

Fluid Integration of Rotation and Translation

Russell Kruger, Sheelagh Carpendale, Stacey D. Scott, Anthony Tang

Department of Computer Science

University of Calgary

Calgary, Alberta, Canada T2N 1N4

{krugerj, sheelagh, sdscott, tonyt}@cpsc.ucalgary.ca

ABSTRACT

Previous research has shown that rotation and orientation of items plays three major roles during collaboration: comprehension, coordination and communication. Based on these roles of orientation and advice from kinesiology research, we have designed the Rotate’N Translate (RNT) interaction mechanism, which provides integrated control of rotation and translation using only a single touch-point for input. We present an empirical evaluation comparing RNT to a common 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 comprehension, coordination and communication roles of orientation.

Author Keywords

Fluid interactions, rotation, translation, orientation, roles of orientation, tabletop collaboration, communicative gestures

ACM Classification Keywords

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

INTRODUCTION

Many different types of software require objects to be translated and rotated. In interfaces, these actions have traditionally been considered as two distinct operations. If one wants to move and rotate an object, one does so sequentially. Yet kinesthetic studies have demonstrated that rotating and translating are inseparable actions in the physical world [25] and that integrated interaction techniques are more appropriate mechanisms for integrated actions [9]. These studies suggest advantages for interaction methods that integrate rotation and translation.

Recently, there has been an increasing interest in tabletop

displays to support small groups collaborating in co-located settings. This interest has arisen from the importance of tables in physical world collaborative activities, coupled with the fact that many collaborative activities now involve digital information and tools. This interest has been reflected in both the research literature [e.g., 10, 13, 15, 17, 18, 19, 20, 21, 22, 23, 26] and industry response through the design of new types of tabletop display hardware [5]. However, developing appropriate interfaces for tabletop displays still presents many design challenges.

While rotation and translation may be important in general interfaces, they are crucial on tabletop displays. This is in part because tabletop displays are horizontal and in part because tabletop displays lend themselves to collaboration. On vertical displays, all collaborators share the same orientation—that is, the top and bottom of the display are the same for everyone. In contrast, collaborators can be seated at different sides of a tabletop and have different views of the display.

Building on earlier work that investigated how people collaborate on tabletop workspaces [24], our recent observational study [10] revealed that manipulation and re-orientation of items is an extremely commonplace action during traditional tabletop collaboration, and identified three major roles that orientation plays during the course of collaboration:

- **Comprehension:** It is easier to comprehend objects when they are the “right way up.” This role is well recognized.
- **Coordination:** Orientation is used to help establish personal and group territories [19] 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.

Motivated by the fact that for extremely common actions, such as rotation on tabletops, a well-designed interaction mechanism can have a dramatic effect, we developed an integrated rotation and translation technique called Rotate’N Translate (RNT). RNT simultaneously rotates and translates a digital object 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

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

CHI 2005, April 2–7, 2005, Portland, Oregon, USA.

Copyright 2005 ACM 1-58113-998-5/05/0004...\$5.00.

only a single touch-point, RNT supports traditional 2D input such as a mouse or touch, incurs minimal overhead for producing rotation, and is technology-independent. Thus, RNT can be used with both indirect and direct input systems and requires no specialized hardware.

In this paper, we present the design, implementation, and evaluation of a novel interaction technique that combines rotation and translation into a single, fluid interaction mechanism RNT. We begin by presenting related work, followed by our design guidelines, which build on the roles of orientation and kinesthetic principles. Then, the behavior of RNT is described, followed by an evaluation that demonstrates the efficacy and efficiency of RNT. We conclude with a discussion of the study results.

RELATED WORK

Rotation is often clumsy in electronic settings, largely because current input devices (mouse, 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 or system selection of a rotation angle, rotation modes accessed through menus or specific controls, and novel input methods that provide extra degrees of freedom with which rotation can be controlled.

Allowing people to manually rotate their information items is conceptually a direct analog for interactions with traditional media. However, the two degrees of freedom offered by common input devices has led to moded controls that increase the interaction overhead. Examples include the "handle-to-rotate" (Microsoft PowerPoint®, Paint®), and the "corner to rotate" method commonly used for rotating digital objects on tabletop displays [20, 21, 22, 23].

Other rotation schemes use an additional degree of freedom provided by a non-traditional input device to gain the capacity for integrated control of rotation and translation. Examples include providing an additional input device to be used by an individual's non-dominant hand [3], providing multi-touch interaction [5, 26], or using an input device that has three or more degrees of freedom [e.g., 1, 6, 8, 12]. However, use of the non-traditional input devices makes these methods technology-dependent.

