Behaviour & Information Technology, vol. 24, no. 5 (2005), 337-346.
Preprint version

TouchGrid: Touchpad pointing by recursively mapping taps to smaller display regions

Morten Hertzum
Computer Science
Roskilde University
Roskilde, Denmark
mhz@ruc.dk

Kasper Hornbæk
Department of Computing
University of Copenhagen
Copenhagen, Denmark
Kash@diku.dk

Abstract. Touchpad devices are widely used but lacking in pointing efficiency. The TouchGrid, an instance of what we term cell cursors, replaces moving the cursor through dragging the finger on a touchpad with tapping in different regions of the touchpad. The touchpad regions are recursively mapped to smaller display regions and thereby enable high-precision pointing without requiring high tapping precision. In an experiment, six subjects used the TouchGrid and a standard touchpad across different numbers of targets, distances to targets, and target widths. Whereas standard touchpad operation follows Fitts' law, target selection time with the TouchGrid is a linear function of the required number of taps. The TouchGrid was significantly faster for small targets and for tasks requiring one tap, and marginally faster for two-tap tasks. Error rates tended to be higher with the TouchGrid than the standard touchpad. All subjects preferred the TouchGrid.

1 Introduction

Pointing is a fundamental, low-level operation performed repeatedly in direct manipulation interfaces. Thus, fast and accurate pointing devices are of considerable importance to users' overall task performance and to their subjective experience of a system. While especially mice but also joysticks, touchpads, trackballs, and other pointing devices have been studied extensively (e.g. Card et al. 1978, Albert 1982, Epps 1986, Karat et al. 1986, Cohen et al. 1993, MacKenzie and Oniszczak 1998, MacKenzie et al. 2001), none of these studies seem to evaluate the small, widely used touchpads of laptop computers. The studies that evaluate touchpads either concern larger touchpads external to the computer or do not report the size of the touchpad.

This study evaluates the TouchGrid, an interaction technique that replaces moving the cursor through dragging the finger along the touchpad surface with taps in different regions of the touchpad (figure 1). The TouchGrid seems particularly suited to small touchpads where performance is hampered by frequent clutching (i.e. lifting the finger from the touchpad surface and repositioning it) during medium and long cursor movements. User performance with the TouchGrid depends on the number of possible targets, rather than the distance moved and the target width. Thus the TouchGrid provides for target selection independently of target size, particularly relevant for high-precision pointing and elderly users. Together with other instances of what we will term cell cursors, the TouchGrid explores an alternative design space compared to interaction techniques whose performance can be modelled with Fitts' law.

Insert figure 1 about here (figures are at the end of the manuscript)

2 Fitts' law and previous work

Fitts' law (Fitts 1954, MacKenzie 1992) establishes that the time required to perform basic aiming movements, such as moving a cursor to a target by means of a mouse, is a function of the distance (D) moved and the width (W) of the target within which the movement must end:

$$MT = a + b \times \log_2(\tfrac{D}{W} + 1) \qquad (1)$$

where a and b are empirically determined constants. Studies generally find that mice and trackballs are superior to joysticks, touchpads, and other pointing devices for which movement time follows Fitts' law (e.g. Card et al.

1978, Epps 1986, Cohen et al. 1993, MacKenzie et al. 2001). However, mice require a flat surface such as a desktop for their operation. The proliferation of laptop computers has created a need for pointing devices that make fewer demands on the surrounding environment. One such device is the touchpad, which has become a standard laptop component. It should be noted that the studies comparing pointing devices have been performed under experimental conditions very dissimilar from the crammed space and unstable support that characterize much use of laptops.

Researchers have attempted to improve touchpad performance by (1) making the touchpad hardware more sophisticated, (2) optimising the control:display gain (C:D gain), and (3) devising new interaction techniques.

Firstly, MacKenzie and Oniszczak (1998) compare two conventional ways of implementing the select operation on touchpads with a pressure-sensing touchpad. On the pressure-sensing touchpad selection occurs when the pressure of the finger exceeds a certain threshold, and this creates both aural and tactile feedback. The pressure-sensing touchpad is 20% faster than lift-and-tap selection and 46% faster than using a physical button for selection. Thus, better exploitation of the touchpad hardware may lead to substantial performance improvements.

