

Under the Table Interaction

Daniel Wigdor^{1,2}, Darren Leigh¹, Clifton Forlines¹,
Samuel Shipman¹, John Barnwell¹, Ravin Balakrishnan², Chia Shen¹

¹Mitsubishi Electric Research Labs
201 Broadway, 8th Floor
Cambridge, MA, 02139, USA

{leigh | forlines | shipman | barnwell | shen}@merl.com

²Department of Computer Science
University of Toronto
Toronto, ON, Canada

{dwigdor | ravin}@dgp.toronto.edu

ABSTRACT

We explore the design space of a two-sided interactive touch table, designed to receive touch input from both the top and bottom surfaces of the table. By combining two registered touch surfaces, we are able to offer a new dimension of input for co-located collaborative groupware. This design accomplishes the goal of increasing the relative size of the input area of a touch table while maintaining its direct-touch input paradigm. We describe the interaction properties of this two-sided touch table, report the results of a controlled experiment examining the precision of user touches to the underside of the table, and a series of application scenarios we developed for use on inverted and two-sided tables. Finally, we present a list of design recommendations based on our experiences and observations with inverted and two-sided tables.

ACM Classification: H.5.2 [User Interfaces]: Input devices and strategies, Interaction styles.

General terms: Design, Human Factors

Keywords: Tabletop interaction, direct touch interaction, touch screen, bimanual input, two-sided, inverted input

INTRODUCTION

Interactive tables have been the focus of a great deal of recent research. In particular, interactive tables offer a compelling platform for shared display groupware, allowing multiple users to simultaneously interact with applications using a direct-touch paradigm. The advantages of tabletops for co-located collaborative groupware have been well explored in the literature [7, 13, 14, 17, 26, 27]. The advantages of using a direct-touch interface have also been investigated [21], including for shared display groupware on tabletops [5].

In this paper, we present an extension to the familiar tabletop interaction paradigm: *under the table* interaction. We have constructed an interactive touch table, designed to receive touch input from both its top and bottom (Figure 1).

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.

UIST'06, October 15–18, 2006, Montreux, Switzerland.

Copyright 2006 ACM 1-59593-313-1/06/0010...\$5.00.

By combining two registered touch surfaces, one facing downward from the underside of the table, the other in the familiar tabletop position, we are able to offer a new dimension of input for co-located collaborative groupware. This design accomplishes a goal that might previously have been thought impossible: increasing the relative size of the input area of a touch table while maintaining its direct-touch input paradigm.

We will explore the design space of two compelling under-the-table designs: an inverted table, which receives input only from the bottom; and a two-sided table, capable of receiving input from both surfaces. Each design affords the development of unique interactions in several ways, including:

- Elimination of occlusion while touching.
- Bimanual input with a natural oppositional posture.
- Concurrent bimanual or two user input to a single area.
- Side separation offering a naturally moded input-space.
- Private input to a shared direct-touch display.
- Dual sides afford a sandwiched 3D virtual space.

We begin by reviewing related work, followed by a brief discussion of our hardware design. Next, we present a suite of new interaction schemes enabled by both an inverted and a two-sided touch table; a study examining the precision of users making touches to the underside of the table; and a series of application scenarios developed for use on inverted and two-sided tables. We conclude by presenting a set of design recommendations derived from the results of the user study and experiences with the usage scenarios.




Figure 1. A two-sided touch table extends the input capabilities of traditional direct-touch surfaces.

RELATED WORK

Recent technological advances in direct touch surfaces have enabled the construction of interactive tables that can be realistically expected to be productized and to have a persistent presence in our daily lives. SmartTech's DViT [24] is a vision-based technology capable of detecting two points of touch, and is widely commercially available as electronic whiteboards. DiamondTouch [5], also commercially available, is a multi-touch technology with the added benefit of user-identification and debris tolerance. It is based on capacitive coupling with arrays of horizontal and vertical transmitters. SmartSkin [16] from SONY's CSL is a research prototype employing a fully-pixelated multi-touch technology, although without user identification capabilities. Frustrated Total Internal Reflectance (FTIR) is another existing multi-touch sensing technology that has received renewed interest recently [9].

These multi-touch technologies have brought forth numerous research efforts studying interactive tables, including user interface and interaction techniques, multi-finger and two-handed gestural interactions, and computer supported collaboration [7, 13, 14, 17, 20, 26, 24, 27, 28]. A few research projects have started to extend the affordances of a direct touch tabletop by varying the viewing angles from four sides [12], or augmenting a table with 3D head-tracked, head-worn displays and 3D-tracked gloves [3]. The research presented in this paper explores the new affordances when a table is made touch sensitive from both the top and the underside surface.

Adding an input surface to the back of a touch-sensitive screen has also been explored by Sugimoto and Hiroki [25]. Their *Hybrid Touch* device is a PDA with a touchpad fastened to the rear of the device. Although similar physically, their device is used only as a relative input pad with variable C/D gain. In the present work, we hope to more broadly explore the potential of two-sided touch input devices.

