Overcoming Limitations of the Trackpad for 3D Docking Operations

David Glesser
Ensimag
Grenoble, FRANCE
glesserd@ensisun.imag.fr

François Bérard
Laboratoire d'Informatique de
Grenoble
Grenoble-INP
Grenoble, France
Francois.Berard@imag.fr

Jeremy R. Cooperstock
Centre for Intelligent Machines
McGill University
3480 University Street
Montreal, QC H3A 0E9
jer@cim.mcgill.ca

Abstract
From notebook trackpads to mobile phones to tabletop surface computing, multitouch input surfaces have become one of the most dominant interfaces for human-computer interaction. Although these are clearly effective for interaction with 2D graphical user interfaces, we suspect that they are not as well suited for interaction requiring greater degrees of freedom (DoF). Here, we consider the possibility of exploiting two such surfaces, one for each hand, as a means of affording efficient control over higher dimensional tasks. We investigate performance on a 6 DoF task, comparing such a two-surface multitouch input device against the results obtained using a standard 2D mouse, a single multitouch surface, and a 6 DoF free-space device. Our results indicate that two multitouch surfaces significantly improve user performance compared to the mouse and to a single surface.

Author Keywords
3D docking; 3D input device; empirical evaluation

ACM Classification Keywords
H.5.2 [**Information interfaces and presentation**]: User Interfaces, Input devices and strategies, prototyping.

General Terms
Human Factors; Design; Measurement.

Introduction

The traditional 2D mouse is often used for performing tasks in 3D virtual environments, for example, in Google Sketchup and Autodesk Maya. However, the growing popularity of touchscreens and surface computing, coupled with new metaphors for manipulating virtual content in 3D, suggests that multitouch devices may challenge the dominance of the mouse for 3D interaction [1, 3, 7, 14]. Furthermore, *multiple* multitouch surfaces can be combined to form an input device, such as the Cubtile [4], that ensures a topology-preserving mapping between input and output. In theory, we expect such a configuration to map more naturally to the user's mental model of the task and exploit parallel input capability between the two hands, thereby allowing greater 3D task performance efficiency than a single surface or a 2D mouse.

The various approaches to interact with 3D content can be challenging to compare directly. Some favor the ease of sketching artistic 3D scenes while others emphasize efficiency of editing precise 3D models. We focus here on the latter case, investigating two orthogonal multitouch surfaces (similar to a Cubtile [4]) by objective measurements of participants' efficiency on a precise 6 degrees-of-freedom (DoF) docking task. Our comparison to popular alternatives finds that two multitouch surfaces significantly improve user performance relative to the mouse and to a single surface.

Literature Review

Early approaches for 3D interaction sought to increase the degrees of freedom of the mouse. These included the now-ubiquitous scroll wheel [16] for control of cursor depth or adding a tilting capability to the base for the Rockin'mouse [2]. Numerous other devices have been proposed, including the Bat[18], the 3Dconnexion elastic SpaceMouse[TM] and isometric SpaceBall,[TM] and the GlobeFish and GlobeMouse [5]. Touchscreen interaction, e.g., with StickyTools [6], Turn&Roll [7], or DS3 [11], can benefit from co-location of the control surface and visualization space. However, in 3D, the mental model of "remaining in contact" with the virtual object becomes problematic when the user wishes to move the object in depth. Indirect multitouch interaction [13, 15], which need not concern itself with co-location between control and display, offers another advantage, in particular for visualization, since the users' fingers do not occlude the display. Moreover, indirect interaction was found to offer greater efficiency and precision [9] for rotation, scaling, and translation tasks [10] using two-finger inputs. For 6 DoF multitouch interaction, the multi-surface Immersion Cubtile provides an easily learned indirect bimanual gestural interaction technique that preserves the topological relationship between the input and output spaces [4]. We were enthusiastic about the possibilities of this device for 6 DoF tasks. However, various limitations including sensor resolution, sensitivity to noise or non-finger contact, and user fatigue, motivated our use of a different design for the experiment described below.

Experiment Design

In general, people are unable to solve a 6 DoF docking task in a coordinated manner. Rather, they switch control between 3 DoF rotations and translations [12]. Although this likely holds regardless of device and interaction technique, we chose to test the docking task because it is a fundamental building block of 3D interaction. Moreover, we want to determine which device and interaction technique afford the best performance, even in an uncoordinated manner.