Secondly, the performance of indirect pointing devices can be optimized by adjusting the relationship between the movement of the device and the ensuing movement of the cursor – that is, the C:D gain. Pointing performance appears to be highest for a medium range of C:D gains. At low gains performance is dominated by long movement times; at high gains by high error rates. Yet, some studies find that the effect of C:D gain is negligible (Jellinek and Card 1990), others that it has an appreciable effect (Graham and MacKenzie 1995), and still others are critical of the gain concept (Accot and Zhai 2001). The movement of the pointing device has been suggested as a related measure more readily linked to changes in performance (Graham and MacKenzie 1995, Accot and Zhai 2001). For small touchpads this movement involves clutching, which slows users down (Jellinek and Card 1990).

Thirdly, the performance of pointing devices can be enhanced by interaction techniques that accelerate long movements or ease high-precision pointing. For example, Drag-and-Pop and Drag-and-Pick use the direction of the initial cursor movement to determine a set of likely candidate targets, and temporarily move these targets to the vicinity of the cursor (Baudisch et al. 2003). Cross-Keys and Precision-Handling accomplish high-precision pointing by introducing a two-step process where the user first sets a selection point by clicking close to the target and then fine tunes the pointing operation using extra handles that require more device movement per unit of movement of the selection point (Albinsson and Zhai 2003). These techniques assume that the cursor is a point. The TouchGrid abandons this assumption.

3 Cell cursors

The TouchGrid is the first instance of a family of interaction techniques we term cell cursors. Cell cursors are motivated by the assumption that for many pointing operations users will be more efficient with an interaction technique that depends on the number of possible targets than with one dependent on the distance to and size of the target.

3.1 Basic ideas

Cell cursors build on three basic ideas:

- The cursor overlays a region of the display and divides that region into a number of cells. For cell cursors in general there are no restrictions on the number and shape of the cells. For the TouchGrid the cursor is a grid with at most three rows and three columns. The row heights and column widths are dynamically set so that the cells contain about the same number of objects.
- Each cursor cell is mapped to a discrete user input. For the TouchGrid the cells are mapped to the corresponding regions of the touchpad. Thus, a tap in the lower right region of the touchpad specifies the lower right cell on the display. For cell cursors in general, input is not restricted to the touchpad.
- Specifying a cell that contains multiple objects causes the cursor to be recursively applied to the cell. Specifying a cell that contains one object is equivalent to selecting that object. Hence, the required number of recursions is a logarithmic function of the number of objects.

The TouchGrid relaxes the requirements for pointing precision by exploiting that the user need not hit the target object, but merely specify a cell sufficiently small to disambiguate which object is being selected. Area cursors (Zhai et al. 1994, Kabbash and Buxton 1995) are also represented by a region (or volume), rather than a point. However, area cursors have a fixed size and are moved around in the same way as conventional point cursors. Zoomable user interfaces (Perlin and Fox 1993, Hornbæk et al. 2002) implement the recursive specification of smaller subsections of, for example, a map. However, zoomable user interfaces have been used for navigation, not object selection, and they magnify the subsections, whereas cell cursors dynamically reduce the size of the cursor.

3.2 TouchGrid example

An example may clarify the operation of the TouchGrid. When the user opens a new window the TouchGrid cursor covers the entire window to enable selection of any object in the window (figure 2a). The cursor is represented as a 3x3 grid, which partitions the screen objects into different cells. In the figure the centre cell contains no objects and attempting to choose it, by tapping the centre of the touchpad, will have no effect. The user, wishing to select the dark circle, taps the middle left part of the touchpad, and because the middle left cursor cell contains multiple objects the grid is recursively applied to the cell. The resulting cursor is a 3-row grid (figure 2b). At this stage dividing the grid into multiple columns would merely result in numerous empty cells. Consequently, the user need merely tap the middle third in a vertical division of the touchpad; the horizontal position of the tap is of no consequence. The result of this second tap is another 3-row cursor (figure 2c). Now the cell with the target object contains no additional objects, and the target object can be selected by tapping the top third of the touchpad. This third tap fine tunes the cursor movement but does not require finer tapping precision than the other taps. Thus, the target object is selected from among the 48 objects by tapping the middle left, the middle, and finally the top of the touchpad.

If users drag their finger along the touchpad surface the TouchGrid disappears and the touchpad operates as a standard touchpad. To resume TouchGrid operation the user clicks any empty region of the display and the TouchGrid reappears fully zoomed out (as in figure 2a). This smooth integration of TouchGrid and standard touchpad also provides the means for undoing erroneous TouchGrid taps. The user simply performs a drag-and-tap stroke and thereby turns the TouchGrid off and back on.