OUR PROTOTYPE

In terms of input bandwidth, an input surface on the underside of a table offers no more input than does positioning additional touch-sensitive input devices elsewhere within arms-reach. The advantage of the positioning of this input surface under the table lies in the extension of the direct-touch input paradigm to a second input area. By carefully positioning the bottom input area below the tabletop input area, the two surfaces can be calibrated together with an overhead projector providing top-projection of the display area. In order to maintain the sense of direct-touch input for a user while touching the underside of the table, it is important to minimize the thickness of the table, so that the actual point of touch is as close to the displayed target as possible.

To build a two-sided touch table, we removed the antenna arrays from two DiamondTouch-88 input surfaces [5], and mounted them in opposition to one another on thin sheets of Lucite plastic. The device was built with a single controller to ensure timing of input data.

We made the table as thin as possible to minimize the distance between the surfaces: the overall thickness is approximately 1.12cm. To create a direct-touch input paradigm, the computer display was projected on to the top surface of the table, and calibrated in software. The same calibration was inverted and used for the bottom of the table to ensure perfect registration of the input surfaces.

PROPERTIES OF UNDER-THE-TABLE DEVICES

Two-sided touch tables present a number of interesting properties that offer new opportunities and challenges to interaction designers. In order to fully explore these properties, we have divided the discussion into two sections. First, we explore the use of an 'inverted touch table', a table with only one input surface, on the bottom of the device. Second, we present a discussion of a fully-enabled two-sided touch table.

Properties of an Inverted Touch Table

An inverted touch table has a single touch surface precisely registered with a display, creating a direct-touch interface. The distinguishing characteristic of an inverted table is that the input area is on the *bottom* of the table, while the display remains on the top. Though intrinsically interesting, this discussion is meant to form part of the larger, two-sided table motivation. It is important to remember that all of the properties discussed here continue to be true of a two-sided touch table.

Direct-Touch Paradigm Maintained

Although the touch and display surfaces are now separated, by ensuring exact registration between the input and display surfaces, an inverted touch table is able to maintain many properties of a direct-touch interface. Figure 2 illustrates multiple touch points on the bottom of the table corresponding with visual elements on the top. Later in the paper, we present the results of a study showing the accuracy of touching targets on an inverted touch table.




Figure 2. A direct-touch input paradigm can be maintained on an inverted touch table. Note: our prototype is not transparent, the hand below the surface is shown for illustrative purposes.

No Occlusion of On-Screen Objects

When interacting with a direct-touch interface, occlusion of the display device is unavoidable: specifying a point on screen requires the user to touch it. By moving the input space to below the display space, we are able to eliminate this occlusion. This is desirable both for individuals working with intricate data and groups where one user may wish to observe displayed imagery currently under manipulation by another user.

Accuracy of Pointing

With traditional tabletop interfaces, touching using a finger is difficult to do with a high degree of precision. Although input devices are typically capable of pixel-level precision, the precise pixel being targeted by a touch is hard for the user to determine or control, since multiple pixels are typically within the bounds of the touch area [18]. It is impossible to offer in-place feedback to the user during the touch, since the selection point is occluded by the finger.

Potter et al. [15] compared three approaches for selection on a touchscreen: *land-on*, *first-contact*, and *take-off*. A land-on strategy selects the object immediately below the finger at the initial point of contact. First-contact selection, similar to *crossing* selection [1], selects the first-on screen object the user's finger touches as they drag around the screen. Finally, take-off selection is done by selecting the object that was last touched before the finger was removed from the screen. Their take-off strategy varied from the others in that its target of influence was not the point of contact of the finger, but rather a cross-hair offset approximately 13mm above the tip of the finger. Potter et al.'s study found that this technique reduced errors. The authors noted, however, that several participants 'took great exception' to the cursor offset, and that selection time was higher than for the other techniques. Users seem to expect that touches will be interpreted in a direct-touch manner. This is likely why the take-off technique has not been widely adopted. An inverted touch table allows for a take-off strategy to be implemented such that a direct-touch input paradigm is maintained, as shown in Figure 3.




Figure 3. Left: Selecting using Potter et al.'s *take-off* strategy, with cursor offset from the touch point. Right: the same accuracy is possible on an inverted table, *without* breaking the direct touch paradigm.

Given the ability to implement the take-off selection such that the point of contact and the point of influence coincide, we argue that pixel-level selection is actually *superior* using an inverted table.

Privacy of Input

On a shared direct-touch system, all three stages of interaction are public – input target specification, the input action, and the consequential change to the system, are all visible to all users. This can be advantageous in circumstances where knowledge of other users' actions is desirable, such as when performing loosely coupled group tasks. There can be circumstances when users of a shared or public display may have a need to interact with the system privately, such as when inputting moves in a game.

To address this, private output on a shared display has been explored [23 and 29], and for touch tables in particular in [13 and 27]. Absent from these discussions, however, has been a review of how to enable private *input* to a shared direct-touch system. With an inverted touch table, both the input gesture and its target can be kept private from other users, preventing others from inferring the private information with which the user is attempting to interact.