Some systems provide the option to automatically orient information items on the tabletop workspace. One approach is to orient items towards the person who has most recently accessed the information. Examples include: InfoTable [17], when a person drags a tabletop item to "their" side of the table, the item automatically rotates toward the closest table edge closest; STARS [13], interaction in the game is turn-based, and items are automatically oriented towards the currently active player; and ConnecTable [23] where passed objects are automatically rotated to be the right-side for the receiver when touched by the receiver. Environment-based strategies typically assume orienting an item to the closest edge of the table is most likely to give

the right person the best view of the item. For example, systems built using DiamondSpin [18, 20, 21] enable automatic re-orientation of digital items towards the outside of the tabletop workspace. Therefore, any item directly in front of a person will always be oriented towards that person, regardless of who is manipulating it. The InteracTable system automatically rotates "tossed" items until they stop on the other side of the table, fully oriented towards the closest table edge [22, 23]. These person-based and environment-based automatic approaches have explored the issues of locating the person manipulating the item and deciding which edge is best. However, even though new technologies can now provide more accurate person-location detection [5, 7]; these strategies assume that the person manipulating the item benefits most from the "best view" of the information. Our earlier work [10] suggests that this assumption may not always be appropriate.

There has also been some exploration of acquiring three-degrees of freedom from two-degree input devices. For example, Chen et al. [4] provides a detailed discussion of variations of moded rotation including 3D rotation options, and Drag [15] allows a person to use a regular mouse to control an object's position while "simulated friction" on the object produces changes in orientation. However, while Drag is conceptually similar to RNT, it was difficult for people to use when evaluated [15]. A more detailed description of the differences between RNT and Drag is included in the discussion.

Note that since our understanding of how orientation contributes to collaboration has only recently been expanded to include coordination and communication [10], the above discussion focuses on the comprehension role of orientation. Some systems provide options of more than one rotation technique. For example Shen et al. [20] offer choices between uniform orientation alignment, manual "corner-to-rotate" options, and an environment-based automatic orientation that begins to address the communication role by offering partial orientation between people at corners of the display.

RNT CONCEPT AND DESIGN GUIDELINES

Our intention is to support all three roles of orientation with one interaction technique. The benefits and requirements for the comprehension role are well understood: items should be oriented correctly for the person viewing them, thereby making them easier to read. The benefits for the coordination role are that collaborators can use orientation to define personal and group territories [19], and to provide non-verbal information about which items are currently in use by whom. To facilitate the coordination role, the full range of freedom for orientation variations must be available for each item and that even minor orientation variations do not get readjusted since they hold information about current use and territoriality [19]. The benefits of the communication role include: initiating a collaborative

communication, conveying support for an idea, and establishing an intended audience. Supporting communication role also requires that small variations in orientation are possible, and that these variations can be achieved casually since communications are best initiated when orientation is imprecise (i.e. when items are oriented between the initiator and the receiver). Summing these up with the advice from kinesthetic studies [25] along with our desire for device independence, we present our design guidelines for an integrated technique.

- **Free rotation:** The technique must support the rotation of an item to any desired angle.
- **Lightweight:** This common interaction goal of minimizing action and cognition overhead will help ensure that the communicative processes afforded by 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. To maintain the long-term communicative effects of orientation, the item's orientation must remain clearly visible after the rotation action has occurred.
- **Integrated:** The interaction technique should integrate rotation with translation since rotating and translating are inseparable actions in the physical world [9, 25].
- **Technology independent:** It is important to support relatively universally available input devices such as those offering 2D, single-contact-point manipulation.

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. While this interaction metaphor comes from the physics of moving an object against friction or through a current, we used a *pseudo-physics* developed and adjusted for interaction ease. The idea of a “current” can be used to illustrate how the rotational behavior of an object relates to its translation. Imagine a current that acts against the object always in direct opposition to the object's movement vector. If the direction of movement changes, so too does the current, maintaining its direct opposition. 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 create a counter-current that is not balanced. In Figure 1(b), contact is initiated in the upper-right and movement is upwards; consequently, the current acts in opposition, placing pressure on the object's left side, and rotating the object counter-clockwise. The object rotates until a balanced relationship with the current is achieved. Once in rotational balance, the object translates at the established