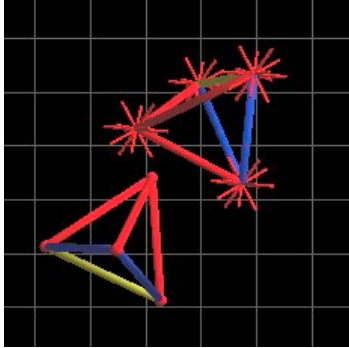We follow the design used by Froehlich et al. [5]. Users

Figure 1: Display for the experimental task. cursor (left) and target (right)
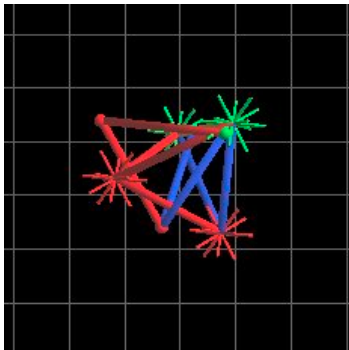

Figure 2: The cursor matches two vertices (green) of the target

are asked to translate and rotate a 3D colored tetrahedral cursor and dock it with a target tetrahedral of identical shape. Spines around the target vertices serve as an indication of docking tolerance. The cursor appears with small spheres at its vertices (Figure 1). When a cursor vertex is within its docking tolerance of the corresponding target vertex, the sphere color changes to green (see Figure 2). For consistency across all experimental conditions and to ensure reasonable depth perception of the scene [8], the display was viewed by test participants on a stereoscopic projection at a resolution of 1024×768 pixels, using polarized glasses. Participants signalled task completion by a click for the mouse and PHANTOM®or a finger-tap for the trackpad-based devices. Three distances to target and three rotation differences were used. They were varied randomly to cover all 3×3 combinations throughout each block. We ensured that translation and rotation always required movement in or about all three axes to avoid trivial conditions. Moreover, to avoid pathologically difficult viewing conditions, we ensured that no targets were presented with one vertex of the tetrahedron hiding another one. Spines around the target vertices indicates a tolerance of 32 *units*, a value chosen empirically to ensure that the task was feasible, but non-trivial, with all four devices tested. We discuss the C/D gain used with all devices below.

The devices that we included in our experiment were (a) two perpendicular multitouch surfaces, reproducing two faces of the Cubtile [4], (b) a traditional 2D mouse, used as a baseline and as the most common device for 3D modeling, (c) a single multitouch surface, as several recent publication have proposed new multitouch 3D interaction, and the focus of our study is to compare one multitouch surface against two perpendicular surfaces, and (d) a free-space 6 DoF stylus, in the form of a

PHANTOM®Omni, as this was recently shown to be the most efficient device for 3D target pointing [17].

All the 3D multitouch interaction literature assumes a *direct* interaction technique [1, 3, 7, 14]: fingers touch the screen where the 3D scene is rendered. Knödel and Hachet, however, showed that indirect interaction favors precision and efficiency in a 3D Rotation, Scaling and Translation task. In addition, an indirect interaction technique allows us to use the same display across input devices, uniformizing conditions. Since no indirect techniques existed as alternatives for touch-screen interaction, with chose to adapt the StickyTools [6] technique for indirect interaction for the purpose of our experiment. To realize our multi-surface multitouch input device, we connected two Apple Magic Trackpads on adjacent faces of a cube with sides of approximately 15 cm, as shown in Figure 3.

Preparatory Trials

We first conducted a number of informal trials with three novice participants to help us decide between various design options for the experiment. For the mouse, we compared two designs: the traditional four-panel display paradigm, with three orthographic projections and one perspective view, and a widget-based interaction paradigm using a stereoscopic perspective projection, inspired by Autodesk Maya. In the orthographic projections condition, rotation was enabled by right-clicking. Trials indicated better performance with the four-panel approach, so we adopted this for our full experiment. Similarly, we compared StickyTools [6] and Turn&Roll [7], adapted for use in a (potentially unfair) indirect control condition, and retained Sticky Tools for the full experience given that the performance differences observed were minor and the possible advantage of its use of bimanual interaction.
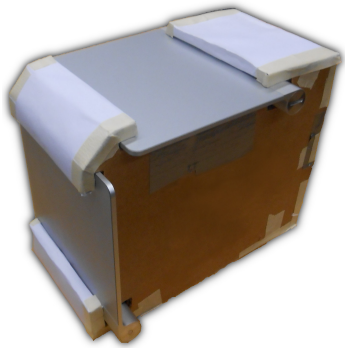
Figure 3: The Magic Cubtile.

Figure 4: Magic Cubtile interaction possibilities. The green ovals represent fingers and the arrows the corresponding movement required to effect the indicated operation.