A potential disadvantage to this type of input privacy is the need to avoid visual feedback to the user making the input. This is especially problematic given that users are unable to see their hands while making the input, thus decreasing their confidence in the accuracy of the system's response. Feedback could be given through one of the privacy techniques described by Shoemaker and Inkpen [23], or through a private audio channel [13]. As we report later, our study found that users are generally able to select within a few millimetres of a target while pointing under the table, so the feedback is necessary only to complete and confirm selections, or for highly precise input.

Accidental Touches Less Likely

As reported in [18], users of shared touch tables often make pointing gestures to objects while discussing them with other users, accidentally triggering input to the system. This is less likely when the display and input surfaces are not the same, since pointing on the display for reference ("look at this!") will not be conveyed as input.

Arm Fatigue

A disadvantage of large direct-touch display surfaces is the need to make elaborate arm movements, increasing fatigue. This might be especially problematic for interaction on the bottom of a table, since the table is not available to serve as an arm rest. In our hardware implementation, we installed the touch table 15cm above a second table, so that the lower base table could be used as an arm rest (Figure 13). Participants in our user study reported universally that touching on the top of the table was significantly more fatiguing than on the bottom, since this lower surface served as an arm rest throughout the entirety of the task.

Reduced Reach

Tables are generally built at a height such that touches below the table cannot comfortably extend beyond the area immediately in front of the user. To increase the range of touch, our inverted table was constructed with a very thin bezel, and placed at a height comfortable for touching the bottom (approximately 80cm). Even with these adjustments, the range of comfortable reach was noticeably reduced from that on the table top. For most users, the comfortable range of reach was less than 45 cm.

It is interesting to note that touches to traditional tabletops are often limited to a range even smaller than this when multiple participants are present [17, 19]. Although a number of factors influence the dimensions of this input area, the social factors may be exaggerated when working under the table, further reducing reach beyond the physical limitations.

Change in Bimanual Posture

With traditional touch tables, bimanual interaction is typically done with the hands flat on the table, thumbs pointing towards one another. Bimanual input to an inverted touch table inverts this, so that the thumbs face away from one another (Figure 4). This may have implications for designers of bimanual interaction. In particular, when not facing one another it may be that the hands are less prone to involuntary complementary movement, which may increase the ease of asynchronous bimanual input.




Figure 4. Bimanual input to an inverted touch table: the relative orientation of the hands is reversed.

In addition to all of the properties of inverted tables listed above, a two-sided touch table adds an extra dimension of input. We now explore how making the top of an inverted touch table touch-sensitive changes its characteristics.

Properties of a Two-Sided Touch Table

The addition of a second input area for direct-touch interaction with a tabletop offers several compelling advantages for the development of interactive systems. In this section, we will review each of the advantages that we have identified, and present design sketches and scenario applications that we developed to exploit them. Of course, a two-sided table is capable of supporting all of the same interactions as a regular touch table, but these are already outlined in several previous research efforts. Our focus here is on interactions that would otherwise not be possible without a two-sided table. It is our hope that the broad strokes we lay out provide inspiration for the development of systems that make use of two-sided tables.

More Input Bandwidth

By adding a second touch surface, the bandwidth of the input device is effectively doubled. This is especially true for multi-user input, since more hands can simultaneously work on the larger surface. Leveraging this bandwidth effectively requires the designer to consider the properties that follow.

Number and Table-Side of Hands has Meaning

The side touched and number of hands being used for interaction is of importance for interaction (Figure 5). Input to the top and bottom can have identical, similar, or completely disparate effects. Additionally, the number and location of the hands has significance: a single hand above can have different semantic implications than a single hand below. Two hands above, two below, or one above and one below all afford potential differences in the semantics of interaction.

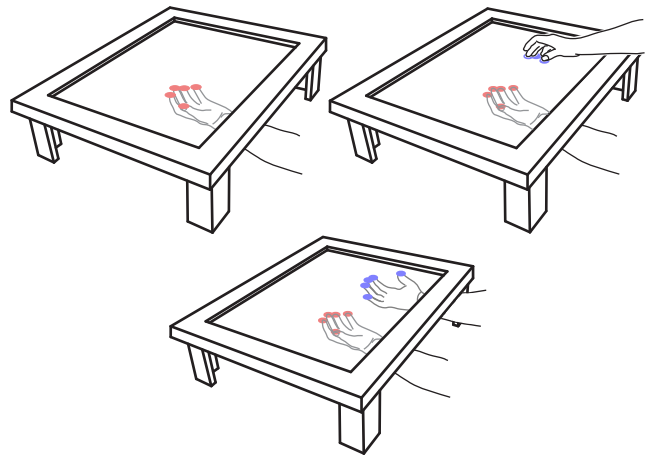


Figure 5. The number of hands and the surface(s) being touched can have differing semantic associations within an application.

Sides Afford Modal Coupling

Bier et al. [4] demonstrated the utility of moding interaction spaces with the use of a click-through toolglass controlled by the non-dominant hand. By tying mode to the side of the table where input is made, the two-sided table affords a similar paradigm. This allows two modes to be maintained continuously, and may reduce errors associated with a moded interface.

