that are seen in existing techniques can be improved by developing new hybrid interaction techniques. New techniques can combine the strengths of direct (shared input space, shared feedback as seen in drag-and-drop) and indirect techniques (local input, as seen in telepointers, pantograph, and radar) to help overcome two main problems: 1) reaching limitations seen in direct techniques, and 2) embodiment limitations seen in indirect techniques. We believe that adding enhanced embodiments, for example, to radar will allow users to reach distant items but also stay aware of others' activities. Pantograph is the only technique that combines aspects of both approaches, but we believe that the embodiment can be enhanced beyond the simple line representation that is used in the technique.

7 CONCLUSIONS AND FUTURE WORK

Designers of tabletop systems have little guidance when attempting to choose interaction techniques that will support both individual work and group awareness. This paper provides designers with two contributions: first, a framework that differentiates between a large number of interaction techniques, and provides three dimensions on which techniques can be compared and evaluated; and second, an exploratory study that provides empirical evidence on how those differences result in changes to group coordination and collaboration in tabletop tasks.

Our results show that there are significant effects of the type of interaction technique on conflict, transfer, reaching, performance, and user preference. These data allow initial recommendations to be made about where and when the different techniques will be effective, and provide a foundation for the design of techniques that can support a wider range of table settings and tasks.

In our future work we want to add further empirical data to the framework that can help generalize to other more specific tasks and to explain some of the phenomena observed in the study. Second, we will develop and test the new techniques suggested by our initial work – in particular, extensions to drag-and-drop to improve reaching, and embodiment enhancements for radar and pantograph to improve group awareness.

8 REFERENCES

[1] Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P., Bederson, B., and Zierlinger, A., Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems, *Proc. Interact 2003*, 57-64.

[2] Bezerianos, A. and Balakrishnan, R. The vacuum: facilitating the manipulation of distant objects, *Proc. CHI 2005*, 361-370.

[3] Biehl, J. and Bailey, B., ARIS: An Interface for Application Relocation in an Interactive Space, *Proc. Graphics Interface 2004*, 107-116.

[4] Brignull, H., Izadi, S., Fitzpatrick, G., Rogers, Y., Rodden, T., The introduction of a shared interactive surface into a communal space, *Proc. CSCW '04*, 49-58.

[5] Davis, J., and Chen, X., LumiPoint: Multi-user laser-based interaction on large tiled displays. *Displays,* 23(5), 2002, 205-211.

[6] Dix, A., Finlay, J., Abowd, G. & Beale, R., *Human–Computer Interaction*, Prentice–Hall, 1993.

[7] Geiβler, J. Shuffle, throw or take it! Working Efficiently with an Interactive Wall, *Proc. CHI 1998*, 265-266.

[8] Gutwin, C. and Greenberg, S. Design for individuals, design for groups: tradeoffs between power and workspace awareness. *Proceedings CSCW 1998.*

[9] Ha, V., Inkpen, K., Mandryk, R., and Whalen, T. Direct Intentions: The Effects of Input Devices on Collaboration around a Tabletop Display, *Proc. IEEE TableTop 2006*, 177-184.

[10] Hascoët, M., Throwing models for large displays, *Proc. HCI 2003*, British HCI Group, 73-77.

[11] Hawkey, K., Kellar, M., Reilly, D., Whalen, T., Inkpen, K., The Proximity Factor: Impact of Distance on Co-located Collaboration, *Proc. GROUP 2005*, 31-40.

[12] Kruger, R., Carpendale, S., Scott, S., and Greenberg, S., How people use orientation on tables: Comprehension, coordination and communication, *Proc. GROUP 2003*, 369-378.

[13] Magerkurth, C., Memisoglu, M., Engelke, T., Streitz, N. Towards the next generation of tabletop gaming experiences, *Proc. Graphics Interface 2004*, 73-80.