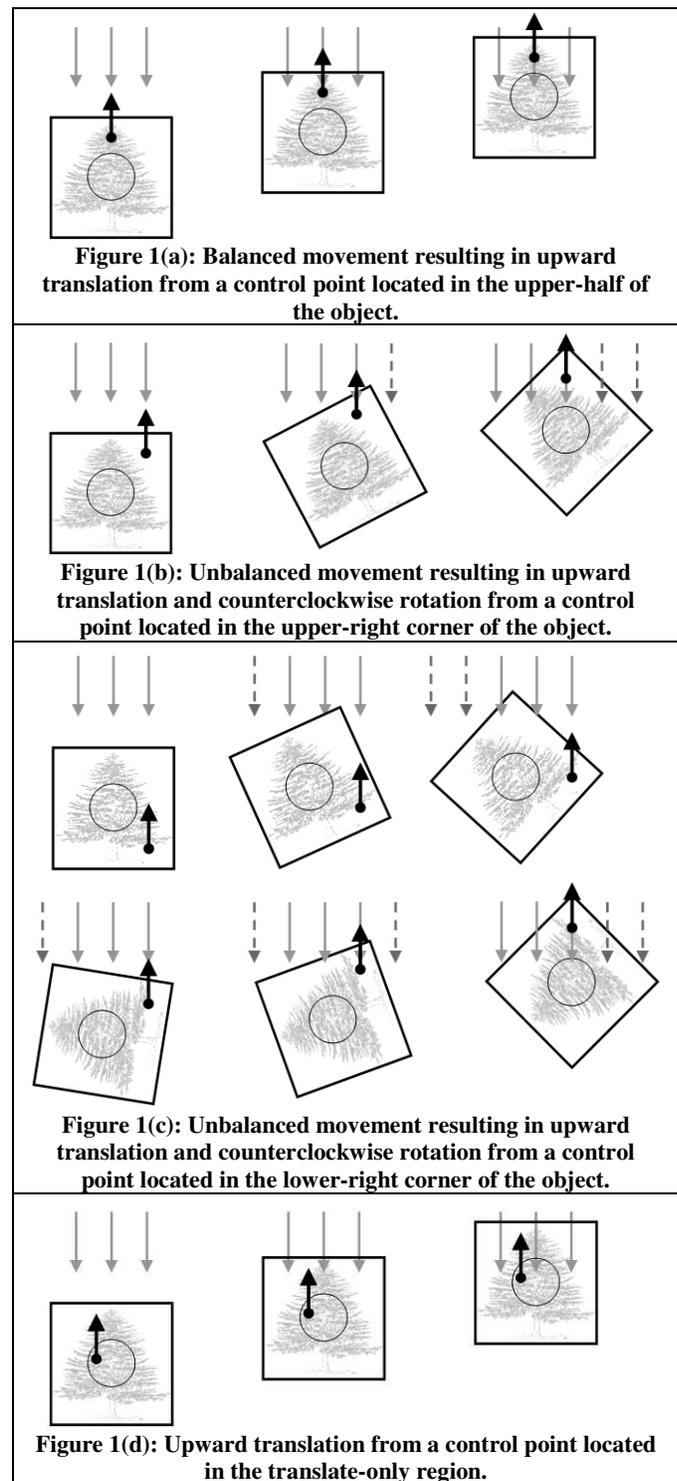


Figure 1: Illustration of integrated rotation and translation.

orientation as long as the movement vector remains unchanged. If the contact-point and movement vector result in an extremely unbalanced situation, rotational changes are more extreme (Figure 1(c)). A greater unbalance causes a greater rotation per unit of movement. All of this implies that the exact center point would provide a translate-only point. People generally seemed to be aware of this, and

through response in pilot studies we expanded the translate-only center point approximately 20% of item's width (25/128 pixels) (Figure 1(d)).

RNT ALGORITHM

Obtaining simultaneous control of an item's position (x and y) and orientation (θ) from 2D (x and y) input proceeds as follows. 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 illustrates 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.

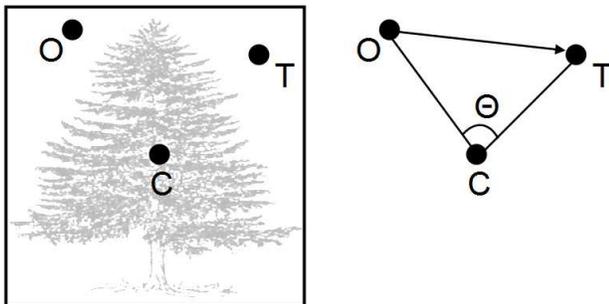


Figure 2: Three required points determine RNT algorithm: C—the Center of the object; O—the Original mouse position; and T—the new touch position (or Target).

EMPIRICAL EVALUATION

To evaluate RNT, we chose a comparison method that could support all three roles of orientation and that did not require special input hardware. This narrowed the possibilities to one of the manual free-rotate moded methods. We choose the traditional-moded "corner to rotate" mechanism (TM) that is currently used in MS PowerPoint® as well as being one of the rotation mechanisms on several tabletop systems [20, 21, 22, 23]. The evaluation starts with a targeting speed and accuracy comparison, continues with a looser targeting or passing activity, and ends with a small group collaboration task. TM behaves as follows: translating an object is simply touch and move; rotating an object involves first selecting the object, touching an indicated corner (12% of object's width 15/128 pixels), and then rotating. No difficulty selecting these handles was either noticed or commented upon. A translated object follows the contact-point vector,

maintaining its existing orientation. A rotating object pivots around its center in location. Figure 3 illustrates TM: in frame (a) the object is stationary, in (b) selected (through a touch), in (c) rotated, and finally after releasing contact in the corner circle, touching and dragging the object is translated in frame (d).