Co-locality of Bimanual Interaction

For our purposes, we define *co-locality* as whether or not the hands are operating within the same virtual space. Balakrishnan and Hinckley [2] found that bimanual input to the same area could be physically separated without a performance penalty, so long as visual feedback of the virtual position of the hands was provided. When working with a direct-touch input table, co-locality of the virtual hands necessitates co-locality of the physical hands. Thus, some forms of bimanual input are not possible, since the hands cannot occupy the same space at the same time. Although techniques have been proposed to decouple the input and display spaces [28], they require multiple steps, and are more complex than simply touching the display.

A two-sided touch table allows both hands to effectively target the same location simultaneously, in a way that is not possible without causing physical interference on a regular touch table. In the case of our two-sided tabletop, co-locality is enabled by one hand operating above, and the other below the surface of the table, as shown in Figure 6.




Figure 6. Both hands of a single user can operate in the same virtual area on two-sided touch table.

Co-locality of Interaction for Multiple Users

Not only does a two-sided touch table enable co-locality for two hands of a single user, but it also enables co-locality of touch for multiple users. As we see in Figure 7, positioning hands on opposite surfaces of the table allows two users to occupy the same virtual space at the same time.




Figure 7. Multiple users can operate on the virtual same space concurrently.

New Type of Symmetry in Bimanual Interaction

Traditional touch table interfaces, where both hands are oriented with the palms down, afford a certain kind of symmetry of bimanual interaction. A two-sided touch table affords two new types of bimanual symmetry (Figure 8). The first is *oppositional translational symmetry*, where the hands face one another and move in step. As well as *rotational symmetry*, where the hands are placed atop one another and rotated in opposition. Each of these types of bimanual symmetry affords a different type of interaction.




Figure 8. Two-sided tables afford new types of bimanual symmetry: (left) oppositional translational symmetry, (right) rotational symmetry.

Potential for Reduction of Physical Interference

Two types of physical interference can occur on a direct-touch interface. The first, addressed earlier, occurs when two users wish to interact in the same space. The other occurs when one or more user attempts to place their hands in such a way that would otherwise cause the arms to pass over one another, or actually collide. A two-sided table allows bimanual interaction without this type of interference, if interactions are designed to make appropriate use of the second surface, such as in the bimanual Pong application scenario presented later.

Affordance of a Three-Dimensional Space

For designers, a regular touch table affords a flat interaction and display area, similar to a desktop computer's display. A two-sided touch table, however, presents a mapping of a third dimension to the application. This might be achieved by visualising the 3D space as a fixed volume sandwiched between the input surfaces.

Alternate Input Mapping or Paradigm by Table Side

As we have discussed, a direct-touch input paradigm can be maintained when working on the underside of the table. Designers may find it desirable, however, to use the physical separation of display and input spaces to break this paradigm. As we see in Figure 9, input to the top surface can be interpreted in a direct touch manner, while at the same time interpreting touches to the bottom of the table as input to a separate space, or input to the same space in a non-direct touch manner.




Figure 9. The bottom surface is used for input to another virtual space while maintaining direct-touch on the top surface.

ACCURACY OF UNDER THE TABLE POINTING

Appropriate visual feedback allows designers of under the table software to implement Potter et al.'s take-off technique [15] while maintaining a direct-touch input paradigm. As such, with appropriate feedback, touches to the underside of the table are actually *more* precise than those on a regular touch table. An important question that arises with inverted and two-sided touch tables is the accuracy with which users can point under the table *without* visual feedback. Although intuition tells us that touching under the table will not be as precise as touches above the table, the accuracy penalty is still of interest for two reasons. First, if the underside of the table is to be used for private input without visual feedback, the accuracy of touching will be of concern to designers. Second, the accuracy will give insight as to how effectively the direct-touch input paradigm is extended to the underside of the table.

To determine the effect of side-of table on initial touch accuracy, we conducted a simple experiment, requiring participants to touch targets using either the top or bottom of the table. We hypothesised that pointing accuracy would be significantly better on the top surface, but that, on the bottom surface, the distance from touch point to the target would be small. In essence, we expect that the direct-touch input paradigm will be effectively carried to the bottom surface of the table, albeit at the cost of reduced accuracy.

Given that an inverted table allows for the implementation of the take-off technique, the goal of our study is not to determine accuracy of selection for direct target acquisition per se, but rather to give designers an idea as to how much they can rely on accuracy of touching for other interaction techniques.

Participants

8 participants (3 male and 5 female) between the ages of 23 and 42 were recruited from the local community and our lab. All were right-handed, had completed at least a Bachelor's degree, and had at least some experience working with a direct-touch interface. Participants were given \$20 for their participation, regardless of performance.

Apparatus

Participants sat at our two-sided table, on to which information was projected with a calibrated projector running at a resolution of 1280x1024 pixels with a display area of approximately 51 x 38 cm. They were instructed to select targets by touching the respective surface of the table. To encourage participants to touch the surface as close to the target as they could, our testing system used the land-on technique, described previously, so that only the point of initial touch of the finger on the table would register as a selection.

