The "Prince" Technique: Fitts' Law and Selection Using Area Cursors

Paul Kabbash
Input Research Group
Computer Systems Research Institute
University of Toronto
Toronto, Ontario,
Canada M5S 1A4
Tel: +1 (416) 978-6619
kabbash@dgp.toronto.edu

William Buxton
Alias Research, Inc.&
University of Toronto
c/o Alias Research, Inc.
110 Richmond St. East
Toronto, Ontario,
Canada M5C1P1
Tel: +1 (416) 362-9181
buxton@alias.com

ABSTRACT

In most GUIs, selection is effected by placing the point of the mouse-driven cursor over the area of the object to be selected. Fitts' law is commonly used to model such target acquisition, with the term A representing the amplitude, or distance, of the target from the cursor, and W the width of the target area. As the W term gets smaller, the index of difficulty of the task increases. The extreme case of this is when the target is a point. In this paper, we show that selection in such cases can be facilitated if the cursor is an area, rather than a point. Furthermore, we show that when the target is a point and the width of the cursor is W, that Fitts' law still holds. An experiment is presented and the implications of the technique are discussed for both 2D and 3D interfaces.

KEYWORDS: Input techniques, graphical user interfaces, Fitts' law, haptic input.

INTRODUCTION

Although the traditional method of selection in direct-manipulation systems is generally effective, there are certain conditions where it breaks down. One of these is when the target is very small. The extreme case of this is when the target is a point. The reason for the problem can best be explained by Fitts' law [3, 8], expressed as:

$$ID = \log_2(A/W + 1). \qquad (1)$$

From this formulation we see that the index of difficulty (ID) of a target acquisition task is a function of the amplitude (A), or distance, of the target from the cursor, and the width of the target (W). The index of difficulty rises as the width of the target gets smaller.

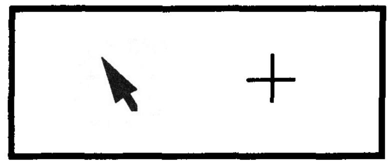

Figure 1: Two typical "point" cursors. Selection with conventional GUIs is literally "pointing," for example, with the point of the arrow or the intersection point of the cross-hair.

With most conventional GUIs, the selection tool is a point, such as represented by the point of an arrow shaped cursor, or the intersection of the lines in a cross-hair cursor (Figure 1). In the extreme case, therefore, we are selecting a point with a point.

The purpose of the research described in this paper is to explore an alternative approach whereby (in the 2D case) the cursor is represented by an *area*, rather than by a point. Just as the area of a fly-swatter makes it easier to swat a small fly, likewise the area of such a cursor should make it easier to select small targets and points.

More formally, it is our claim that selecting a small target with an area cursor can be modeled by a slight twist of Fitts' law, namely, that the W term now applies to the width of the cursor, rather than the width of the target. Figure 2 illustrates the approach using an area cursor, as well as the traditional approach.

In what follows, we report on an experiment that demonstrates the applicability of Fitts' law to selecting point targets with an area cursor. We follow this with a discussion of the design implications of our findings to other 2D and 3D tasks.

Finally, due to the similarity of their benefits, we name the use of area cursors after the first manufacturer of oversized tennis rackets: the *Prince* technique.

THE EXPERIMENT

Subjects

Twelve students from the University of Toronto participated as paid volunteers. All had experience using the mouse and were strongly right handed based on the Edinburgh Handedness Inventory [10].

Apparatus

Equipment was an Apple *Macintosh IIfx* with 13-inch RGB monitor. Subjects performed the tasks using their right hand and a standard mouse. The control/display ratio of the mouse was adjusted to the second fastest setting on the *Macintosh* Control Panel. Since even small lags (75 ms) in display response have been found to degrade human performance on Fitts' law tasks [5, 9], the software was optimized to ensure that drawing updates did not delay movement of the Prince cursor.

