Fluid Orientation on a Tabletop Display: Integrating Rotation and Translation

Russell Kruger, Sheelagh Carpendale, Tony Tang, Stacey Scott
Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4
{krugerj, sheelagh, tonyt, sdscott}@cpsc.ucalgary.ca

ABSTRACT

Previous research has shown that rotation and orientation of items plays three major roles during the course of tabletop collaboration: comprehension, coordination and communication. Based on these roles of orientation, we designed a novel tabletop rotation mechanism, Rotate 'N Translate (RNT), which provides integral control of rotation and translation using only a single touch-point for input. We present an empirical evaluation comparing RNT to the standard rotation mechanism that separates control of rotation and translation. Results of this study indicate RNT is more efficient than the separate mechanism and better supports the communicative and comprehensive roles of orientation, without sacrificing support for coordination.

Categories and Subject Descriptors

H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces – *collaborative computing, computer-supported cooperative work, evaluation/methodology.*

General Terms

Design, Experimentation, Human Factors, Performance

Keywords

Tabletop collaboration, orientation, roles of orientation, rotation, communicative gesture.

1. INTRODUCTION

Recently, there has been an increasing interest in tabletop displays to support small groups collaborating in co-located settings. This has arisen naturally from the fact that tables have long been important in collaborative activities, coupled with the fact that many collaborative activities now involve digital information and tools. This interest has been reflected in both research literature [e.g., 13, 14, 15] and industry response through the design of new types of tabletop display hardware [3]. However, developing appropriate interfaces for tabletop displays still presents many new challenges.

One of these challenges arises because tabletop displays are horizontally oriented. On vertical displays, all collaborators share the same orientation—that is, the top and bottom of the display are the same for everyone. This assumption, however, is not necessarily the case for groups using a tabletop display. Collaborators seated at different sides of a table have very different views of the display. Although earlier research has recognized that orientation in collaborative tabletop settings is a complicated issue [16], the common response from tabletop interface designers has been to provide system support for reorienting objects "right way up" for individuals [9, 11, 13]. However, a recent observational study [7] has identified three major roles that orientation plays during the course of tabletop collaboration:

- **Comprehension:** It is easier to comprehend objects when they are "right way up." This is the commonly recognized role.
- **Coordination:** Orientation is used to help establish personal and group spaces and to signal ownership of objects.
- **Communication:** Orientation is useful in initiating communicative exchanges and in continuing to inform group members about collaborative work patterns.

Because the manipulation of tabletop objects is an extremely common activity in tabletop collaboration [7], well-designed interaction mechanisms can have a dramatic effect on the collaborative process. In this paper, we present the design, implementation, and evaluation of a novel tabletop interaction technique that combines rotation and translation into a single, fluid interaction mechanism (RNT). We begin by presenting our design guidelines, which build on the previously described roles of orientation and additional kinesthetic principles, and then describe the behavior of the new rotation mechanism. Following this, we present a user study that demonstrates the efficacy and efficiency of the new technique, and conclude with a discussion of the study results.

2. ROTATION DESIGN GUIDELINES

To maintain the benefits that tabletop environments afford for collaboration, it is important to support the various kinds of interactions common to tabletops. Object reorientation and translation are ubiquitous tabletop actions, and are often used in collaboration in ways identified by the three roles of orientation. As a result, three of our design guidelines for a new rotation mechanism build on implications drawn from these roles [7]. The fourth is based on kinesthetic studies about rotation and translation [17], and the last is based on a review of tabletop interface literature [12]. Our design guidelines are as follows:

- **Free rotation:** The technique must support the rotation of an item to any desired angle.
- **Lightweight:** The technique must place minimal overhead on a rotation action. Since reorientation of objects is a common tabletop action, a lightweight rotation mechanism will help ensure the collaborative processes of orientation transfer into digital workspaces.
- **Feedthrough:** To preserve the communicative role of orientation, it must be obvious to others when an individual is rotating an object. As well, to maintain the long-term communicative effects of orientation, the item's orientation must be clearly visible after the rotation action has occurred.
- **Integrated:** The interaction technique should integrate rotation with translation. Kinesthetic studies of the manipulation of physical objects have demonstrated that rotating and translating are inseparable actions in the physical world [17]. Integral interaction techniques are more appropriate for manipulating integral attributes, and separable interaction techniques are more appropriate for manipulating separable attributes [6].
- **Technology independent:** An important secondary goal for tabletop interaction techniques is platform and technology independence. Collaborative interfaces for tabletop displays are still actively being researched, and it is unclear that any one tabletop system will be suitable for all types of collaborative tasks [12]. Thus, developing interaction mechanisms that require specialized technology may place premature limits on desirable tabletop configurations.

Providing integrated control of rotation and translation may provide better support for the three roles of orientation than a separated interaction mechanism. For instance, the intentional communicative role of orientation, where an individual rotates and moves an object towards a collaborator to initiate communication, may be more fully supported with an integrated, free rotation mechanism. This type of rotation is of particular interest to the collaborative process, as rotating an object to favor one's collaborator is a compelling communicative gesture. In an observational study of tabletop collaboration [7], we saw that this action always initiated discussion and was a well-understood method for starting immediate collaboration. Moreover, we have noticed this communicative gesture in a variety of situations during our day-to-day interactions. Thus, providing a mechanism that allows people to draw on their familiarity with this gesture may be particularly appropriate for tabletop interfaces.

RNT allows a digital tabletop object to be simultaneously rotated and translated using only a single touch-point to control the object. It provides clear feedthrough of rotation actions, indicates item position and orientation throughout the action, and maintains the orientation of rotated items when the manipulation is complete. By using only a single touch-point, RNT supports one-handed interaction and incurs minimal overhead for producing rotation. As well, RNT is technology-independent—it can be used with both indirect and direct input systems, and requires no specialized hardware.

3. RELATED WORK

Rotation is often clumsy in electronic settings, largely because current input devices (e.g., mouse, keyboard, and stylus) provide few degrees of freedom compared to the range of manipulations possible with one's hand on a physical object. Several approaches to digital rotation have been considered, including manual selection of a rotation angle, specific rotation "modes" accessed either through menus or specific controls, and novel input methods that provide an extra degree of freedom (which can then be used to control rotation). Unfortunately, these approaches do not satisfy the design requirements given above.

Manually selecting a rotation angle, as is done in graphics applications such as Microsoft Paint, allows an object's degree of rotation to be precisely controlled and maintained by the user. However, these schemes require significant overhead to produce a rotation action, and thus cannot be considered lightweight.