For the multi-surface multitouch configuration, we compared three interaction techniques. MagicComplex is designed to be similar to a conventional trackpad, but using the two surfaces to effect translation or rotation control about the three axes. MagicArc uses two fingers for translation and one for rotation. This was motivated by the results of Masliah et al. [12], who showed that users tend to perform these operations sequentially, rather in parallel. MagicDual is similar, but enforces a bimanual component for rotations, as one finger must "hold" the virtual object while the other rotates it. These interaction techniques are illustrated in Figure 4.

MagicArc and MagicDual yielded similar results in our informal trials. MagicComplex was abandoned as it required considerable mental effort to use, which resulted in poorer performance. Users could easily *rotate* the cursor about all three axes simultaneously with either Magic Arc or MagicDual and could *translate* in all three axes simultaneously with MagicArc, but not with MagicDual. We thus opted to use MagicArc in the experiment. We attempted to find appropriate C/D gain values for each device by empirical testing. For the mouse, the default Ubuntu 12.04 behaviour was employed. For all other conditions, we chose a constant gain that was high enough to reduce clutching but low enough to facilitate the precision docking.

Formal Experiment

We recruited 16 participants (12 male, 4 female) ranging in age from 20-30, drawn from the student population of our university. Participants were provided with an introduction to the experiment and given an opportunity to familiarize themselves with each input device. At the end of the experiment, they were asked to complete a brief questionnaire for additional feedback and
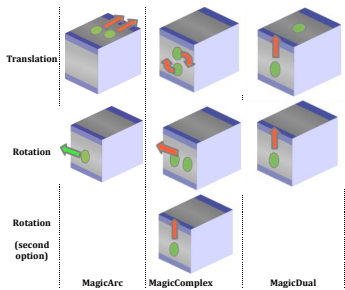
compensated for their time. All were right-handed with reasonable stereo viewing acuity, and had prior experience using a mouse and a multitouch surface. Two participants had previously used a PHANTOM.

Participants performed 10 trials for each of the 4 devices, leading to a total of $16 \times 10 \times 4$ trials. For balanced order, the devices were presented according to a *Williams design* [19]. Error rate was recorded, but was not analyzed as it remained below 4% across all conditions. This was largely confined to the trackpad and Magic Cubtile, resulting from accidental brushing of the hand against the trackpad surface. A maximum of three errors was allowed per target, which proved sufficient to ensure successful docking in all trials. Docking time was recorded from the last validation action, at which point, the target is refreshed, to the validation of the current target.
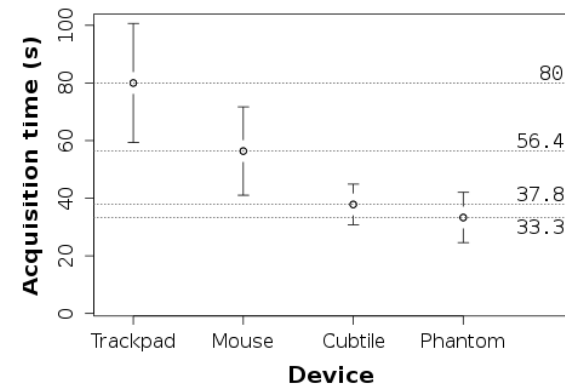

Figure 5: Mean docking time (s) by device with 95% confidence intervals.

The docking time per device is shown in Figure 5. A repeated measure ANOVA revealed a significant effect of

the device on the docking time (F(3,60)=10.39, p = 1.333e-05). Post-hoc comparisons between conditions using paired t-tests, with holm adjustments, revealed that the Magic Cubtile and PHANTOM conditions were significantly faster ($T(15) > 3.93, p < 0.01$) than the Mouse and Trackpad condition for all comparisons, except for Magic Cubtile and Mouse ($T(15) = -3.06, p < 0.05$). The mouse was also significantly faster than the trackpad ($T(15) = 3.81, p < 0.01$). We did not measure a significant difference between Magic Cubtile and PHANTOM.

Discussion and Future Work

The results confirm our hypothesis that a second multitouch surface improves efficiency of the docking task over a single trackpad. The single trackpad was clearly the worst device. We observed that all participants first attempted to rotate the cursor, then translate it toward the target. However, because the cursor can also rotate while being translated with StickyTools, the separation of these two operations proved challenging. The surprising outcome is that the Magic Cubtile outperformed the most commonly used device for 3D interaction, the standard mouse, by a large margin (33%).

The Magic Cubtile had similar performance to that of the PHANTOM, which was recently shown as the most efficient device for 3D target acquisition. The fact that these achieved comparable performance is in itself an important result, as it suggests a much less expensive alternative to the latter for 3D interaction, using commodity hardware.