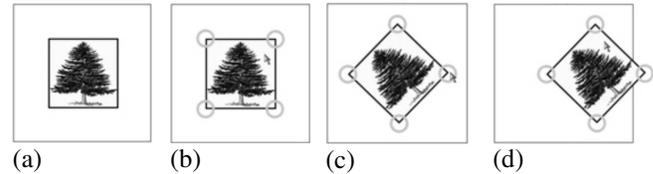


Figure 3: Sequence of images illustrating the "corner to rotate" traditional-moded 'TM' mechanism

Method

Participants

Eighteen paid participants (9 male, 9 female, 15 right-handed) were recruited from the university population. All were proficient computer users, although most were not computer science majors, and only three had used a tabletop display before. Most of the participants rated themselves as beginners with existing digital rotation mechanisms.

Apparatus

The experiment was conducted on a top-projected, touch-sensitive tabletop display powered by a Pentium IV PC running Windows XP. A projector with a resolution of 1280 x 1024 pixels was used to project images onto a horizontal, 72 inch diagonal, SMART DVIT board (Figures 4, 5, 6, 7). Although SMART's DVIT system allows for simultaneous touch detection, the technology does not currently indicate which person initiated which touch. Thus, the collaborative group task was designed so that only one person needed to interact with the table at a given time.

Experimental Design

Participants completed three tasks during the course of the study. 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 was partnered with two other participants for the third task. The order of two rotation mechanisms for each task was counterbalanced.

Tasks

The first task was designed to evaluate the efficiency of the two techniques as a lightweight, free-rotation mechanism as required by the comprehension 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 rotation, and thus is inherently supported by the design of RNT.

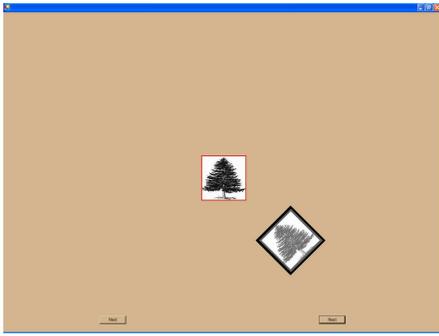


Figure 4: Task 1 involved positioning the central tree image over the target image (thickly outlined)

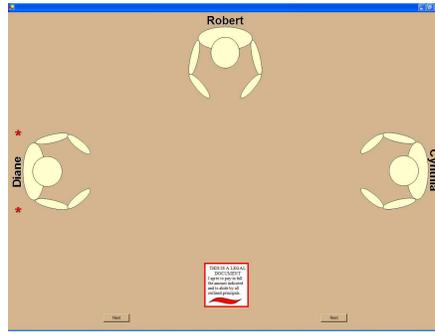


Figure 5: Task 2 involved passing the document image to representations of people seated at different sides of the table

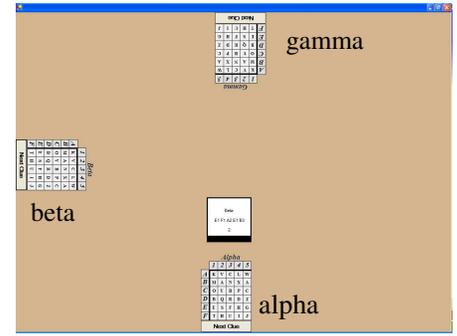


Figure 6: Task 3 involved passing a clue (black edge, thick black bottom edge) to participants at different sides of the table

Task 1: Precision Targeting. This task compared RNT and TM 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 needs to be able to support precise movements to address comprehension needs as well as to support other orientation adjustment activity. In this task, users were required to reposition and reorient a small image of a tree (5.1x5.1 inches, 128x128 pixels), which initially appeared upright and in the centre of the workspace, to a new location and orientation (Figures 4 and 7(a)). The target, the same size as the tree image and outlined with a 10 pixel border, appeared at one of eight different positions circularly located around the middle object at a distance of 275 pixels (10.9 inches). The target's orientation was one of sixteen (0° , 22.5° , 45° , ..., 292.5° , 315° , and 337.5°) possible orientations.

Participants were instructed to, as quickly as possible, position the object atop the target, matching the target orientation. Participants completed 32 trials for Task 1, with the target appearing 4 times at each of the 8 different target locations, and twice at each of the 16 different target orientations. The ordering and pairing of the locations and orientations was random. Participants advanced to the next trial by touching a 'Next' button at the bottom of the screen.

The comparison metrics for Task 1 included time, accuracy, number of touches, touch distance (captured in a log file), and 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 is the total length of all paths touched by the user during the trial (measured in pixels). Each participant was asked to rate the techniques on learnability, ease of use, enjoyment of use, and suitability of the technique for the task, and to indicate their overall preference for one of the techniques for the task performed.

Task 2: Document Passing. This task compared RNT and TM for a less-precise rotation and translation task that attempted to mirror a real-world collaborative activity—the passing of documents. During the task, participants were asked to imagine themselves as legal assistants for a law

firm. Their job was to pass documents to lawyers positioned around the table in such a way that the receiving lawyer could read the document. For each trial, the participants passed a document, which initially appeared upright at the bottom of the screen, to digital avatars located at three sides of the display (Figures 5 and 7(b)).