Before the experiment, we tested for animation delays by making the system "arrow" pointer visible in the center of the paddle and found that it was not possible to shake the paddle from the arrow, even at movement speeds likely to be much faster than those encountered during the experiment.

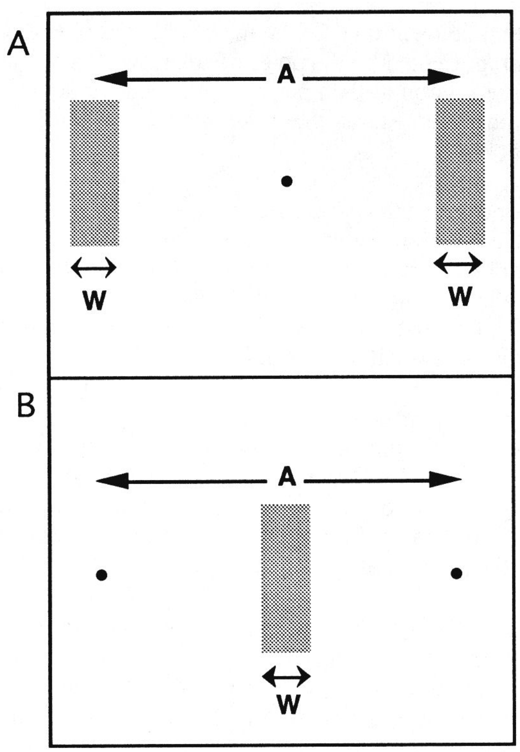

Figure 2: Two versions of Fitts' reciprocal aiming task. The conventional approach is shown in 2(a). Here a standard "point" cursor is moved amplitude *A* between two targets of width *W*. In 2(b) we see the variation, where an area cursor of width *W* is moved between two point targets separated by distance *A*.

Procedure

Subjects performed a reciprocal point-select task using both the Fitts and Prince techniques (Figure 2). They were given written instructions and several warm-up trials prior to data collection. In addition, they performed one practice session on each technique.

For each technique, two targets appeared on either side of the monitor. Subjects moved the cursor back and forth between the targets and selected each target by pressing and releasing the mouse button. They were instructed to balance speed and accuracy for an error rate around 4%, and an error beep sounded if selection occurred outside the target. Results of movement time and error rate were given to subjects at the end of each session.

The cursor and the target were represented in the two techniques using different objects. In the Fitts condition the cursor was a small black dot with radius 2 pixels, and the targets were rectangles having width *W* and height 200 pixels. The Prince condition reversed these objects exactly (see Figure 2b), so that subjects controlled a rectangular cursor (width *W*, height 200 pixels) and used it to capture two target dots (radius 2 pixels). In both conditions, the rectangles were unframed and shaded light blue. The Prince cursor was transparent, so that the targets could be clearly seen beneath it. When the target and cursor overlapped, their appearance in the two techniques was nearly indistinguishable, the primary difference being which object moved or was stationary.

Design

A fully-crossed, within-subjects factorial design with repeated measures was used. Factors were movement amplitude (64, 128, 256, and 512 pixels), target or cursor width (8, 16, 32, and 64 pixels), and technique (Fitts and Prince). The amplitude and width conditions yielded seven levels of task difficulty, ranging from 1 bit to 6.02 bits. The *A-W* conditions were presented in random order with a block of ten trials performed at each condition. A session consisted of a sequence of sixteen blocks covering all *A-W* conditions. After training, ten sessions were performed in all, alternating between the Fitts (five sessions) and the Prince (five sessions) techniques. The order of techniques was counterbalanced, with half of the subjects beginning with the Prince technique and the other half with the Fitts technique. Subjects took about one hour to complete the experiment.

Dependent variables were movement time (*MT*), error rate (*ER*), constant error (*CE*), and variable error (*VE*). The latter two measures were used to describe the quality of placement of the response selections [11]. *CE* was measured in the horizontal axis, as the signed distance between target and cursor centers at the moment a selection occurred, and was used to detect systematic trends towards undershooting or overshooting the target center. *VE* captures the endpoint variability of responses and corresponds to effective target width ($W_e = 4.133 \times SD_x$; see [8]).