Insert figure 2 about here (figures are at the end of the manuscript)

3.3 Target selection time

For keyboard-operated pointing techniques Card et al. (1978) find that target selection time is a linear function of the number of keystrokes required to move the cursor to a target. Analogously, we assume that target selection time with the TouchGrid is a linear function of the number of taps (N) required to select a target:

$$T_{TouchGrid} = a + b \times N \qquad (2)$$

where a and b are empirical constants to be determined by regression analysis. The reciprocal of the regression line slope (1/b) gives the tapping rate, which reflects the time required for both mental preparation and physically performing a tap.

3.4 Other cell cursors

The TouchGrid combines a touchpad and a grid-shaped cell cursor. However, some devices may not have a touchpad, or characteristics of users or tasks may necessitate other input techniques. For grids with three rows of three cells alternative input techniques include mapping the grid to the 3x3 digit keys (digits 1 through 9) on the numeric keypad. Here zooming out/undoing can be mapped to the 0 key. Cell phones have the same arrangement of digit keys, and cell cursors may be suited to their new graphical displays with objects in two dimensions (e.g. Robbins et al. 2004).

4 Experimental method

To investigate cell cursors empirically, we conducted an experiment comparing the TouchGrid with a standard touchpad. The experiment employed a two (interaction techniques) times three (number of targets) times two (distances to targets) times two (target widths) within-subjects design with six subjects performing 1440 target selections each.

4.1 Subjects

The experimental subjects were 4 women and 2 men, ranging in age from 21 to 56 years with an average of 31.7 years. All subjects had used a touchpad before, one using it daily and two subjects using it only on previously owned laptops. Five subjects used their right hand to control the touchpad; one subject used his left hand.

4.2 Tasks

The basis for the experimental tasks is the multi-directional tapping test (ISO 9241-9 2000), in which subjects are required to select a sequence of targets arranged in a circle (see figure 2). The target to be selected by the subject is highlighted, and the targets are sequenced so that the distance from one highlighted target to the next is nearly equal to the diameter of the circle. The reasons for choosing the multi-directional tapping test are twofold. First, this test has been widely used in evaluations of both non-touchpad and touchpad input. Second, we intend mainly to investigate situations with a limited number of targets to choose from, such as on desktops or in dialog boxes, of which the test appears to be representative.

The multi-directional tapping test was varied on three dimensions expected to influence performance:

- *Number of targets (three levels)*. 8, 24, or 48 targets were equally spaced around the circumference of a circle. For the TouchGrid we hypothesize that target selection time depends on the number of taps required. Whereas one tap is sufficient to select a target in the 8-target tasks, two taps are required in the 24-target tasks, and the 48-target tasks require two or three taps depending on which target the user is to select. According to Fitts' law target selection time with the standard touchpad does not depend on the number of targets.
- *Distance to targets (two levels)*. Targets were arranged in either a small circle with a diameter of 12 cm or a large circle with a diameter of 22 cm. According to Fitts' law, target selection time for the standard touchpad increases with increasing distance to targets. Distance to target is assumed not to affect performance with the TouchGrid.
- *Target width (two levels)*. Small targets had a diameter of 2 mm and large targets a diameter of 6 mm. According to Fitts' law, target selection time for the standard touchpad increases with decreasing target width. Target width is assumed not to affect performance with the TouchGrid.

The combination of the above dimensions gives the 12 tasks used in the experiment.

4.3 Design

The experiment employed a within-subjects design where all subjects used both interaction techniques to complete three blocks of 12 tasks. Each task consisted of 20 trials. Half of the subjects used the TouchGrid for the first half of the session and the standard touchpad for the second half of the session. The other half of the subjects used the standard touchpad first, then the TouchGrid. For each of the three blocks of tasks, subjects received the 12 tasks in random order. The number of target selections performed by a subject totalled:

2 interaction techniques x
3 blocks x
12 tasks x
20 trials =
1440 target selections per subject

4.4 Procedure

The experiment was administered individually for each subject. After a brief presentation of the experiment, subjects filled out a background questionnaire. Next, they were explained the standard touchpad (with which all subjects had some experience) and the TouchGrid (which was new to them). Then, subjects tried the experimental software on some sample tasks. Subjects were instructed to work as quickly as possible, while maintaining high accuracy. They were also told to use only one hand for completing the experiment and were offered a hand rest. Pilot testing showed that the hand rest consistently improved performance.