A common method of producing rotation on computers is to use a rotation mode—that is, use a specific rotation control (accessed directly or through a menu) to rotate an object around a fixed point. For example, in Microsoft PowerPoint, a rotation handle is provided to allow an object to be rotated around its center. This use of a rotation control point is one way of providing rotation when using a two-dimensional pointing device such as a mouse. In general, a two-dimensional device (providing x and y) is insufficient to simultaneously control both the object's position (x and y) and orientation (θ). One method of overcoming this limitation is to use the location of interaction to determine whether the object is to be rotated or translated. For instance, the "corner to rotate" method uses the following convention: initiating contact at the object's center is used to translate the object, while initiating contact at a corner is used to rotate the object. This method is commonly used for rotating digital objects on a tabletop display [14, 15].

Other rotation schemes that use an additional degree of freedom have the capacity to provide integrated control of rotation and translation. Examples of these schemes include providing an additional input device to be used by an individual's non-dominant hand [2], providing multi-touch interaction [18], or using an input device that has three or more degrees of freedom [e.g., 1, 4, 5, 8]. However, all these schemes require non-traditional input devices, and are therefore technology-dependent.

While a single two-dimensional pointing device is unable to simultaneously control the location (x and y) and orientation of an object (θ) as separate functions, it is possible to develop a mechanism that integrates control of these attributes. A recent technique developed by Mitchell [10] called "Drag," allows a person to use a regular mouse to control an object's position while "friction" on the object produces changes in orientation. As the object "drags" across the surface, digitally-simulated frictional forces cause it rotate.

Our RNT mechanism also uses a simulated force, in our case a "current," to integrate rotation and translation. However, the behavior of RNT and Drag are sufficiently different to produce noticeable differences in interaction. Furthermore, separate evaluations of these two techniques produced considerably different empirical results. During Mitchell's evaluation, participants had significantly more difficulty operating Drag compared to the traditional mechanism that provides separate control of rotation and translation. Users felt that they did not have sufficient control of Drag and that they could not adequately predict its behavior—results which stand in sharp contrast to results of the empirical evaluation of RNT presented in this paper.

4. RNT BEHAVIOUR

RNT allows an object to be simultaneously rotated and translated in a single fluid motion, controlling *x* and *y* position plus orientation (θ) using a single 2D contact-point. The metaphor of a "current" can be used to illustrate the rotational behavior of the item as it is manipulated. As an object is moved, a current acts against the object in opposition to the movement vector. If the direction of movement changes, so too does the current. When the object is stationary, no current exists. As the object is manipulated, the current acts against the object to produce rotational changes, while the movement vector yields positional changes.

If the contact-point of the object combined with the movement vector creates a balanced counter current, the object does not rotate (Figure 1(a)). In contrast, the object will rotate if the contact-point and movement vector creates a counter current that is imbalanced against the edges of the object. In Figure 1(b), contact is initiated in the upper-right corner and movement is directly upwards. The current acts in opposition, placing more "pressure" on the left side of the object, thereby rotating the object counterclockwise. In producing rotational changes, the current always acts against the full-width of the object as it is moved. In Figure 1(b), the object rotates counterclockwise until a balanced relationship with the current is achieved. In this diagram, the dashed arrows represent the increased width of the current acting on the object edges. As the object rotates, additional dashed arrows act on the right-hand side of the object until they balance the arrows on the left-hand side. Once rotational balance is achieved, the object continues to translate at the established orientation as long as the movement vector remains unchanged.

If the contact-point and movement vector result in extremely unbalanced current pressure against the object, rotational changes are more extreme (Figure 1(c)). The greater the unbalance, the greater the rotation per unit of movement. An exception to this rule is a central translate-only region depicted by a circle in the middle of the object. If the contact-point is located inside this region, the object is strictly translated without being affected by the current (Figure 1(d)). Prior to commencement of the empirical evaluation, it was determined that a translate-only region enabled users to exercise more precise control over the behavior of objects, and was therefore deemed necessary to include.

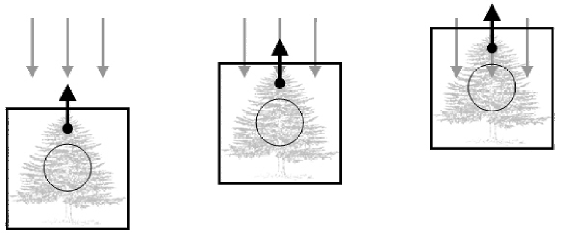

Figure 1(a): Balanced movement resulting in upward translation from a control point located in the upper-half of the object.

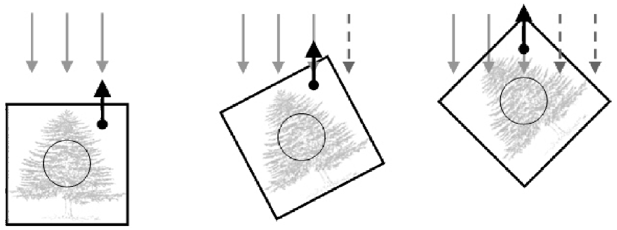

Figure 1(b): Unbalanced movement resulting in upward translation and counterclockwise rotation from a control point located in the upper-right corner of the object.

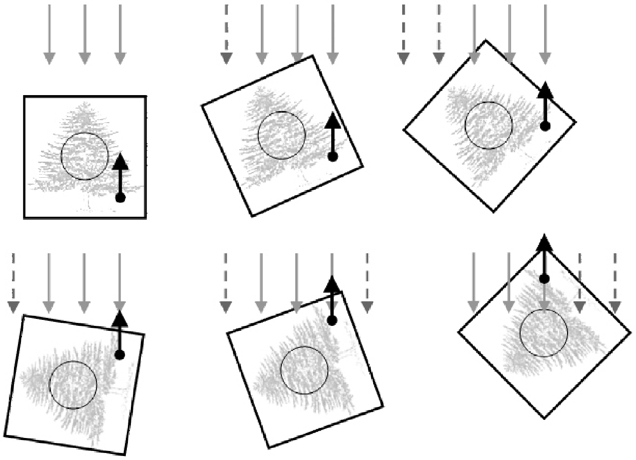

Figure 1(c): Unbalanced movement resulting in upward translation and counterclockwise rotation from a control point located in the lower-right corner of the object.