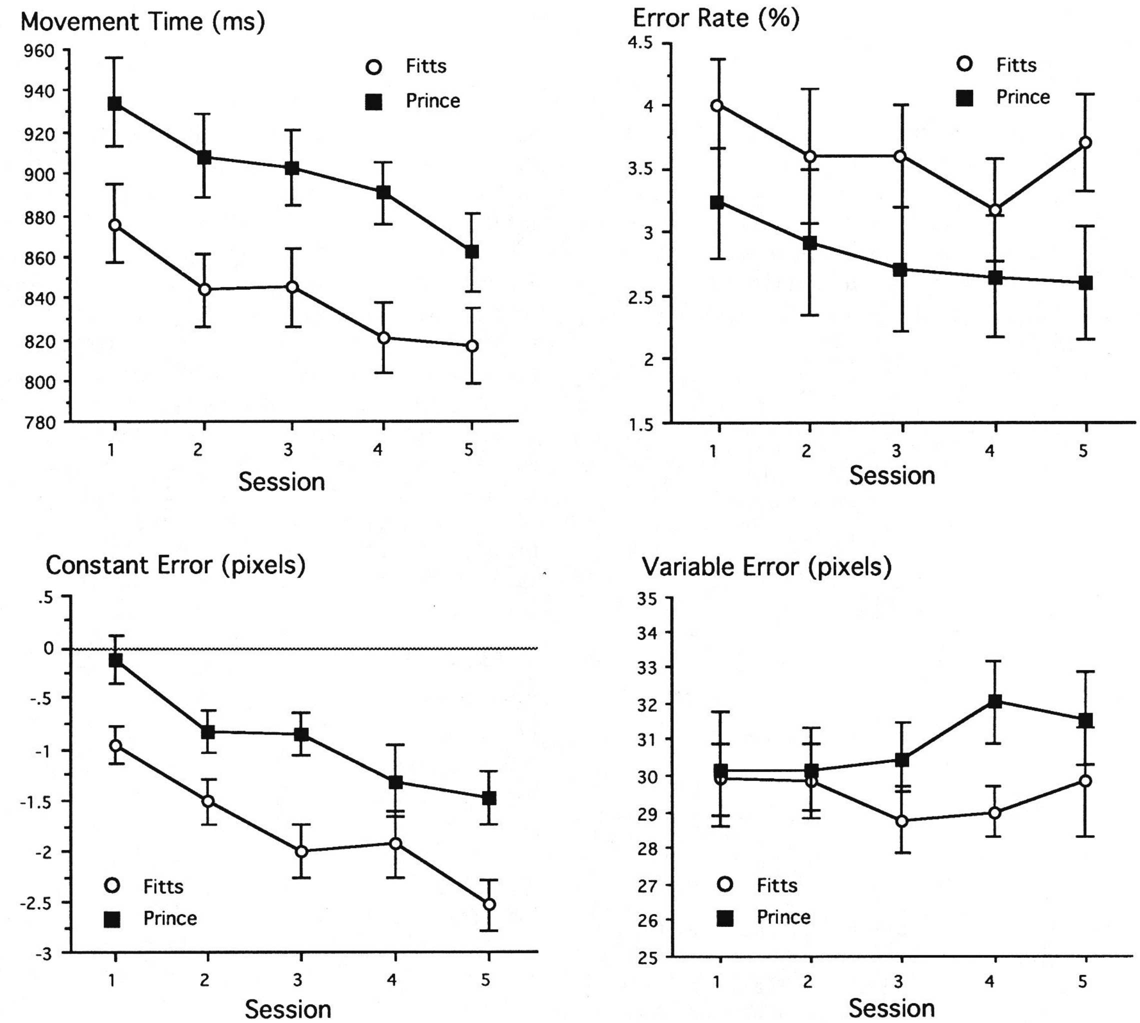

Figure 3: Means (with standard error bars) for movement time, error rate, constant error and variable error, decomposed by technique and session.