The point of touch registered by the system was visualized with the display of a black dot. When touching on the top surface, this dot was projected onto the top of the participant's finger, and so was of little use to refine the selection. When working on the bottom, the dot would allow the participant to understand the point of contact.

Task and Procedure

Each participant performed half of the experiment on each side of the table: first on one side, then switched to the other for the duration of the experiment. The order of input surface was balanced between participants.

Selections were divided into pre-computed 'walks' of 5 target selections, which began with a red target presented on the screen. Participants would select the red target to begin the walk, and then select each of the four remaining green targets of the walk in sequence. Each target within a walk was of identical size, and the distance from a selection point to the next target was fixed within the walk. This meant that each selection of a green target within a walk had an identical Fitt's law index of difficulty (ID) [6, 11]. Three target sizes of 10mm (25 pixels), 20mm (50 pixels), and 30mm (75 pixels) in diameter were fully crossed with three distances of 60mm, 120mm, and 180mm, resulting in nine size/distance combinations.

The direction of selection was one of the 8 compass directions (north, north/west, west, etc). Figure 10 illustrates:




Figure 10. Experimental procedure: first target in a set of 5 is red, the following 4 are green. All targets in a set have identical ID's.

Before beginning the experiment, the procedure and apparatus were explained, and participants were allowed to practice the task until they felt comfortable with it, which usually occurred within 10 practice trials.

Participants were required to repeat the selection if they did not successfully select the target the first time. This ensured that they did not try to race through the experiment by touching anywhere. However, our analysis only considered the first touch of each trial as being significant in terms of accuracy measurement. Since our goal was to determine the accuracy of under-the-table touching, we were not particularly interested in the speed of selections.

Design

Each participant completed 75 selections for each of the 9 size/distance combinations, on each side of the table. Although each walk consisted of 5 selections, only the last 4 were included in our analysis, since these 4 all had identical ID's. The purpose of the first, red trial was only to position the participant to begin the walk.

The overall design of the experiment was:

8 participants x
2 sides of the table x
3 target sizes x
3 target distances x
60 valid (green circle) selections =
8640 total target selections

Hypotheses

We hypothesized that land-on selection would be less accurate on the bottom of the table, where users are forced to rely on proprioception alone to determine the position of the hand. We expect, though, that pointing inaccuracy on this surface will be sufficiently small so as to allow the use of a direct-touch input paradigm for interactions. We also expect that both target size and distance between targets will have an effect on accuracy. In addition to the usual effects of these factors [6, 11], an additional effect in this circumstance is the recency of feedback of hand location. When the distance between subsequent targets is small, the amount of movement from the previous position is minimized, minimizing the distance that the hands must travel without feedback. We expect that this will increase the effect of distance on the accuracy of pointing, and hypothesize a corresponding asymmetry of the effect of distance on accuracy across touched-side conditions.

Results

Our primary dependent measure was *first touch error*: the distance in mm from the first touch to the nearest pixel of the target. If the first touch was within the target, the first touch error was 0. Such zero values are included in the calculation of all means in the following analyses and in Tables 1 and 2. As anticipated, a repeated measures ANOVA showed significant effects on *first touch error* for *table-side* ($F_{1,8}=199.13$, $p<.001$, $\mu(\text{top}) = 0.07\text{mm}$, $\mu(\text{bottom}) = 2.94\text{mm}$), *target distance* ($F_{2,8} = 31.15$, $p<.001$, $\mu(25\text{mm}) = 2.72\text{mm}$, $\mu(50\text{mm}) = 1.22\text{mm}$, $\mu(75\text{mm}) = 0.55\text{mm}$), and *target size* ($F_{2,8}=271.11$, $p<.001$, $\mu(150\text{mm}) = 1.18\text{mm}$, $\mu(300\text{mm}) = 1.51\text{mm}$, $\mu(450\text{mm}) = 1.75\text{mm}$).

Post hoc tests revealed that the means of *first touch error* for each target distance were significantly different from one another at all levels of *target size* on the underside of the table, but that targets of size 20 and 30 mm did not result in significantly lower *first touch errors* when working on the top of the table. Also as predicted, there was a significant *distance x table-side* interaction ($F_{4,4}=47.6, p=.002$), supporting our hypothesis that a lack of feedback of hand position during movement compounds the difficulty of the selection task.

When touching on top of the table, average *first touch error* was 1 mm. On the bottom of the table, the average *first touch error* was 3 mm. Table 1 shows how often the target was missed on the first attempt, and Table 2 shows the average *first touch error* for each target size and distance on both surfaces of the table.

		Target Distance					
		60 mm	120 mm	180 mm	60 mm	120 mm	180 mm
Target Size	10 mm	14%	10%	12%	66%	77%	78%
	20 mm	1%	1%	1%	30%	43%	43%
	30 mm	1%	1%	1%	12%	17%	22%
Table-top				Table-bottom			

Table 1. Missed target rate: percentage of trials on which the target was missed on the first touch, by target size and distance, for both table sides.