[14] Marwood, D. and Greenberg, S. Real time groupware as a distributed system: concurrency control and its effect on the interface. *Proceedings CSCW 1994*, pp. 207-217.

[15] Myers, B., Bhatnagar, R., Nichols, J., Peck, C., Kong, D., Miller, R., and Long, A., Interacting At a Distance: Measuring the Performance of Laser Pointers and Other Devices, *Proc. CHI 2002*, 33-40.

[16] Nacenta, M., Aliakseyeu, D., Subramanian, S., and Gutwin, C., A Comparison of Techniques for Multi-Display Reaching, *Proc. CHI 2005*, 371-380.

[17] Parker, J., Mandryke, R., and Inkpen, K., TractorBeam: seamless integration of local and remote pointing for tabletop displays, *Proc Graphics Interface '05*, 133-140.

[18] Pinelle, D., Gutwin, C., Greenberg, S. Task Analysis for Groupware Usability Evaluation: Modeling Shared-Workspace Tasks with the Mechanics of Collaboration. *TOCHI 10(4)*, 2003, 281-311.

[19] Pinelle, D., Gutwin, C., & Subramanian, S. Designing Digital Tables for Highly Integrated Collaboration. *Technical Report HCI-TR-06-02*, Computer Science Department, University of Saskatchewan, 2006.

[20] Rekimoto, J. Pick-and-Drop A Direct Manipulation Technique for Multiple Computer Environments, *Proc. UIST 1997*, 31-39.

[21] Rekimoto, J. and Saitoh, M. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments, *Proc. CHI 1999,* 378-385.

[22] Ringel, M., Ryall, K., Shen, C., Forlines, C., Vernier, F. Release, relocate, reorient, resize: fluid techniques for document sharing on multi-user interactive tables, *Proc. CHI 2004*, 1441-1444.

[23] Segal, L., Designing Team Workstations: The Choreography of Teamwork, in *Local Applications of the Ecological Approach to Human-Machine Systems*, P. Hancock, J. Flach, J. Caird and K. Vicente ed., Lawrence Erlbaum, Hillsdale, NJ, 1995, 392-415.

[24] Segal, L. Effects of Checklist Interface on Non-Verbal Crew Communications. *Contractor Report 177639.* NASA Ames Research Center, 1994.

[25] Scott, S., Carpendale, S., and Inkpen, K., Territoriality in collaborative tabletop workspaces, *Proc. CSCW 2004*, 294-303.

[26] Scott, S., Grant, K., & Mandryk, R., System Guidelines for Co-located Collaborative Work on a Tabletop Display, *Proc. ECSCW 2003*, 159-178.

[27] Shen, C., Everitt, K., Ryall, K., UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces, *Proc. UbiComp 2003*, 218-288.

[28] Shen, C., Lesh, N.B., Vernier, F., Forlines, C., Frost, J. Sharing and building digital group histories, *Proc. CSCW 2002*, 324-333.

[29] Shen, C., Vernier, F., Forlines, C., and Ringel, M., DiamondSpin: an extensible toolkit for around-the-table interaction, *Proc. CHI 2004*, 167–174.

[30] Swaminathan, K. and Sato, S. Interaction design for large displays. *Interactions*, 4, 1, January 1997, 15-24.

[31] Ullmer, B., Ishii, H., Glas, D., mediaBlocks: physical containers, transports, and controls for online media, *Proc. SIGGraph 1998*, 379-386.

[32] Vernier, F., Lesh, N., Shen, C. Visualization techniques for circular tabletop interfaces, *Proc. AVI 2002*, 257-263.

[33] Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large high-resolution displays, *Proc. UIST 2005*, 33-42.

[34] Voida, S., Podlaseck, M., Kjeldsen, R., and Pinhanez, C., A study on the manipulation of 2D objects in a projector/camera-based augmented reality environment Proc. CHI 2005, 611-620.

[35] Wu, M., Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays, Proc. UIST 2003, 193-202.