Results

There were significant main effects of technique on all four dependent variables. While subjects were slower using the Prince technique (900 ms vs. 841 ms; $F_{1,10} = 273.5$, $p < .001$), they produced fewer errors (2.8% vs. 3.6%; $F_{1,10} = 15.7$, $p < .005$) and tended to aim nearer to the target centers ($CE = -.922$ vs. -1.787 pixels; $F_{1,10} = 27.0$, $p < .001$). Subjects therefore appear to have been more careful when making selections with the Prince technique. This did not however provide an advantage in terms of motor response variability. With both techniques, *VE* was very close to the average nominal width of 30 pixels. However, *VE* was 30.9 pixels in the Prince technique and 29.5 pixels in the Fitts technique, so subjects were somewhat less variable in endpoint placement using the Fitts technique ($F_{1,10} = 9.40$, $p < .02$).

This suggests that the bias towards accuracy rather than speed in the Prince technique was due to its unfamiliarity. We investigated this possibility further, by examining performance in relation to learning phase over the five sessions (Figure 3). The analysis revealed a main effect of session on *MT* ($F_{4,40} = 11.8$, $p < .001$) representing a small improvement for both techniques, in total less than 8% from sessions 1 to 5. A significant effect of session on *CE* ($F_{4,40} = 10.4$, $p < .001$) suggests that subjects also increasingly undershot the targets as they grew more confident with the task. They were able to do so without incurring greater errors or endpoint variability, as evidenced by the lack of session effects on *ER* ($F_{4,40} = 1.07$, $p > .05$) and *VE* ($F_{4,40} = .543$). However, the two-way interaction of technique x session was not significant for any of the four dependent variables ($F_{4,40} < 1.06$, $p > .05$), implying that the speed-accuracy tradeoff in the Prince technique did not change relative to the Fitts technique even as subjects progressed through the trials. Thus, it is unclear from the present data to what extent further practice with the Prince technique would have altered its accuracy bias.

Fit of the Model

Since there were no significant interactions on *MT* between session effects and those of the other independent variables, regression lines were fitted to the *MT* data averaged over sessions 1 to 5 (Figure 4). The regression analyses were performed on the data normalized for errors, using the method described in [9].

As predicted by Fitts' law, there was a linear relationship and high correlation between MT and ID for both techniques. The fitted equations were $MT = 198 + 204\ ID$ ($r = .978$) for the Fitts technique and $MT = 267 + 203\ ID$ ($r = .965$) for the Prince technique, with the model explaining 95.6% of the variability in the Fitts technique ($F_{1,15} = 302.5$, $p < .001$) and 93.2% in the Prince technique ($F_{1,15} = 191.4$, $p < .001$).

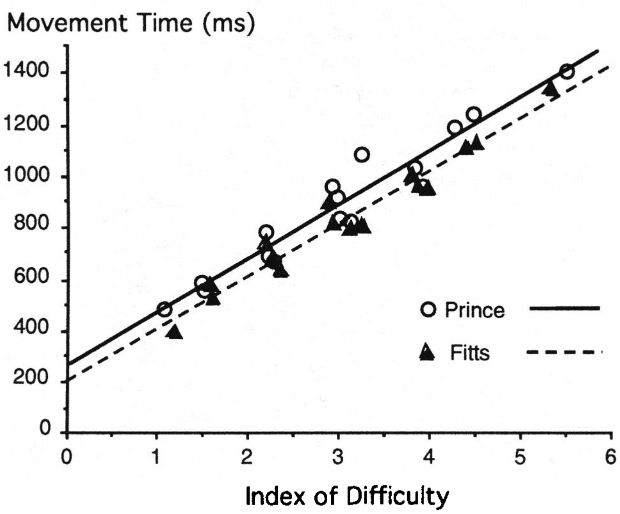

Figure 4: Scatter-plots of the *MT-ID* relationship in the Fitts and Prince techniques. The equation fitted was $MT = a + b\ ID$, where $ID = \log_2(A/W_e + 1)$. The reciprocal of the slope of each line gives the index of performance ($IP = 1/b$) for the technique.