Participants were free to decide which orientations satisfied the readability condition and the appropriate proximity of the passed document to the lawyer avatars. 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 a little farther, approximately 630 pixels (25.0 inches). The document was to be passed as quickly as possible to the lawyer indicated by stars next to their avatar (see Figure 5). Pressing a "Next" button advanced them to the next trial. Task 2 had 30 trials, with 10 document passes to each of the three lawyers. The ordering of these passes was random.

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 to pass to the general vicinity of a lawyer avatar 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. This task compared RNT and TM for a document-passing task in a collaborative 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 receiver to whom the clue was to be passed. Upon receiving the clue, the receiver would copy

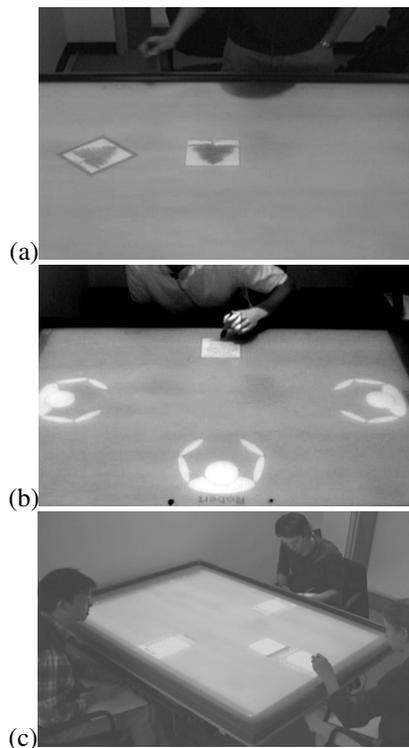


Figure 7: The 3 tasks in progress: (a) Task 1; (b) Task 2; (c) Task 3

down letter-number pairs (such as “E1 F1 A2 E1 B3”) contained on the clue using pen and paper on a clipboard. Double-touching the ‘Next Clue’ button at the bottom of their decoder lookup table dismissed the current clue and caused the next clue to appear at a location randomized independently of the location of the dismissed clue.

When not passing or receiving clues, participants were instructed to decode the clues (Figure 7(c)). 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 were as follows: Alpha to Beta and Beta to Gamma – 615 pixels (24.4 inches); and 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. 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. We recommended to participants that they work on decoding and unscrambling their own clues when not passing or receiving clues. Thus, the participants would be busy when

clues were passed to them, allowing us to compare the feedthrough generated by the techniques.

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. The task was deemed complete when a member of the group read the decoded sentence out loud.

This task was set in a collaborative activity context to provide an opportunity for higher-level evaluation of the techniques. This collaborative task was also intended to provide 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, video data was collected from each of the six groups and later analyzed to aid in a high-level evaluation of the techniques. The time to pass clues was not measured due to the imprecision inherent in measuring time for this task.

Procedure

The participants filled out a consent form, a pre-session questionnaire, and were given an introduction to the tabletop display. The experimenter then explained Task 1 and each rotation technique separately. Thirty-two practice trials were performed followed by 32 task trials for each rotation technique. After both sets were 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. Participants completed 5 practice trials before doing 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.

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 organized by task.

Task 1: Precision Targeting

Results show that RNT is faster and more efficient than TM for precision targeting, and is equivalently accurate (Table 1). The results of a repeated-measures ANOVA showed that RNT took significantly less time to complete the precision targeting task ($F(1,17)=11.967, p<0.005$) and was equivalently accurate (location offset $p=0.703, ns$; and rotation offset $p=0.866, ns$). The ANOVA also showed that RNT required significantly fewer touches ($F(1,17)=76.754, p<0.001$) and a significantly shorter touch distance ($F(1,17)=17.935, p<0.005$).

Participants were asked to rate each technique on a Likert-type scale of 1 to 5 (1=very difficult/not at all and 5=very easy/very much) for ease of learning, ease of use,

enjoyment level, and task suitability. The participants were split as to their overall preferred technique, with 8 users preferring RNT, and 10 users preferring TM. This difference was not statistically significant.

	Measure				
	Time (sec)	Location offset (pixels)	Rotation offset (degrees)	Touches	Touch distance (pixels)
TM	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)

Table 1: Result means and standard deviations for Task 1

Task 2: Document Passing

RNT outperformed TM on all measures for Task 2 (see Table 2). A two-way repeated measures ANOVA showed that RNT produced quicker completion times than TM ($F(1,17)=97.999$, $p<0.001$). It also showed a significant difference for completion times for target positions ($F(2,34)=6.475$, $p<0.005$) with no interaction between rotation technique and target position. Post-hoc analysis showed that passing the document to the right lawyer was significantly faster than to the middle lawyer for both rotation techniques. This is likely due to the longer distance to the middle lawyer's avatar and to the proximity of the right lawyer to the participant's right hand, given that the majority of our participants were right-handed.