		Target Distance					
		60 mm	120 mm	180 mm	60 mm	120 mm	180 mm
Target Size	10 mm	0.16 (0.67)	0.16 (0.83)	0.27 (0.75)	4.05 (5.79)	5.78 (6.27)	6.69 (6.52)
	20 mm	0.02 (0.05)	0.08 (1.50)	0.00 (0.04)	2.01 (5.01)	2.41 (4.52)	2.98 (5.03)
	30 mm	0.02 (0.23)	0.07 (0.91)	0.01 (0.17)	1.02 (4.95)	1.00 (3.07)	1.21 (3.28)
Table-top				Table-bottom			

Table 2. First touch error: mean distance in mm from first touch to the nearest pixel of the target, by target size and distance, for both sides of the table. Standard deviation in parentheses. 1 mm equals approximately 2.5 pixels.

It is clear from these results that, with sufficiently large targets, users are able to target an on-screen object without visual feedback. The results also point to the ability for the direct-touch input paradigm to be effectively extended to the bottom of the table. It is important to note that although the participants were not able to track their hands while moving on the ‘table bottom’ condition, the position of the hand prior to beginning each selection was known, because visual feedback of the previous touch-point was provided. To ensure accuracy consistent with these results, a known resting point should be provided to the user, similar to the nubs on the home-row of most keyboards.

Of interest to designers is that, on average, when participants missed smaller targets on the first touch, the point of contact was actually *within* the area that would have been occupied by the next-larger target. This suggests that, when accuracy of the first, blind touch is important, a system that renders targets smaller than their actual area of influence would do better than a system which simply has larger targets. This is distinct from systems where visual feedback of the position is given prior to selection, where it was found that accuracy was identical for these two conditions [8].

Having examined the accuracy of pointing without feedback, and confirmed that inverted and two-sided tables are able to effectively retain a direct-touch input paradigm, we now turn our attention to various interaction scenarios we have implemented.

INTERACTION SCENARIOS

The following interaction scenarios are meant to demonstrate the potential of two-sided touch tables to enhance applications developed for tabletops.

Synchronous Turn-Taking Games

There is a class of board games, one of which is *Diplomacy* [30], where players secretly write down their moves, and then all moves are revealed and executed simultaneously. To exemplify this class of games while making it immediately accessible to novice players, we modified the game play of *Checkers* such that it required similar input: both players specify their moves to the system without revealing them, and the system then executes both moves simultaneously and renders the appropriate outcome.




Figure 11. Checkers under the table. Left: beginning state of the board, users both enter their moves secretly by touching under the table. Right: the system evaluates both moves simultaneously and renders the consequence.

Input to the system used a version of Potter et al.’s take-off technique [15], modified so that it had no visual feedback or offset between finger and target location. To provide feedback to the user, a coordinate system was drawn around the board, and private confirmation of each input, including the current location of the finger prior to selection, was provided using a text-to-speech engine giving output to an earphone. It sometimes occurred that two users would make contact under the table while giving input. We reduced this by giving them the option of rendering the board twice, allowing each to give input simultaneously without making contact.

Two-Handed Map Browsing

Similar conceptually to that presented by Forlines and Shen [7], our map browsing application allows users to define arbitrary regions of zoom while viewing geographical data. In our version, (Figure 12), the location of the zoom window is determined by the placement of one hand above the table. To control the zoom level, we leverage the same ‘grab and separate’ gesture used in [7], but with two fingers of the same hand below the table. The focus of the zoom can be moved by dragging a single finger on the lower surface, without occluding display.




Figure 12. Two-handed two-sided map browsing application. A hand above the table controls the window location and size, the hand below the table controls zoom level and focus.

The zoom and focus controls could not be implemented on a traditional touch table without breaking the direct-touch input paradigm, temporarily occluding the area of interest with the hands, or requiring off-target controls.

3D Object Modelling and Painting

We modified a 3D modelling and painting application to accept input from our two-sided table. The application, *Smooth Teddy*, is a newer version of the tool presented by Igarashi et al. [10]. It allows the user to build smooth 3D models using two-dimensional strokes and paint. The interface is intended to be mouse driven, and requires frequent mode switches between model modification (painting / modelling) and model manipulation (rotation / translation). To adapt the interface for use on a two-sided table, we first mapped all painting and modelling functions to touches above the table, and all positioning functions to the bottom of the table.



Figure 13. 3D object modeling and manipulation. Left: top-view of image being edited. Right: the user makes use of both hands simultaneously to draw and control orientation and position of the model.

We envisioned that users would work with the application with their dominant hand working on the top, while keeping their non-dominant hand positioned on the bottom surface.

For 3D rotation, *Smooth Teddy* uses the familiar ArcBall technique [22], where the user drags a single point to rotate the model under the cursor. In order to apply the metaphor of the model existing in a volume between the top and bottom of the table, the ArcBall manipulations were reversed, so that, for example, dragging the finger to the left caused the model to rotate to the right. This maps to the ArcBall metaphor in that manipulations under the table are to the *back* of the ArcBall, and so moving the back of the ball to the left moves the front of the model to the right. An evaluation of whether this reversal is preferred might help to reveal how well the two-sided table affords a 3D space.