As is evident in Figure 4, the slopes of the two regression lines did not differ. Each yielded an index of performance (*IP*) of about 4.9 bits/s, which is comparable to earlier experiments [7]. Because the techniques had different intercepts, however, the *MT-ID* relationship in the Fitts technique was displaced slightly downward with respect to the Prince technique.

Individual subject regressions also were computed. These exhibited similar trends with no differences between techniques for slope ($F_{1,10} = .416$), but a significantly greater positive intercept in the Prince technique ($F_{1,10} = 22.1$, $p < .001$). The correlation coefficients obtained for individual subjects were, with one exception, greater than .90 and revealed no differences between techniques ($F_{1,10} = 1.04$, $p > .05$). Hence the results for the aggregated regressions given in Figure 4 appear to have been consistently exhibited within each subject.

Discussion

While the results clearly demonstrate that Fitts' law applies to the Prince technique, a question remains as to why the performance differences between techniques were reflected in the intercept of the regression lines, rather than the slope.

One interpretation of the intercept is that it represents time spent on the targets rather than time spent moving between them [13, p. 146]. Considered in this way, "time on target" includes only the time the cursor is held motionless over the target.

For the reciprocal aiming task used in this experiment, time on target would include the time necessary for the subject to verify that the cursor is over the target, the time to execute the button press itself, as well as preparation time to program the next movement, as in [1]. There is evidence that the verification component, in particular, is sensitive to the accuracy demands and objectives of the task [1, 6, 12]. Thus, if subjects in our experiment were in fact being somewhat more careful with the Prince technique than the Fitts technique, this may have increased their verification time and hence the intercept.

DESIGN IMPLICATIONS

"The Prince and the Pointer"

Having established that the area and standard cursors follow similar prediction models, we now consider some of the properties of using an area cursor as a positioning and selecting tool.

There appear to be two main benefits of applying the Prince technique. The first is illustrated in Figure 5 and contrasts the difficulty of acquiring a small target using the standard cross-hair cursor and a rectangular area cursor. For such tasks, the area cursor approach is clearly much easier.

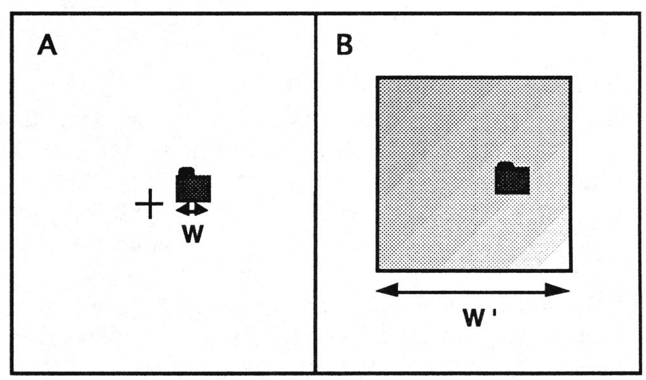

Figure 5: In 5(a) the target is selected using a standard "cross-hair" cursor. The difficulty of the task is limited by the size of the target (W). In 5(b), an area cursor with width W' surrounds the target to select it. The difficulty of this task is a function of W'.

We can use the results from the experiment to quantify the difference between the two approaches. Assume the target has width W, the area cursor width W', and the distance moved in both cases is A. Then, when A/W is large (i.e., the task is hard), the difference in index of difficulty (ID, equation 1) for the two tasks will approach $\log_2(W'/W)$.