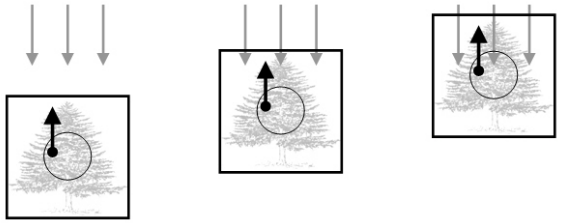

Figure 1(d): Upward translation from a control point located in the translate-only region.

Figure 1: Illustration of integrated rotation and translation.

5. RNT ALGORITHM

There are three points of interest for determining the algorithmic behavior of RNT: *C*—the Center of the object; *O*—the Original mouse position or contact-point, and *T*—the new mouse position or Target. Figure 2 represents the case where the initial contact-point is located near the upper-left corner and movement is slightly down and to the right. In this case, the object is translated by the vector *OT* and rotated about point *O*. The angle of rotation is the angle θ, formed by *O*, *C*, and *T*. Note that the figure shows a large distance between *O* and *T*; however, since the algorithm processes at 30 frames per second, in practice, the vector *OT* tends to be only a few pixels long. After translation and rotation changes are applied, the original target point (*T*) becomes the new original mouse position (*O*); point C remains the center position of the object, and a new target point (*T*) is determined by the next directional movement.

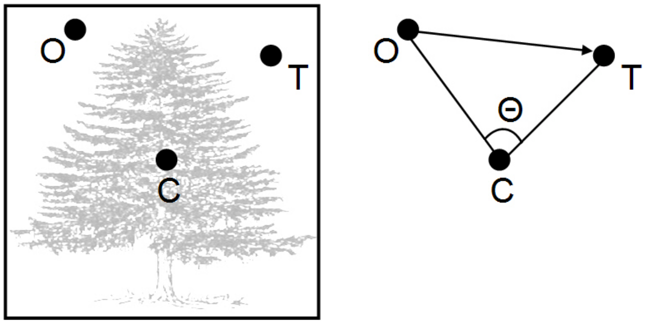

Figure 2: Three required points determine the RNT algorithm: C—the Center of the object; O—the Original mouse position; and T—the new touch position (or Target).

6. EMPIRICAL EVALUATION

To evaluate RNT, we conducted an empirical study comparing RNT against the standard "corner to rotate" mechanism. For this experiment, the "corner to rotate" mechanism behaved as follows: to translate an object, it was simply required to touch the object and move it in the desired direction. The object would follow the contact-point, maintaining its existing orientation. For rotation, the object first had to be "selected" by touching the object to invoke four corner circles. Then, the object had to be touched inside one of these circles and rotated by moving circularly either in a clockwise or counterclockwise direction. During this motion, the center of the object was used as the pivot point around which the object rotated. Figure 3 shows an example of interaction using this mechanism: the object is stationary in frame (a), selected in frame (b) (as a result of having touched the object), rotated in frame (c) (after having touched and rotated a corner circle), and finally translated in frame (d) (after releasing contact in the corner circle and touching and dragging the object elsewhere).

6.1 Method

6.1.1 Participants

Eighteen paid participants (9 male, 9 female), fifteen of whom were right-handed, were recruited from the university population to take part in the study. All were proficient computer users (although most were not computer science majors), and only three had used a tabletop display before. The majority of participants were self-rated beginners in terms of experience with existing digital rotation mechanisms.

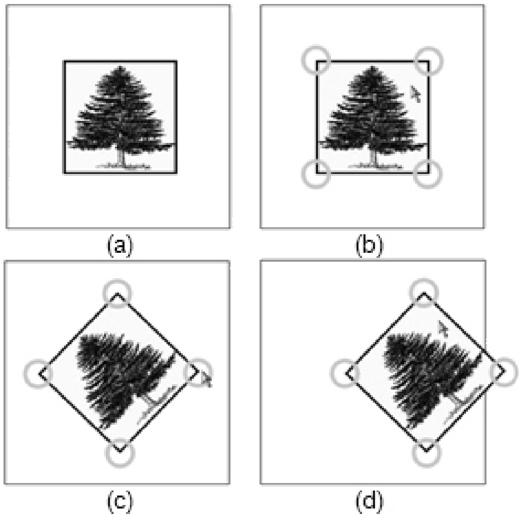
(a) (b) (c) (d)

Figure 3: Sequence of images illustrating the "corner to rotate" mechanism.

6.1.2 Apparatus

The experiment was conducted on a top-projected, touch-sensitive tabletop display powered by a Pentium IV-based PC running Windows XP. A projector running at a resolution of 1280 x 1024 pixels was used to project images onto a horizontal SMART DViT board (Figure 4). The visible display area on the board measured approximately 40.6 inches wide x 50.8 inches long. Although SMART's DViT system allows for simultaneous touch detection, the technology does not currently allow a touch to be matched to the person who initiated the touch. Thus, the collaborative group task was designed so that only one person would need to interact with the table at any give time.

6.1.3 Experimental Design

Participants completed three tasks during the course of the study: two involving object translation and rotation, and the third involving a collaborative document-passing task. The study used a within-participants design in which all eighteen participants completed each of the three tasks using both rotation techniques. Each participant performed the first two tasks individually, and then performed the third task together with two other participants (who had already completed Tasks 1 and 2). The order in which participants used the two rotation mechanisms for each task was counterbalanced.

6.1.4 Tasks

The first task was designed to evaluate the efficiency of the two techniques as a lightweight, free rotation mechanism as required by the comprehensive role of orientation. The second and third tasks were designed to determine the extent to which the respective mechanisms support the communicative role of orientation as embodied in the act of passing and rotating tabletop objects. Tasks were not designed to address the coordinative role of orientation, since this role requires only that a mechanism maintain the position and orientation of an object after it has been rotated.