Participants were more efficient at document passing using RNT compared to using TM. On average, participants used almost three times as many touches with TM. A one-way repeated-measures ANOVA showed that RNT needed significantly fewer touches ($F(1,17)=224.253$, $p<0.001$) and a shorter touch distance than TM.

A two-way repeated measures ANOVA showed that RNT required a shorter touch distance ($F(1,17)=295.184$, $p<0.001$). It also showed a significant effect of target position on touch distance ($F(2,34)=72.321$, $p<0.001$), with no interaction between these variables. Post-hoc analysis showed that the touch distance was significantly longer

Technique	Measure		
	Time (seconds)	Touches	Touch Distance (pixels)
TM (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)

Table 2: Result means and standard deviations for Task 2

when passing the document to the middle lawyer for both rotation techniques. This result is likely because the middle lawyer is farther from the starting position.

For this task, 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.

Participants' comments clearly reflected this preference. Some examples include: "Both are easy to learn. But the simultaneous one [RNT] is very easy to use, very quick and direct"; "I couldn't wait for the separate rotation and translation [TM] to end!"; "[RNT] More like real life situation and method of doing task"; "The simultaneous method [RNT] was easier to use as it required less touches. Felt like a more natural movement -> more like how a paper is really passed -> 1 motion not 2 motions"; "[TM] Tired your arm out. More, more effort was needed to complete task"; "[RNT] Easier and faster, less energy needed to complete task".

Task 3: Collaborative Document Passing

Results from Task 3 further reinforce the effectiveness of RNT compared to TM. RNT required fewer touches and a shorter touch distance (Table 3). A one-way repeated measures ANOVA showed that RNT required significantly fewer touches ($F(1,5)=79.489$, $p<0.001$), and a shorter touch distance ($F(1,5)=29.972$, $p<0.005$) than TM.

Participant ratings for the two techniques were similar to those for Task 2, showing a strong preference for RNT. 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 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. Participants' comments were similar in flavour to Task 2 comments. Examples include: "Combined rotation [RNT] definitely was easier to use – it did not distract me as much when it was my turn to pass the clue"; "[TM] Took more time"; "[TM] Could be a more fluid movement"; "[RNT] Just one movement, moved faster"; "[RNT] Felt very natural"; and "[TM] Seemed a bit more difficult for some group members to master".

Video Analysis

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

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

Table 3: Result means and standard deviations for Task 3

to the categories listed below. Where appropriate, we consulted the logfile data to clarify specific rotation and translation information for this analysis.

- **Passing orientation and proximity:** Did participants pass the clues so that the receiver could view them “right way up?” How accurately (orientation and proximity) did the passer leave the clue for the receiver? Did inaccuracy affect the collaboration in any noticeable way? Were adjustments made by the person receiving the clue?
- **Receiving clues:** At what point did the clue receiver notice the clue was intended for them? At what point during the passing motion were the clues copied down?
- **Order of rotation:** For TM, 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 the two 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?

Passing Orientation and Proximity

The participants generally felt comfortable orienting the clues. In all cases, and without instruction from the experimenter, participants passed the clue so that its orientation was “right way up” for the receiver. Furthermore, very few adjustments were made after the clue was passed. Most (17 of 216) of these were performed by the passer of the clue. Only 3 of the 216 total were preformed by the receiver. All post-pass adjustments were minor in nature (i.e. less than 10 degree rotation and less than 2 inch translation).

Though clues were passed “right way up” for the receiver, the final orientations were imprecise with respect to the tabletop edge. The average final angle offset was 9.59° for TM and 13.25° for RNT. Note that in usage this difference of 3.6° is only 0.01% of 360° and is barely noticeable. In fact many (42 of 216) of the offset angles for either TM or RNT were more than +/-20°. Similarly, proximity was also casual. However, this imprecision did not appear to affect either the collaboration or the task.

Receiving 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. Perhaps due to the pressure of the game, clues were often acknowledged when they appeared, as well as during the pass. This applied to both RNT and TM. Results of the video analysis indicate no clear difference between the techniques as to when a receiver noticed the passing gesture. We believe the act of a collaborator’s arm moving towards them was sufficient to attract their attention regardless of the technique being used.

Order of Rotation

For TM, 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 to rotate the clue in front of them; the second approach was to first rotate the clue, and then to pass it to the receiver. The majority of passes were done according to the second strategy (78 of 108 passes, or 72.2%). The first strategy, which involved rotating directly in front of the receiver, produced noticeable awkwardness. 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. The rotation of the clue by the passer in front of the receiver appeared to be awkward both in terms of reach and in terms of imposing on the receiver’s personal space. Since the “move then rotate” strategy comprised over a quarter of all passes using TM, TM appears to generate more socially awkward situations.

For RNT, it was possible to use only the integrated technique to perform a pass, as well as to use both the translate-only region and the integrated technique to perform a pass. However, the video analysis revealed that participants rarely used the translate-only feature when passing (4 of 108 passes, or 3.7%). Thus, it appears participants were willing and able to use the integrated RNT mechanism as it was intended.