For example, if W is 6 pixels and W' is 96 pixels, the Prince technique represents a savings of about 4 bits as rated by Fitts' law. Given the performance level arrived at in the experiment (4.9 bits/s), this translates to a movement time savings of roughly .75 s per mouse selection. In the case of $A = 384$ pixels, for instance, this is a 93% reduction. (Of course, the movement time savings will be even larger with a device that does not perform as well as the mouse; e.g., using $IP = 1.5$ bits/s, reported in [7] for trackball performance during a dragging task, the predicted savings are well over 2 s per selection.)

A second capability of the area cursor is that it may function as a "net." Used in this way, an area cursor can group and select a collection of points or small objects with a single pointing movement, much as the "lasso" tool is used in drawing applications like MacDraw. This capability, however, also serves to illustrate a drawback with using the area cursor as the only selection tool in a GUI. This is that the Prince technique is inappropriate for fine positioning tasks, because selections may become ambiguous when displays are cluttered.

Our belief is that an effective way of exploiting the Prince technique is to combine it with the traditional point-cursor approach. Where fine positioning is not required, it may be possible to replace it by coarse positioning and the Prince technique. Furthermore, by dynamically switching between Prince and point-cursor positioning techniques, the difficulty of positioning tasks can be matched more closely to task context.

We illustrate this using three examples.

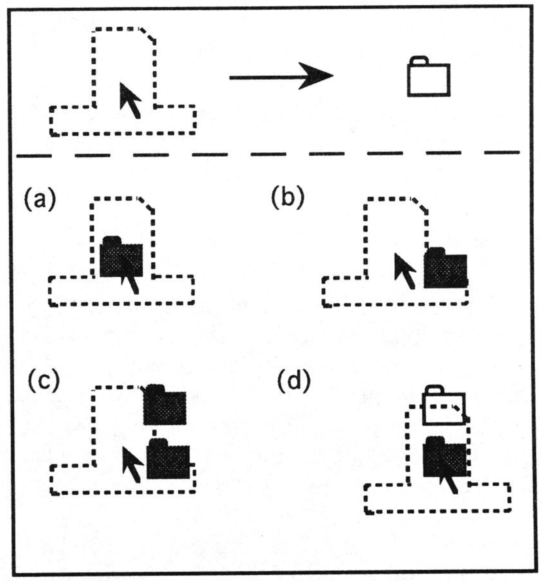

Figure 6: Moving a file into a folder. The user acquires a file and begins to drag its outline toward the folder. The folder will be highlighted to indicate selection. This occurs in (a) when the pointer moves inside the folder, and in (b) when the folder and the file overlap. In (c), the selection is ambiguous because the file overlaps with two folders. This can be resolved as in (d), by repositioning the pointer in one of the folders.

Example 1: Dragging a File into a Folder

Figure 6 gives an example of how use of the area and point cursors can be combined in traditional 2-dimensional GUIs. It illustrates how task difficulty can be matched to the accuracy demands of the task, through only a small modification of current practice. In this case, the cursor switching is effected by the system rather than the user.

The task is moving a file to a folder, as in the *Macintosh* Finder (Figure 6). Typically, an outline of the file is displayed beneath the mouse pointer as it is being dragged towards the folder. The drag outline can be interpreted as an area cursor. We can then define the task of acquiring the folder in two ways, with the system responding either to the location of the pointer (Figure 6a) or the icon of the file being dragged (Figure 6b).

The second representation of the task (Prince technique) will be easier whenever the file icon is larger than the folder icon. (The second task may also be more intuitive. e.g., Why do people miss the Trash icon so often? Perhaps it's because we're attending to the file we're moving, rather than the location of the pointer.) If ambiguity results with this technique, as in Figure 6c, users can either reposition the file cursor to remove the ambiguity or they can revert to positioning with the point cursor. That is, fine positioning need only be invoked as a last resort (when displays are cluttered), not as the default.

Example 2: Toolglass & Magic Lenses

The Prince and point-cursor techniques may be used in combination. This is seen, for example, in the Toolglass and Magic Lens techniques introduced by Bier, Stone, Pier, Buxton, & DeRose [2]. These techniques illustrate an elegant solution to the "display clutter/ambiguity" problem raised in the previous example.