Task 1: Precision Targeting. The purpose of Task 1 was to compare RNT and the separate technique for a precise rotation and translation task. Such a task is an important benchmark from which to compare the two mechanisms, since a rotation technique for a tabletop display needs to be able to support precise movements to address comprehensive needs as well as to support general tabletop activity. In this task, users were required to reposition and reorient an image of a tree that initially appeared upright and in the centre of the tabletop display to a new location and orientation (Figure 4). The tree object was square, measuring 128 x 128 pixels (or 5.1 x 5.1 inches). The target, outlined with a 10 pixel-wide border, appeared at one of eight different positions circularly located around the middle object at a distance of 275 pixels (10.9 inches), and was oriented at one of sixteen different orientations (0º, 22.5º, 45º, … , 292.5º, 315º, and 337.5º). Participants were instructed to position the object such that part of the black border was visible on all sides. Participants completed 32 trials for Task 1, with the target object appearing four times at

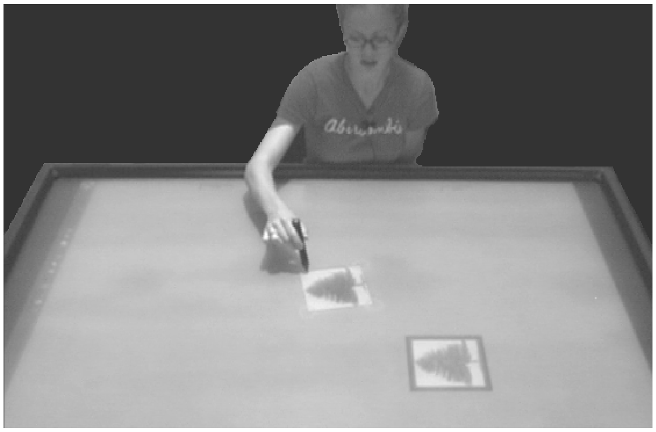

Figure 4: Task 1 involved positioning the central tree image over the target image (thickly outlined).

each of the eight different target locations, and twice at each of the sixteen different target orientations. The ordering and pairing of the locations and orientations were randomly chosen.

Participants were instructed to overlay the tree object on top of the target as quickly and accurately as possible. After the object had been placed on the target to the participant's satisfaction, they could manually advance to the next trial by touching either of the two "Next" buttons at the bottom of the screen. Allowing users to choose when the trial ended allowed us to compare the accuracy of the two techniques.

The metrics used to compare the techniques for Task 1 included time, accuracy, number of touches, required touch distance (each of which was programmatically captured by the application), as well as user preferences. Accuracy was measured in pixels for the location offset (distance between the image centers), as well as the rotation offset (difference in degree of rotation). Touches included all contacts made with the board during the trial. Touch distance measured (in pixels) the amount of contact participants had with the table through the course of the trial. Each participant was asked to rate the techniques on learnability, ease of use, enjoyment of use, and suitability of the technique for the task, as well as to indicate an overall preference for one of the two techniques for the task performed.

Task 2: Document Passing. The purpose of Task 2 was to compare RNT and the separate technique for a less-precise rotation and translation task that attempted to mirror a real-world collaborative tabletop activity—the passing of documents. In this task, users were required to pass an image of a document that initially appeared upright at the bottom of the screen to representations of individuals at three different sides of the table (Figure 5). During the task, participants were asked to imagine themselves as legal assistants for a law firm. As legal assistants, their job was to pass documents to lawyers positioned around the table in such a way that the receiving lawyer could read the passed document.

The "readability" stipulation arose out of the observational study [7], which suggested that objects tended to be passed in collaborative tabletop settings to facilitate readability. In passing the objects "right way up," participants were free to decide orientations that satisfied this condition. Participants were also free to decide the appropriate proximity of the passed document to the lawyer representations. The lawyers at the left and right positions were located approximately 535 pixels (21.2 inches) from the original document location, while the lawyer opposite the participant was located approximately 630 pixels (25.0 inches) from the original document location.

At the start of each trial, the document appeared in an upright orientation at the bottom of the screen. Participants were instructed to pass the object as quickly as possible to the lawyer indicated (the one with the stars on either side of their name). After each pass was made to the participant's satisfaction, pressing either of the two "Next" buttons advanced them to the next trial. Task 2 required thirty trials, with ten document passes to each of the three lawyers. The ordering of these passes was randomly chosen.

The metrics used for Task 2 were similar to those used for Task 1: time, required number of touches, required touch distance, and user preferences—the only difference being accuracy. In this case, the goal of the task was not to pass the document to a precise target location or orientation, but rather to the general vicinity of a lawyer representation and at an angle suitable for reading. In so doing, we sought to determine whether participants thought "inaccuracy" in passing was acceptable for this type of task. Orientation and proximity measurements were also collected and compared to measures taken from a real collaborative setting (Task 3) to determine whether participant behavior for the second task was ecologically valid. The remaining metrics were gathered in the same fashion as they were for Task 1.

Task 3: Collaborative Document Passing. The purpose of Task 3 was to compare RNT and the separate technique for a document-passing task in a collaborative tabletop setting. Groups of three participants were responsible for completing a word puzzle by passing and decoding clues to form a completed sentence. The task proceeded as follows: a clue object would appear at one of the three collaborator's positions (Alpha, Beta or Gamma as seen in Figure 6), oriented "right way up" for that position. Each clue specified a clue receiver to whom the object was to be passed. Upon receiving the clue, the clue receiver would copy down letter-number pairs (such as "E1 F1 A2 E1 B3") contained on the

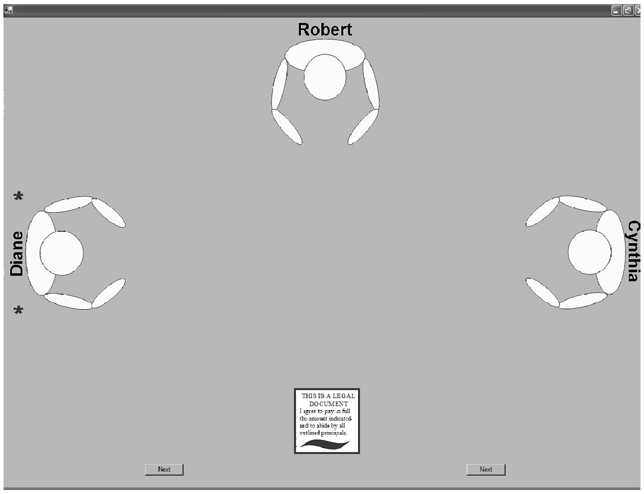

Figure 5: Task 2 involved passing the document image to representations of people seated at different sides of the table.