Issues/Problems

Participants had no major problems with either rotation technique. For TM, 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 rotation or over-rotation was a sporadic issue, also occurring for 8 of the 108 trials (7.4%). In these cases, the person attempted to pass the clue too quickly, and as a result, applied too much “force” (and thus rotation) to the clue.

Another interesting issue 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 TM, and 4 for RNT). Of particular interest here is that RNT was able to more fluidly support the required change of direction—that is, the same passing gesture could be used to complete the pass in these cases. In contrast, TM typically required release and reacquisition of the clue, especially when the “rotate first” passing strategy mentioned above was being used.

Discussion

For translation and rotation tasks, our study demonstrates that RNT is superior or comparable to TM on every measure we considered. Since TM requires one more touch than RNT to operate, the number of touches measure (average number of touches TM/RNT for Task 1: 3.67/2.09; Task 2: 2.9/1.2; Task 3: 5.2/3.0) may not provide a clear

comparison. However, RNT outperforms on time, and touch distance. In addition, users overwhelmingly preferred RNT for Tasks 2 and 3, consistently rating it higher on the key variables—learnability, ease of use, and task suitability.

The equivocal preferences for the precision targeting task (Task 1) are worth discussing. For this task, participants reported some difficulty using RNT to make slight rotational adjustments in-place. The fact that they perceived their accuracy to be less is important, even though the accuracies were comparable. Tasks 2 and 3 indicate that RNT encourages casual gestures, which are not commonly associated with accuracy. There is a possibility that with familiarity and feedback, a user might gain confidence in RNT's accuracy.

One possibility is to adjust RNT to more fully allow separable interaction. Currently, rotation mechanisms such as TM require separate interactions. RNT requires rotation to be integrated with translation and allows translation to be performed separately. A simple extension would be to provide integrated interaction and also allow separable interaction for both rotation and translation. Incorporating a rotation only option—without interfering with the general integrated interaction—would be worth evaluating.

RNT has clear potential as a general rotation mechanism. RNT is faster, more efficient, and as accurate as TM. It also supports the communication role of orientation, as embodied in the task of document passing and reorienting, to a greater extent than TM. Equally important, RNT more fluidly supports the comprehension role of orientation by allowing more efficient rotation of tabletop objects to any desired orientation. In more fully supporting the communication and comprehension roles of orientation, RNT does not sacrifice support for the coordination role, which requires that a mechanism maintain the position and orientation of an object after it has been rotated.

At a higher-level, the video analysis revealed the casual nature of collaborative passing tasks is both familiar and acceptable to collaborators, and such interactions were better supported by RNT than by TM. This analysis also revealed that the “pass then rotate” strategy sometimes used with TM may be somewhat awkward socially for the task of passing documents. This is because the rotation took place in another person's space and caused hesitation for both participants. RNT fluidly supports a change of direction during the passing of an object. With TM, such a direction change can involve stopping and re-starting, whereas with RNT, one simply changes direction while proceeding.

The RNT interaction technique has been incorporated into two larger software projects. One is software designed to support a magazine layout task. The other has been designed to support the organizational and scheduling tasks associated with agile software development [2] (Figure 9). While the items in the study were all squares, note that in Figure 9 RNT has been applied to rectangular items. RNT seems applicable to at all rectangular shapes and perhaps all

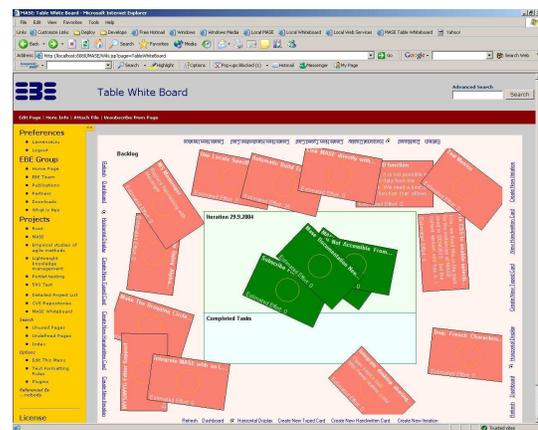


Figure 9: Agile software development application

shapes in general but this is still to be studied. In Figure 9 one can also see the casual multi-orientation organizational structure more common in real world planning activities.

Our RNT mechanism uses a simulated force to integrate rotation and translation. However, our *simulated force* relates more to ideas of *alternate interface physics* [16] and *cartoon physics* [14] than to the more accurate physics of Drag [15]. In adhering to a physics-based model, Drag models *changes* in friction depending on the location of the contact-point according a damping function. As a result of the damping function, less movement is required to produce changes in orientation towards the edges of the object. No such damping function is implemented in RNT. The separate evaluations of these two techniques produced considerably different empirical results. During Drag'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. In our study, users found RNT easier, faster, and just as accurate as traditional-moded rotation and translation interactions.