Both devices offer a natural mapping of user's physical motions into the virtual cursor's motion, contrary to more traditional 2D devices. While the PHANTOM may have had a slight performance advantage over the Magic Cubtile in our experiment, the extensive use of the wrist for rotation may be problematic. Indeed, in the post-test questionnaire, one user reported that this hurt his wrist. The Magic Cubtile also offers a smaller footprint on the desk, and the afforded stability is likely superior to that of the PHANTOM. We expect that the allowed tolerance for docking plays a substantial role in the relative performance of these two devices, a question that is the focus of our follow-on study.

In the future, it would be interesting to compare results on targets of varying index of difficulty and examine performance of more expert users, after acquiring several hours of practice with each device. Expert 3D modelers, for example, are highly adept at performing their tasks. Can a skilled user outperform the PHANTOM or Cubtile with sufficient practice?

Acknowledgments

References

[1] Au, O. K.-C., Tai, C.-L., and Fu, H. Multitouch gestures for constrained transformation of 3d objects. *Computer Graphics Forum 31*, 2 (2012), 651–660.

[2] Balakrishnan, R., Baudel, T., Kurtenbach, G., and Fitzmaurice, G. The Rockin'Mouse: integral 3D manipulation on a plane. In *CHI*, ACM (New York, 1997), 311–318.

[3] Cohé, A., Dècle, F., and Hachet, M. tBox: A 3D transformation widget designed for touch-screens. In *CHI*, ACM (2011), 3005–3008.

[4] de la Rivière, J.-B., Kervégant, C., Orvain, E., and Dittlo, N. Cubtile: a multi-touch cubic interface. In *VRST*, ACM (New York, 2008), 69–72.

[5] Froehlich, B., Hochstrate, J., Skuk, V., and Huckauf, A. The globefish and the globemouse: two new six degree of freedom input devices for graphics applications. In *CHI*, ACM (New York, 2006), 191–199.

[6] Hancock, M., Cate, T. T., and Carpendale, S. Sticky tools: full 6dof force-based interaction for multi-touch tables. In *Interactive Tabletops and Surfaces*, ACM (New York, 2009), 133–140.

[7] Herrlich, M., Walther-Franks, B., and Malaka, R. Integrated rotation and translation for 3D manipulation on multi-touch interactive surfaces. In *Smart Graphics*, Springer-Verlag (Berlin, Heidelberg, 2011), 146–154.

[8] Hubona, G. S., Shirah, G. W., and Jennings, D. K. The effects of cast shadows and stereopsis on performing computer-generated spatial tasks. *Trans. Sys. Man Cyber. Part A 34*, 4 (July 2004), 483–493.

[9] Knödel, S., and Hachet, M. Multi-touch RST in 2D and 3D Spaces: Studying the impact of directness on user performance. In *3DUI*, IEEE (2011).

[10] Kurtenbach, G., Fitzmaurice, G., Baudel, T., and Buxton, B. The design of a gui paradigm based on tablets, two-hands, and transparency. In *CHI*, ACM (New York, 1997), 35–42.

[11] Martinet, A., Casiez, G., and Grisoni, L. Integrality and separability of multitouch interaction techniques in 3D manipulation tasks. *IEEE Transactions on Visualization and Computer Graphics 18*, 3 (2012), 369–380.

[12] Masliah, M. R., and Milgram, P. Measuring the allocation of control in a 6 degree-of-freedom docking experiment. In *CHI*, ACM (New York, 2000), 25–32.

[13] Moscovich, T., and Hughes, J. F. Indirect mappings of multi-touch input using one and two hands. In *CHI*, ACM (New York, 2008), 1275–1284.

[14] Reisman, J. L., Davidson, P. L., and Han, J. Y. A screen-space formulation for 2d and 3d direct manipulation. In *UIST*, ACM (New York, 2009), 69–78.

[15] Schmidt, D., Block, F., and Gellersen, H. A comparison of direct and indirect multi-touch input for large surfaces. In *INTERACT*, Springer-Verlag (Berlin, Heidelberg, 2009), 582–594.

[16] Venolia, D. Facile 3D direct manipulation. In *CHI*, ACM (New York, 1993), 31–36.

[17] Wang, G., McGuffin, M. J., Bérard, F., and Cooperstock, J. R. Pop-up depth views for improving 3D target acquisition. In *Graphics Interface* (2011), 41–48.

[18] Ware, C. Using hand position for virtual object placement. *Vis. Comput. 6*, 5 (Nov. 1990), 245–253.

[19] Williams, E. J. Experimental designs balanced for the estimation of residual effects of treatments. In *Australian Journal of Scientific Research*, Ser. A 2 (1949), 149–168.