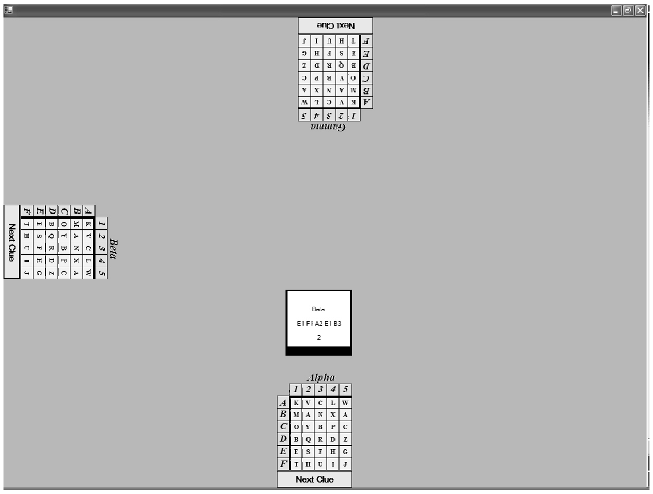

Beta
E1 F1 A2 E1 B3
2

	1	2	3	4	5
A	K	V	C	L	W
B	M	A	N	X	A
C	O	Y	B	P	C
D	B	Q	R	D	Z
E	E	S	T	H	G
F	T	H	U	I	J

Figure 6: Task 3 involved passing a clue (outlined in black with a thick black bottom border) to participants seated at two different sides of the table.

clue using pen and paper on a provided clipboard. They would then dismiss the clue by double-touching on the "Next Clue" button located at the bottom of their decoder table (which was used for converting letter-number pairs into a decoded letter). This would cause the next clue to randomly appear at one of the three positions, though not necessarily in front of the person who just dismissed the clue.

When not passing or receiving clues, participants were instructed to decode the clues using the decoder table located in front of them. Once the letters were decoded, the word remained to be unscrambled. For instance "E1 F1 A2 E1 B3" decoded to form the letters "E T V E N" which unscrambled to form the word "EVENT."

To complete the task, each participant passed three clues to each of the other collaborators (for 18 clue passes in total). The order of these passes was randomly chosen. The approximate distances between the positions (from the center of each person's decoder table) were as follows: Alpha to Beta (and vice versa) – 615 pixels (24.4 inches); Beta to Gamma – 615 pixels (24.4 inches); Gamma to Alpha – 745 pixels (29.6 inches). In passing the clues, participants were free to decide both the proximity and orientation that was appropriate for the receiving person. If needed, the receiving person was also free to adjust the clue as desired.

Participants were told that they as a group were competing against five other groups to complete the task as quickly as possible. Thus, they were instructed to adopt any strategy that would help speed up their completion time. To allow us to compare the techniques on the basis of their communicative role, we recommended to participants that they work on decoding and unscrambling their own clues while not passing to or receiving clues from others. Thus, the participants would be busy when clues were passed to them, allowing us to compare the feedthrough generated by the techniques in causing clue receivers to notice the passing gesture.

Once all clues had been passed and dismissed, the group was allowed to work collectively on decoding and unscrambling words, as well as piecing together the final sentence. Participants were encouraged to place their clipboards on the tabletop display to facilitate their work. The task was deemed complete when a member of the group read the decoded sentence out loud.

Whereas Task 2 gathered precise time measurements for passing gestures made by each of the two techniques, Task 3 examined these gestures in the context of a collaborative activity so as to provide an opportunity for higher-level evaluation of the techniques. This collaborative task was also intended to provide additional data regarding the use of orientation as a communicative gesture. To this end, the metrics gathered for Task 3 were in part similar to the metrics used for Task 2, including required touches, touch distances, and user preferences. In addition to these metrics though, video data was collected from each of the six groups and later analyzed to aid in a high-level evaluation of the techniques. Time to pass clues was not measured due to the imprecision inherent in measuring time for this task.

6.1.5 Procedure

The experimental sessions were run as follows. First, individuals filled out a consent form and a pre-session questionnaire, and were given an introduction to the tabletop display. The experimenter then explained Task 1. Next, for each rotation technique, the experimenter explained the technique allowing the participant to perform 32 practice trials followed by 32 task trials. Once both set of trials had been completed, participants were given a short rest. The experimenter then explained Task 2, allowing 30 practice trials followed by 30 task trials for each rotation technique.

Task 3 was completed in a separate session on a later date, in which participants completed a reduced number of practice trials (five) before completing Task 3. After each task, participants completed a post-task questionnaire which gathered preference data for the two techniques. At the end of the third task, participants were debriefed and paid for their involvement in the experiment.

7. RESULTS

In total, 1152 trials were completed for Task 1, 1080 trials were completed for Task 2, and 216 trials were completed for Task 3. Results of these trials are presented in the sections that follow, organized by task.

7.1 Task 1: Precision Targeting

Our results show that RNT was faster and more efficient than the separate mechanism for precision targeting, while maintaining equal levels of accuracy. Users required an average of 8.18 seconds to complete the task with the separate mechanism, compared to only 7.40 seconds with RNT (Table 1). An ANOVA showed a main effect of rotation technique on completion time ($F_{1,17}$=11.967, $p < 0.005$). Accuracy, both in terms of location and rotation offset, was not significantly different between the two techniques (p=0.703 and p=0.866, respectively).

Similarly, RNT outperformed the separate mechanism on both the number of required touches, as well as the required touch distance. Using the separate mechanism, users needed 3.67 touches on average, compared to an average of only 2.09 touches for RNT. Likewise, users required longer touch distances on average for the separate technique (732.86 pixels or 29.07 inches) as compared to RNT (607.21 pixels or 24.09 inches). An ANOVA

Table 1: Result means and standard deviations for Task 1

Technique	Measure				
	Time (seconds)	Location Offset (pixels)	Rotation Offset (degrees)	Touches	Touch Distance (pixels)
Separate	8.18 (1.16)	2.74 (0.65)	1.33 (0.42)	3.67 (0.69)	732.86 (98.84)
RNT	7.40 (1.19)	2.79 (0.54)	1.34 (0.27)	2.09 (0.52)	607.21 (130.46)

showed a main effect of rotation technique on both the number of touches ($F_{1,17}$=76.754, $p < 0.001$) and touch distance ($F_{1,17}$=17.935, $p < 0.005$).

Participants were asked to rate each mechanism on a scale of 1 to 5 (5 being very preferable) for four different measures. Ratings for the two techniques were similar, though surprisingly, the separate mechanism was rated slightly better than RNT: 4.6 vs. 4.3 for ease of learning; 4.2 vs. 3.8 for ease of use; 3.8 vs. 3.7 for enjoyment level; 4.1 vs. 3.9 for task suitability. Participants were split as to their overall preferred mechanism, with eight users preferring RNT, and ten users preferring the separate mechanism.