With them, the area and point cursors are represented as separate objects and distributed between the two hands. The area cursor—that is, the Toolglass and Magic Lens widgets themselves—is controlled by the nondominant hand, while the dominant hand manipulates the point cursor.

When a task does not need fine positioning, the widgets can be used alone. An example is using a magnifying magic lens. Because of the widget's size, it is easy to position over the desired regions of the screen with the nondominant hand, and there are few if any negative consequences of "spill over" to other objects. However, when fine positioning is required, the technique can take advantage of the interaction between the two hands. An example is *clicking-through* the magnification widget with the point cursor in order to select one of the small objects being magnified. In this example, the synergies work to our advantage since the magnification widget itself increases the width W of the target object, therefore reducing the index of difficulty in its selection.

The head prop and slicing tool [4] illustrates a similar distribution of coarse and fine positioning tools between hands.

Example 3: Silk Cursor

All of our examples thus far have involved 2D selection using an area cursor. The Prince concept can be applied in 3D as well. In this case, it takes the form of a *volume* cursor. This was seen in the silk cursor study of Zhai, Buxton, & Milgram [14]. This was our first study employing the Prince technique, although its main objective was to investigate the effectiveness of occlusion cues in 3D selection, rather than Fitts' law.

A 3D dynamic target acquisition task, "virtual fishing," was designed for the experiment. In each trial of the experiment, an "angel fish" with random size and color appears swimming randomly within a 3D virtual environment. The subjects were asked to move a 3D volume cursor to envelop the fish and "grasp it" when the fish was perceived to be completely inside of the cursor.

In the experiment, only one fish was presented at a time since the technique breaks down in crowded waters.

A partial solution worth studying is to give the net a size operation (e.g., perhaps performed by the other hand). If you're swimming in open water, the net can be large. But when you swim into a school of fish, then you'd scale the net to make it smaller. This makes it easier to swim between the fish. At some point, however, two fish are going to be so close that they can't easily be distinguished even with a small net. The size operation by itself is also not ideal, since important occlusion information about the targets is lost when the net is too small.

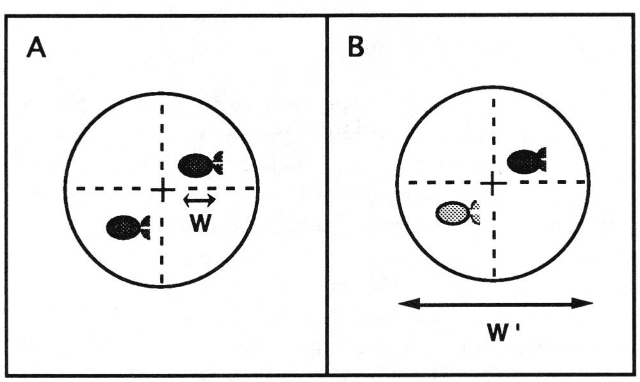

Figure 7: The silk cursor with a homing grid. The depth occlusion cues in 7(a) tell the user that two fish are caught in the net. Hence, the homing grid must be used to select one of the fish. In 7(b), one fish is shown within the net while the other is in front.

The software can help here by drawing a "homing grid" inside the net (with a cross-hair cursor showing the center of mass of the net). The homing grid can always be shown, even when the net is large. Presumably, there will be a cross-over point in terms of what the user pays attention to. When tasks are coarse, selections can be made with reference to the surface of the net. When the display is cluttered, attention shifts more to the homing grid. But even here the silk net can help with positioning by providing occlusion information about the environment (see Figure 7).

CONCLUSION

An alternative approach to pointing, called the *Prince* technique, was investigated and found to be comparable to traditional pointing methods. Because the Prince technique uses a cursor of large area or volume, it is suitable for tasks that are normally difficult with the standard pointer, such as acquiring small targets or points. We feel that the Prince technique may be especially valuable when used in conjunction with traditional pointing techniques, where it can be used to tailor task difficulty more closely to the accuracy demands of the task. The examples presented three distinct methods suggesting how this might be accomplished.