Two-Handed High-Speed Games: Two-Paddle Pong

An advantage of using both the top and bottom of the table for interaction is that asymmetric bimanual actions do not interfere with one another, either through collision or requiring path changes to avoid collision: both arms can be moved anywhere within reach without colliding, or needing to change the path input to the application in order to avoid the other arm. To take advantage of this property, we implemented a bimanual version of the classic game *Pong*. In the traditional game, each player controls a single ‘paddle’, capable of movement in only one axis, to deflect a ball toward the other player’s paddle. In our version, each player controls two paddles, one with each hand, placed on either surface of the table. In our version, bouncing the ball between the player’s own paddles will cause the ball to accelerate, creating a game play advantage.




Figure 14. Two-paddle *Pong*. Left: playing on the two-sided table. Right: Each paddle is operated with a different hand. Trapping the ball between the paddles causes it to accelerate.

Since the paddles could move only along a line, we mapped only the horizontal position of the hands within the game-area to the position of the paddles. This allowed users to run their hands along the border of the table closest to them, providing a more comfortable position, while maintaining one dimension of absolute control.

Without a two-sided table, this game could not have been played on a touch table. The asynchronous actions of the arms would cause them to collide. By separating the input between the two surfaces of the table, complete freedom of movement is afforded to each arm.

DESIGN RECOMMENDATIONS

The following design recommendations, drawn from our experiences in developing interaction scenarios and from observations and results of the user study, are included to aid designers attempting to implement inverted or two-sided touch table systems.

Choose Appropriate Subset of Properties

Many of the properties of inverted and two-sided tables delineated at the beginning of this paper are mutually exclusive. For example, to leverage the higher accuracy of direct-touch input, visual feedback must be provided, making privacy of input impossible; bimanual and multi-user co-locality cannot be used simultaneously; and mousing input by table side cannot be done while also trying to use both sides to reduce physical interference between users. Designers must take care to choose an appropriate subset of the listed properties in order to build a usable system.

Ergonomics

Three general ergonomic factors are important when designing a two-sided table. First, the thickness must be minimized, in order to reduce parallax between bottom-side input space and top-side display space. Second, the height of the table must be carefully set to allow appropriate comfort for the arms while performing input. Lastly, we recommend the inclusion of a second surface below the input device to allow the arms to rest while performing under the table input (Figure 13).

Design Interactions for Comfort

Designers must take care when designing for inverted or two-sided tables that they address issues of physical comfort that are not of concern when designing for a traditional tabletop. Although technically possible and potentially semantically desirable, requiring, for example, simultaneous direct-touch input by a single user to both sides of the table at opposite corners is likely to cause physical discomfort. Likewise, requiring users to constantly move their hands between the top and bottom of the table for input should be avoided. Designers should carefully consider the physical requirements on users of their system when developing new interaction methods.

Table Sides Afford Different Uses

As we have discussed, the top and bottom of the table offer distinctly different advantages and disadvantages for interaction designers. Some advantages of two sided tables, such as the avoidance of physical interference, may seem to afford interaction techniques that treat input to either side of the table identically. It is important, however, to remember that physical construction and ergonomics afford different interaction schemes, and should be considered carefully when designing interfaces that leverage an inverted or two-sided table.

Cultural and Social Issues

As mentioned previously, the private space discussed by Scott et al. [19] is likely to be larger on the underside of a table. There may also exist cultural norms of varying significance that would make some forms of under the table interaction inappropriate for a particular user population.

CONCLUSIONS

Adding an interactive touch surface to the bottom of a table adds a new dimension of affordances to touch table interaction. In this paper, we have explored the advantages and characteristics of two types of tables: an *inverted table*, with the touch surface on the bottom and display on the top, and a *two-sided* touch table, with both surfaces fully enabled for interactive input. Leveraging these technologies, a new genre of bimanual input, such as oppositional input, and asymmetric overlapping two-handed operations can now be envisioned and designed. Our user studies have shown the feasibility of using under-the-table interaction and input space to extend the direct-touch interaction paradigm. In all, we believe that under the table interaction is a promising avenue for development.

ACKNOWLEDGEMENTS

We thank Jacquelyn Martino and Mike Wu for assistance with figures, Takeo Igarashi for sharing the source code of Teddy and Smooth Teddy, William Yerazunis & Paul Dietz for construction assistance, and our participants.