7.2 Task 2: Document Passing

RNT outperformed the separate mechanism on all measures for Task 2. In terms of speed, RNT produced quicker completion times than the separate mechanism, with users requiring on average only 3.55, 3.63, and 3.42 seconds to pass the document to the left, middle and right lawyers respectively (Table 2). In contrast, users needed an average of 4.61, 4.79 and 4.50 seconds to complete the task with the separate mechanism. An ANOVA showed a main effect of rotation technique ($F_{1,17}$=97.999, $p < 0.001$) and target position ($F_{2,34}$=6.475, $p < 0.005$) on completion time, with no interaction between these variables. Post-hoc analysis showed significant differences between the times required to pass the document to the middle vs. right lawyer for both rotation techniques. The lower completion times for the right-most lawyer are likely due to the close proximity of this position (both the left and right lawyers were slightly closer than the middle lawyer as described above), combined with a handedness effect due to most participants being right-handed.

As with precision targeting, users were more efficient at document passing using RNT compared to the separate mechanism. On average, users needed only 1.21, 1.15, and 1.19 touches with RNT for passes to the left, middle and right lawyers respectively. In contrast, users required almost three times as many touches using the separate mechanism (on average, 2.93, 3.03, and 2.96 touches). An ANOVA showed a main effect of rotation technique on the number of touches ($F_{1,17}$=224.253, $p < 0.001$), with no effect for position ($p = 0.947$). Similarly, users required greater touch distances for the separate mechanism (772.13, 928.72, and 764.10 pixels or 30.63, 36.84, and 30.31 inches for the left, middle and right lawyers respectively) compared to RNT (502.42, 618.88, and 493.98 pixels or 19.93, 24.55, and 19.59 inches for these same positions). An ANOVA showed two main effects of rotation technique ($F_{1,17}$=295.184, $p < 0.001$) and position ($F_{2,34}$=72.321, $p < 0.001$) on touch distance, with no interaction between these variables. Post-hoc analysis showed significant differences between the touch distance required to pass the document to the middle vs. right lawyer as well as the middle vs. left lawyer for both rotation techniques. This result is very likely due to the further distance between the initial starting position of the document and the middle lawyer as described above.

Table 2: Result means and standard deviations for Task 2

Technique	Measure		
	Time (seconds)	Touches	Touch Distance (pixels)
Separate (Left, Middle, Right)	4.61 (0.10)	2.93 (0.46)	772.13 (66.28)
	4.79 (0.84)	3.03 (0.68)	928.72 (96.06)
	4.50 (0.84)	2.96 (0.57)	764.10 (85.07)
RNT (Left, Middle, Right)	3.55 (0.77)	1.21 (0.24)	502.42 (71.66)
	3.63 (0.82)	1.15 (0.30)	618.88 (81.05)
	3.42 (0.74)	1.19 (0.29)	493.98 (73.46)

For document passing, participants showed a strong preference for RNT. RNT was rated higher on average for each of the four preference categories: (4.9 vs. 4.7 for ease of learning; 4.8 vs. 3.9 for ease of use; 4.6 vs. 3.2 for enjoyment level; 4.7 vs. 3.2 for task suitability). As well, an overwhelming 16 of the 18 participants preferred RNT for this task.

7.3 Task 3: Collaborative Document Passing

Results from Task 3 further reinforce the effectiveness of RNT compared to the separate mechanism. In terms of required touches, passers needed only 2.99 touches on average to complete the task using RNT, compared to the 5.20 touches that were needed on average with the separate mechanism (Table 3). An ANOVA showed a main effect of rotation technique on the number of required touches ($F_{1,5}$=79.489, $p < 0.001$). Similarly, users required longer touch distances for the separate technique compared to RNT. For the separate mechanism, the average required touch distance was 1093.04 pixels (43.36 inches), while RNT required only 817.64 pixels (32.43 inches). Again, an ANOVA showed a main effect of rotation technique on touch distance ($F_{1,5}$=29.972, $p < 0.005$).

Participant ratings for the two techniques were similar to those for Task 2, with participants showing a strong preference for RNT for this task. RNT was rated higher on average for all four rating categories: (4.7 vs. 4.4 for ease of learning; 4.5 vs. 3.7 for ease of

Table 3: Result means and standard deviations for Task 3

Technique	Measure	
	Touches	Touch Distance (pixels)
Separate	5.20 (1.04)	1093.04 (154.18)
RNT	2.99 (1.06)	817.64 (145.35)

use; 4.2 vs. 3.2 for enjoyment level; 4.5 vs. 3.2 for task suitability). Also, 14 of the 18 participants preferred RNT for Task 3.

7.4 Video Analysis

A detailed video analysis of the collaborative task was conducted to supplement the low-level statistical data, and to aid in a high-level evaluation of the two techniques. To this end, each of the 216 passes were analyzed according to the following categories:

- **Passing orientation and proximity:** Did participants pass the clues so that the receiver could view them "right way up?" How accurate (in terms of orientation and proximity) did the passer leave the clue for the receiver? Did inaccuracy affect the collaboration in any noticeable way? Were any adjustments made by the person receiving the clue?
- **Receiving of clues:** At what point did the clue receiver notice the clue was intended for them? At what point during the passing motion were the clue contents copied down?
- **Order of rotation:** For the separate mechanism, if participants passed the clue according to the orientation of the receiver, did they rotate the clue before or after it was translated? For RNT, did participants actually use the simultaneous rotation and translation capability of the technique, or did they separate out these actions?
- **Issues/problems:** Were there any problems with either of the techniques for the passing task? Did people ever change the direction in which a clue was being passed, and if so, did the mechanism support this change?

7.4.1 Passing Orientation and Proximity

In all cases and without instruction from the experimenter, the person responsible for passing the clue left it at an orientation that was "right way up" for the receiver of the clue. This likely indicates that the requirement of Task 2, namely passing the document so that it was "right way up" for the lawyer representations, was appropriate, thereby lending further ecological validity to results obtained from that task.

After a clue was passed, the receiver rarely made adjustments to the orientation or proximity of the clue. 196 of the 216 passes (90.7%) involved no adjustments by either the passer or the receiver after the document was passed. Of the adjustments that were made, most were performed by the passer of the clue (17 adjustments, or 7.9% of the 216 passes), and were minor in nature (i.e. less than 10 degree rotation and less than 2 inch translation). Only three adjustments were made by the receiver of the clue over the course of the 216 trials, and these were also minor in nature. Thus, participants generally felt comfortable with the orientation and proximity of the passed clues. This is significant because the clues were oriented on average at a final angle of 9.59° offset from an orientation orthogonal to the table edge for the separate technique, and at a final angle of 13.25° for RNT (rotation angles were programmatically recorded during task completion and were used to assist in the video analysis process). Moreover, there were numerous examples of final orientations (including those after adjustments) in which the degree of rotation was 20°+ off from an orthogonal orientation (42 of the 216 passes or 19.4%). Yet, people were comfortable with reading clues at angles of this nature. In terms of proximity, data was collected that measured the distance from the center of the passed clue to the center of the receiving person's table. For Task 3, the clues on average ended up a distance of 149.14 pixels (5.92 inches) away for the separate technique, and 159.59 pixels (6.33 inches) away for RNT. Thus, people were willing to read and copy down clues that were not positioned directly in front of them.