The current study is an initial probe into a rich design space. Many questions and issues remain. We investigated selection tasks involving one width parameter, either the target or the cursor. What happens when there are two width parameters, defined by moving and stationary objects? The whole issue of "grasping" isolated objects from among a close cluster requires much more investigation. Likewise, the 3D case of the volume cursor deserves study. It would also be worthwhile to compare and/or combine the technique with gravitational "snapping" techniques. Finally, for the full potential of the technique to be realized, it is likely that new affordances (such as supporting "grasping") need to be built into input devices, such as mice. This also requires further study.

ACKNOWLEDGMENTS

This research has been undertaken under the auspices of the Input Research Group (IRG) of the University of Toronto. The authors gratefully acknowledge the contribution of the members of the group to this work. We would especially like to acknowledge the contribution of Shumin Zhai. We would also like to thank Abigail Sellen of Rank Xerox EuroPARC, Scott MacKenzie of the University of Guelph, and Mark Tapia of the University of Toronto for their contributions.

The IRG is generously supported by the Information Technology Research Centre of Ontario (ITRC), the Natural Sciences and Engineering Research Council of Canada (NSERC), Xerox PARC, Alias Research Inc. and Hitachi Corp. This support is gratefully acknowledged.

REFERENCES

1. Adams, J. J. (1992). The effects of objectives and constraints on motor control strategy in reciprocal aiming movements. *Journal of Motor Behavior, 24*, 173–185.

2. Bier, E. A., Stone, M. C., Pier, K., Buxton, W., & DeRose, T. D. (1993). Toolglass and magic lenses: The see-through interface. *Proceedings of SIGGRAPH '93* (pp. 73–80). New York: ACM.

3. Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology, 47*, 381–391.

4. Hinckley, K., Paush, R, Goble, J. C., & Kassell, N. F. (1994). Passive real-world interface props for neurosurgical visualization. *Proceedings of the CHI '94 Conference on Human Factors in Computing Systems* (pp. 452–458). New York: ACM.

5. Hoffman, E. R. (1992). Fitts' Law with transmission delay. *Ergonomics, 35*, 37–48.

6. Jagacinski, R. J., Repperger, D. W., Moran, M. S., Ward, S. L., & Glass, B. (1980). Fitts' law and the microstructure of rapid discrete movements. *Journal of Experimental Psychology: Human Perception and Performance, 6*, 309-320.

7. MacKenzie, I. S., Sellen, A., & Buxton, W. (1991). A comparison of input devices in elemental pointing and dragging tasks. *Proceedings of the CHI '91 Conference on Human Factors in Computing Systems* (pp. 161–166). New York: ACM.

8. MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction, 7*, 91–139.

9. MacKenzie & Ware. (1993). Lag as a determinant of human performance in interactive systems. *Proceedings of the CHI '93 Conference on Human Factors in Computing Systems* (pp. 488–493). New York: ACM.

10. Oldfield, R. C. (1971). The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia, 9*, 97–113.

11. Schutz, R. W., & Roy, E. A. (1973). Absolute error: The devil in disguise. *Journal of Motor Behavior, 5*, 141–153.

12. Walker, N., Meyer, D. E., & Smelcer, J. B. (1993). Spatial and temporal characteristics of rapid cursor positioning movements with electromechanical mice in human-computer interaction. *Human Factors, 35*, 431–458.

13. Welford, A. T. (1968). *Fundamentals of skill.* London: Methuen.

14. Zhai, S., Buxton, W., & Milgram, P. (1994). The "silk cursor": Investigating transparency for 3D target acquisition. *Proceedings of the CHI '94 Conference on Human Factors in Computing Systems* (pp. 459–464). New York: ACM.