REFERENCES

1. Accot, J. and Zhai, S. (1997). Beyond Fitts' law: models for trajectory-based HCI tasks. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 295-302.
2. Balakrishnan, R. and Hinckley, K. (1999). The role of kinesthetic reference frames in two-handed input performance. *Proceedings of the ACM UIST Symposium on User interface Software and Technology*. pp. 171-178.
3. Benko, H., Ishak, E., and Feiner, S. (2005). Cross-dimensional gestural interaction techniques for hybrid immersive environments. *Proceedings of the IEEE Virtual Reality Conference*. pp. 209-216.
4. Bier, E. Stone, M., Pier, K., and Buxton, W. (1993). Toolglass and magic lenses: The see-through interface. *Computer Graphics*. pp. 73-80.
5. Dietz, P. and Leigh, D. (2001). DiamondTouch: a multi-user touch technology. *Proceedings of the ACM UIST Symposium on User interface Software and Technology*. pp. 219-226.
6. Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47. pp. 381-391.
7. Forlines, C. and Shen, C. (2005). DTLens: multi-user tabletop spatial data exploration. *Proceedings of the ACM UIST Symposium on User interface Software and Technology*. pp. 119-122.
8. Grossman, T. and Balakrishnan, R. (2005). The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 281-290.

9. Han, J. (2005). Low-cost multi-touch sensing through frustrated total internal reflection. *Proceedings of the ACM UIST Symposium on User Interface Software and Technology*. pp 115-118.
10. Igarashi, T., Matsuoka, S., and Tanaka, H. (1999). Teddy: a sketching interface for 3D freeform design. *Proceedings of the ACM SIGGRAPH Conference on Computer Graphics and Interactive Techniques*. pp. 409-416.
11. MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human Computer Interaction*, 7, pp 91-139.
12. Matsushita, M., Iida, M., Ohguro, T., Shirai, Y., Kakehi, Y., and Naemura, T. (2004). Lumisight table: a face-to-face collaboration support system that optimizes direction of projected information to each stakeholder. *Proceedings of the ACM CSCW Conference on Computer Supported Cooperative Work*. pp. 274-283.
13. Morris, M., Morris, D., and Winograd, T. (2004). Individual audio channels with single display groupware: effects on communication and task strategy. *Proceedings of the ACM CSCW Conference on Computer Supported Cooperative Work*. pp. 242-251.
14. Morris, M., Huang, A., Paepcke, A., and Winograd, T. (2006). Cooperative gestures: Multi-user gestural interactions for co-located groupware. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 1201-1210.
15. Potter, R., Weldon, L., and Shneiderman, B. (1988). Improving the accuracy of touch screens: an experimental evaluation of three strategies. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 27-32
16. Rekimoto, J. (2002). SmartSkin: an infrastructure for freehand manipulation on interactive surfaces. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 113-120.
17. Ryall, K., Forlines, C., Shen, C., and Morris, M. (2004). Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. *Proceedings of the ACM CSCW Conference on Computer Supported Cooperative Work*. pp. 284-293.
18. Ryall, K., Morris, M., Everitt, K., Forlines, C., and Shen, C. (2006). Experiences with and observations of direct-touch tabletops. *Proceedings of IEEE TableTop the International Workshop on Horizontal Interactive Human Computer Systems*. pp. 89-96.
19. Scott, S., Carpendale, M.S.T., and Inkpen, K. (2004). Territoriality in collaborative tabletop workspaces. *Proceedings of the ACM CSCW Conference on Computer Supported Cooperative Work*. pp. 294-303.
20. Shen, C., Vernier, F., Forlines, C., and Ringel, M. (2004). DiamondSpin: an extensible toolkit for around-the-table interaction. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 167-174.
21. Shneiderman, B. (1991). Touchscreens now offer compelling uses. *IEEE Software*, 8(2). pp 157-165.
22. Shoemake, K. (1992). ARCBALL: a user interface for specifying three-dimensional orientation using a mouse. *Proceedings of the Graphics Interface Conference*. pp.151-156.
23. Shoemaker, G. and Inkpen, K. (2001). Single display privacyware: augmenting public displays with private information. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. pp. 522-529.
24. Smart Technologies Inc. Digital Vision Touch Technology. <http://www.smarttech.com/dvit/>
25. Sugimoto, M., Hiroki, K. (2006). HybridTouch: An Intuitive Manipulation Technique for PDAs Using Their Front and Rear Surfaces. *Proceedings of MobileHCI 2006, Espoo, Finland*. (to appear).
26. Wilson, A. (2005). PlayAnywhere: a compact interactive tabletop projection-vision system. *Proceedings of the ACM UIST Symposium on User interface Software and Technology*. pp. 83-92.
27. Wu, M. and Balakrishnan, R. (2003). Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *Proceedings of the ACM UIST Symposium on User interface Software and Technology*. pp. 193-202.
28. Wu, M., Shen, C., Ryall, K., Forlines, C., and Balakrishnan, R. (2006). Gesture registration, relaxation, and reuse for multi-point direct-touch surfaces. *Proceedings of IEEE TableTop - the International Workshop on Horizontal Interactive Human Computer Systems*. pp 183-190.
29. Yerazunis, W., Carbone, M. (2001). Privacy-enhanced displays by time-masking images. *Australian Conference on Computer-Human Interaction*.
30. [http://en.wikipedia.org/wiki/Diplomacy_\(game\)](http://en.wikipedia.org/wiki/Diplomacy_(game))