When proximity and orientation results from Task 3 are compared to those from Task 2, it is apparent that participants treated the lawyer representations much the same way as they did their real collaborators in Task 3. Rotation differences from an orientation orthogonal to the table edge for Task 2 were 8.29° for the separate technique and 6.68° for RNT; both of which are smaller but comparable to those seen in Task 3. It may be noted that rotation differences for RNT are actually smaller than those for the separate mechanism in Task 2—results which are not mirrored by Task 3. This is likely due to the fact that with repetition of the passing gesture, RNT allowed participants to be more consistent and "accurate" in their passes than the separate mechanism due to use of the same single-motion passing gesture. Since Task 3 did not involve a highly repetitious task, this effect was not as strong for the third task. Offset differences for Task 2 (measured from the center of the passed document to the center of the lawyer representations) were 141.92 pixels (5.63 inches) for the separate technique, and 152.95 pixels (6.07 inches) for RNT. These numbers are only slightly smaller than those seen in Task 3. Thus, the manner in which participants passed documents to the lawyer representations in Task 2 mirrors their behavior towards their real collaborators in Task 3, thereby lending further ecological validity to the results obtained from Task 2.

In general, results from the video analysis of Task 3 indicate a tendency of individuals to be slightly less-precise with RNT for the task of passing a document in a collaborative setting as compared to the separate technique. For the separate mechanism, it appears that more attention is paid to the orientation of the clue and the document proximity, which may be due to the separate rotation and translation phases. On the other hand, because rotation and translation are performed simultaneously using RNT, it appears that individuals focus more on the high-level passing gesture with RNT than on specific rotation angles and document locations. Results of the video analysis indicate a willingness of individuals to work with documents aligned at these imperfect angles and distances without any noticeable effects on the collaboration. Thus, the somewhat less-precise nature of RNT for the task of passing a document is unlikely to hinder collaboration; rather, since this imprecision is comfortable and familiar to collaborators, it may actually enhance collaboration by allowing individuals to exploit their ability to be imprecise (as is currently the case with paper-based media [7]). Additionally, this imprecision has implications for the design of tabletop interfaces: although many systems typically orient tabletop objects orthogonally to the table edge [e.g., 9, 11], results from the video analysis indicate that this strategy is unnecessary and may even hinder collaboration by requiring individuals to work with orientations that are less familiar and/or less comfortable to them.

7.4.2 Receiving of Clues

While Task 2 focused solely on the passer in the context of a passing gesture, Task 3 provided an opportunity to focus on the receiver of the passed object as well. Results of the video analysis indicate no clear difference between the techniques as to when a receiver noticed the passing gesture. There were numerous examples for both techniques of the receiver noticing the clue when it first appeared in front of the passer (i.e. the receiver read the clue upside down or sideways); when the clue was being rotated towards the receiver (either in place as with the separate technique or as it was being moved with RNT); and when the clue was being moved towards the receiver (regardless of its orientation). Although we initially believed the reorientation of the clue as it was being passed (as is the case with RNT) might result in increased feedthrough and thus a quicker response by the receiver, it appears that the nature of the task prevented this from occurring. That is, since participants were instructed to complete the task as quickly as possible, it was in their best interests to be vigilant in noticing when a clue was passed towards them. The simple act of a collaborator's arm moving towards them was sufficient to attract their attention. It appears that the consequential communication provided by the passer's body was a more salient a cue than the feedthrough generated by either technique. Thus, this task did not reveal significant differences between the techniques in terms of the receiver's role in intentional communicative gestures.

7.4.3 Order of Rotation

The video analysis revealed that RNT was used in an integrated fashion. That is, only 4 of 108 passes (3.7%) involved separate control of rotation and translation using the translate-only region of the mechanism in conjunction with rotating the clue in-place. Thus, it appears participants were willing and able to use the mechanism as it was intended.

For the separate mechanism, it was possible to adopt two primary strategies to complete the passing task. The first was to pass the clue in its initial orientation to the receiver, and then rotate the clue in front of them; the second approach was to first rotate the clue, and then pass it to the receiver. The majority of passes were done according to the second strategy (78 of 108 passes, or 72.2%), although a quarter of passes were done according to the first (29 of 108, or 26.9%). This result is important because the first strategy (pass, then rotate) appeared to be somewhat awkward for collaborators in that the clue passers had more difficulty in rotating the object at this increased distance. As well, this rotation took place in an area more accessible to clue receivers than passers, which produced noticeable awkwardness for clue receivers in these situations. For instance, on a number of occasions (6 of 29 or 20.7%) after the clue was passed, the receiving collaborator attempted to rotate the clue themselves, but stopped when the passer did so. Thus, it would seem that the first strategy may be more awkward in terms of both reach and personal space issues. Since it comprised over a quarter of all passes using the separate mechanism, the separate mechanism appears to be more awkward for the purpose of passing a document on a tabletop.

7.4.4 Issues/Problems

Participants did not have major problems with either of the two rotation techniques. For the separate technique, there were sporadic problems associated with trying to initiate contact in the corner circles and failing to do so (this occurred for 8 of the 108 trials, or 7.4%). For RNT, inadvertent or over-rotation was a sporadic issue, also occurring for 8 of the 108 trials (7.4%). In these cases, the person attempting to pass the clue tried to do so too quickly, and as a result, applied too much "force" (and thus rotation) to the clue.

Another sporadic issue of interest involved participants passing a clue to a particular position, only to realize mid-pass they were passing it to the wrong person (thus requiring a change of direction). This occurred 6 times in total (2 for the separate mechanism, and 4 for RNT). The interest here is that RNT was seen to much more fluidly support the required change of direction—that is, without requiring release or separate reorientation of the clue, the same passing gesture was used to complete the pass in these cases. The separate mechanism did not support fluid change of direction, but rather required release and reacquisition of the clue.

8. DISCUSSION

Our study demonstrates that RNT is superior to the separate mechanism for translation and rotation tasks on tabletop displays. RNT was superior or comparable to the separate mechanism on every measure we considered, both in terms of efficiency and accuracy. In addition, users overwhelmingly preferred RNT for Tasks 2 and 3, consistently rating it higher on key variables such as learnability, ease of use, and task suitability.