CONCLUSIONS AND FUTURE WORK

The main contribution of this paper is the introduction and evaluation of RNT, an interaction mechanism that provides fluid, integrated control of rotation and translation using a single contact-point for input. We have empirically shown that RNT is more efficient and effective than the traditional “corner to rotate” mechanism (TM), and provides better support for the communication and the comprehension roles of orientation.

RNT requires only a single contact-point for input makes it technology-independent—an important characteristic for its utility in general, and its general applicability in other interaction environments, such as walls and desktops.

The next stage of this project will be to explore the utility of RNT beyond tabletop environments. Though this research was largely motivated by the need to provide better

orientation support on digital tabletop workspaces, the overwhelming success of RNT during evaluation suggests that RNT may benefit any application involving rotation and translation operations, such as traditional desktop graphical design packages and image editors.

ACKNOWLEDGMENTS

We gratefully acknowledge support from Natural Sciences and Engineering Research Council (NSERC), Informatics Circle of Research Excellence (iCORE), Alberta Ingenuity, and Canadian Foundation for Innovation (CFI). We thank anonymous reviewers and colleagues at Interactions Lab, University of Calgary insightful comments on this work.

REFERENCES

- Balakrishnan, R., Baudel, T., Kurtenbach, G. and Fitzmaurice, G. The Rockin' Mouse: integral 3d manipulation on a plane. In *Proceedings of CHI '97*, ACM Press, (1997), 311-318.
- Beck, K. *Extreme Programming Explained - Embrace Change*, Addison Wesley 2000.
- Buxton, W. and Myers, B. A study in two-handed input. In *Proceedings of CHI '86*, (1986), 321-326.
- Chen, M., Mountford, S.J., and Sellen, A. A study in interactive 3-D rotation using 2-D control devices. In *Proceedings of Computer Graphics and Interactive Techniques*, (1988), 121-129.
- Deitz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. In *Proceedings of UIST '01*, ACM Press, (2001), 219-226.
- Fitzmaurice, G. W. *Graspable User Interfaces*. Ph.D. Thesis, University of Toronto, Toronto, ON, 1996.
- Hancock, M. S., and Booth, K. S. (2004). Improving menu placement strategies for pen input. In *Proceedings of Graphics Interface* (2004), 221-230.
- 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*, ACM Press, (1999), 103-112.
- 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 (1994), 3-26.
- Kruger, R., Carpendale, S., Scott, S. D., Greenberg, S. Roles of orientation in tabletop collaboration: comprehension, coordination and communication. *Journal of Computer Supported Collaborative Work*. (in press).
- MacKenzie, I., S. Fitts' law as a research and design tool in human-computer interaction. *Human Computer Interaction*, 7, (1992), 91-139.
- MacKenzie, I., Soukoreff, R. and Pat, C. A two-ball mouse affords three degrees of freedom. In *Extended Abstract of CHI '97* ACM Press, (1997), 303-304.
- Magerkurth, C., Stenzel, R. and Prante, T. STARS – a ubiquitous computing platform for computer augmented tabletop games. In *Extended Abstract of UbiComp '03*, Springer, (2003), 267-268.
- McCloud, S. *Understanding Comics: The invisible art*. Harper Collins, New York, New York, 1993.
- Mitchell, G. D. *Orientation on Tabletop Displays*. M.Sc. Thesis, Simon Fraser University, Burnaby, BC, 2003.
- Perlin, K., and Fox, D. Pad: An alternative approach to the computer interface. In *SIGGRAPH '93 Proceedings*, ACM Press, (1993), 57-64.
- Rekimoto, J. and Saitoh, M. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *Proceedings of CHI '99*, ACM Press, (1999), 378-385.
- Ringel, M., Ryall, K., Shen, C., Forlines, C. and Vernier, F. Release, relocate, reorient, resize: fluid techniques for document sharing on multi-user interactive tables. In *Extended Abstracts of CHI'04*, ACM Press, (2004), 1441-1444.
- Scott, S.D., Carpendale, M.S.T, & Inkpen, K.M. Territoriality in collaborative tabletop workspaces. In *Proceedings of CSCW'04*, (2004), 294-303.
- Shen, C, Vernier, F., Forlines, C. and Ringel, M. DiamondSpin: An extensible toolkit for around-the-table interaction. In *Proceedings of CHI'04*, ACM Press, (2004), 167-174.
- Shen, C., Lesh, N., Vernier, F., Forlines, C. and Frost, J. Sharing and building digital group histories. In *Proceedings of CSCW '02*, (2002), 324-333.
- 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*, ACM Press, (1999), 120-127.
- Tandler, P., Prante, T., Müller-Tomfelde, C., Streitz, N. and Steinmetz, R. ConnecTables: dynamic coupling of displays for the flexible creation of shared workspaces. *Proceedings of UIST '01*, ACM Press, (2001), 11-20.
- Tang, J. Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34, 2 (1991), 143-160.
- 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*, ACM Press, (1998), 312-319.
- Wu, M., Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. In *Proceedings of UIST '03*, ACM Press (2003), 193-202.