At a higher-level, the video analysis revealed that imprecision in collaborative passing tasks is both familiar and acceptable to collaborators, and such imprecision was better supported by RNT than by the separate mechanism. This analysis also revealed that the "pass then rotate" strategy of the separate mechanism may be somewhat awkward for the task of passing documents, and that RNT fluidly supports changing the direction of a passed object.

A slight blemish on RNT's overwhelming user support was a slim majority preferring the separate mechanism for Task 1. For this task, participants reported some difficulty using RNT to make slight rotational adjustments in-place. To address these concerns and to further increase the technique's utility, it may be desirable to augment the mechanism with a "rotate-only" region, similar to the four corner circles used in the separate mechanism. Augmenting RNT in this way would allow integral control of rotation and translation to be maintained for document passing and other less-precise tasks, while allowing separate control of rotation and translation for precise adjustments as needed.

With these modifications in mind, RNT has clear potential as a general rotation mechanism for tabletop displays. RNT supports the communicative role of orientation, as embodied in the task of document passing and reorienting, to a greater extent than the separate mechanism. Equally importantly, RNT more fluidly supports the comprehensive role of orientation by allowing more efficient rotation of tabletop objects to any desired orientation. In more fully supporting the communicative and comprehensive roles of orientation, RNT does not sacrifice support for the coordinative role, which requires only that a mechanism maintain the position and orientation of an object after it has been rotated.

9. CONCLUSION

The main contribution of this paper is the presentation and evaluation of RNT, a tabletop rotation mechanism that provides fluid, integral control of rotation and translation using a single contact-point for input. An empirical evaluation demonstrated that RNT more fully supports the collaborative roles of orientation than the standard "corner to rotate" mechanism, which provides separate control of an object's position and orientation. RNT was more efficient and it was highly preferred for the task of document passing—an intentional communicative gesture important for tabletop collaboration. RNT was also superior for a precise targeting task, a result that has positive implications for supporting the comprehensive role of orientation as well as the mechanism's general utility for tabletop interaction.

Furthermore, RNT requires only a single contact-point for input, thus making it technology-independent—an important characteristic given the current state of tabletop research. RNT also allows an object to be freely rotated to any desired angle in a lightweight fashion, needing only a single motion to produce rotational changes. Such changes in orientation are immediately visible, thereby providing clear feedthrough of rotation actions.

While the roles of orientation and design criteria derived from them suggest RNT may be useful as a general tabletop rotation mechanism, it may also be important to consider the mechanism in the context of specific tabletop applications. To this end, future investigations are planned in which RNT is integrated into tabletop applications that support specific collaborative tabletop tasks, such as image sharing and organization, and urban planning and design.

10. ACKNOWLEDGMENTS

This research was supported in part by NSERC, iCORE, Alberta Ingenuity and Intel Inc. We thank John Light and members of the Interactions Lab at the University of Calgary for their insightful comments.

11. REFERENCES

[1] Balakrishnan, R., Baudel, T., Kurtenbach, G. and Fitzmaurice, G. The Rockin'Mouse: integral 3d manipulation on a plane. In *Proceedings of CHI '97* (Atlanta, GA, March 1997), ACM Press, 311-318.

[2] Buxton, W. and Myers, B. A study in two-handed input. In *Proceedings of CHI '86* (Boston, MA, April 1986), ACM Press, 321-326.

[3] Deitz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. In *Proceedings of UIST '01* (Orlando, FL, November 2001), ACM Press, 219-226.

[4] Fitzmaurice, G. W. *Graspable User Interfaces*. Ph.D. Thesis, University of Toronto, Toronto, ON, 1996.

[5] Hinckley, K., Sinclair, M., Hanson, E., Szeliski, R. and Conway, M. The VideoMouse: a camera-based multi-degree-of-freedom input device. In *Proceedings of UIST '99* (Asheville, NC, November 1999), ACM Press, 103-112.

[6] Jacob, R., Sibert, L., McFarlane, D. and Preston Mullen, M. Integrality and separability of input devices. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 1, 1 (March 1994), 3-26.

[7] Kruger, R., Carpendale, S., Scott, S. and Greenberg, S. How people use orientation on tables: comprehension, coordination and communication. In *Proceedings of GROUP '03* (Sanibel Island, FL, November 2003), ACM Press, 369-378.

[8] MacKenzie, I., Soukoreff, R. and Pat, C. A two-ball mouse affords three degrees of freedom. In *Extended Abstract of CHI '97* (Atlanta, GA, March 1997), ACM Press, 303-304.

[9] Magerkurth, C., Stenzel, R. and Prante, T. STARS – a ubiquitous computing platform for computer augmented tabletop games. In *Extended Abstract of UbiComp '03* (Seattle, WA, October 2003), Springer, 267-268.

[10] Mitchell, G. D. *Orientation on Tabletop Displays*. M.Sc. Thesis, Simon Fraser University, Burnaby, BC, 2003.

[11] Rekimoto, J. and Saitoh, M. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *Proceedings of CHI '99* (Pittsburgh, PA, May 1999), ACM Press, 378-385.

[12] Scott, S., Grant, K. and Mandryk, R. System guidelines for co-located, collaborative work on a tabletop display. In *Proceedings of ECSCW '03* (Helsinki, Finland, September 2003), Kluwer Academic Publishers, 159-178.

[13] Shen, C., Lesh, N., Vernier, F., Forlines, C. and Frost, J. Sharing and building digital group histories. In *Proceedings of CSCW '02* (New Orleans, LA, November 2002), ACM Press, 324-333.

[14] Streitz, N., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P. and Steinmetz, R. i-Land: an interactive landscape for creativity and innovation. In *Proceedings of CHI '99* (Pittsburgh, PA, May 1999), ACM Press, 120-127.

[15] Tandler, P., Prante, T., Müller-Tomfelde, C., Streitz, N. and Steinmetz, R. ConnecTables: dynamic coupling of displays for the flexible creation of shared workspaces. In *Proceedings of UIST '01* (Orlando, FL, November 2001), ACM Press, 11-20.

[16] Tang, J. Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34, 2 (February 1991), 143-160.

[17] Wang, Y., MacKenzie, C., Summers, V. and Booth, K. The structure of object transportation and orientation in human-computer interaction. In *Proceedings of CHI '98* (Los Angeles, CA, April 1998), ACM Press, 312-319.

[18] Wu, M. and Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. In *Proceedings of UIST '03* (Vancouver, BC, November 2003), ACM Press, 193-202.