Within tasks the next target appeared as soon as the previous target had been correctly selected. Subjects could not proceed until the correct target had been selected. Between blocks the subjects could relax for a moment before they started the next block. Upon completing the three blocks for an interaction technique, subjects filled out a questionnaire based on the ISO 9241-9 standard (using the modifications given by Douglas et al. 1999). After performing with both interaction techniques, subjects completed a comparison of them, using questions adapted from Douglas et al. (1999).

All experimental sessions were conducted on the same 1 GHz Dell laptop with a 6x4.5 cm Synaptics touchpad and a 15-inch screen. The C:D gain was set at the middle value in Windows 2000. A test application presented the tasks to the subjects and logged their input. On average, the experimental sessions took 1.25 hours to complete.

5 Results

Below we use repeated measures analysis of variance to analyse the data obtained. Because the first selection in each task requires that subjects move their hands from the enter key to the touchpad, we removed those trials from the below analysis, leaving 19 trials per task in each block.

5.1 Learning effect

Figure 3 shows the change over blocks for target selection time in error-free trials and for error rates. Helmert contrasts indicate a significant difference between all blocks (all with $p<0.001$), so the remainder of the analysis is done using only the data from block 3.

Insert figure 3 about here (figures are at the end of the manuscript)

5.2 Target selection time

Figure 4 summarizes target selection times for error-free trials in block 3. We find interactions of distance to target with interaction technique, $F(1,5)=8.48$, $p<0.05$, and of target width with interaction technique, $F(1,5)=128.00$, $p<0.001$. In both cases, subjects are faster with the TouchGrid.

We also find a significant interaction between number of targets and interaction technique, $F(2,10)=75.37$, $p<0.001$. Planned linear contrasts show that subjects are faster with the TouchGrid for 8 targets, $F(1,5)=39.02$, $p<0.01$, and marginally faster for 24 targets, $F(1,5)=4.92$, $p<0.08$. For 48 targets, subjects are faster with the touchpad, but this is not significant, $F(1,5)=0.66$, $p>0.4$.

As expected, the TouchGrid performs equally well for different distances to targets, $F(1,5)=0.80$, and different target widths, $F(1,5)=1.02$, both with $p>0.3$. Conversely, target selection time with the touchpad does not depend on the number of targets, $F(2,10)=1.36$, $p>0.3$.

Overall, when we compare performance for the 12 tasks, six tasks are significantly faster with the TouchGrid, and no tasks are significantly faster with the touchpad (significance levels use Bonferoni adjustment for multiple comparisons). The only task where the touchpad approaches significantly faster performance is with 48 large targets, close together.

Insert figure 4 about here (figures are at the end of the manuscript)

5.3 Error rates

Data on errors are summarized in table 1. With the touchpad, a trial contains an error if the subject clicks outside of a target. In no trial did a subject in the touchpad condition click on a wrong target. With the TouchGrid, we consider a trial to contain an error if the subjects (1) select a wrong target, (2) tap a cell that contains no selectable object (e.g. the centre cell of figure 2a), or (3) switch to conventional use of the touchpad by dragging their finger on the touchpad surface. These three kinds of errors occur with about the same frequency: 65, 75, and 58 times, respectively. Surprisingly, the impact of errors on target selection time is similar with the touchpad and the TouchGrid. The error rates cover large individual differences; interestingly, one subject in the TouchGrid condition had an average error rate of only 3.94%.

Insert table 1 about here (tables are at the end of the manuscript)

Figure 5 shows the error rates for all trials in block 3. Distance to target does not interact with interaction techniques, $F(1,5)=0.06$, $p>0.9$. We find a significant difference in how the interaction techniques are affected by target width, $F(1,5)=21.92$, $p<0.01$. For small targets, the touchpad and the TouchGrid both have an error rate of 10.82% (SD=3.1). For large targets, however, the error rate with the touchpad is 2.05% (SD=14.2), significantly lower than the 11.26% (SD=31.6) with the TouchGrid, $F(1,5)=57.52$, $p<0.001$.

The error rates for different numbers of targets are illustrative of some aspects of TouchGrid performance. With 8 targets to choose from, subjects make about as many errors with the TouchGrid (M=5.48%) as with the touchpad (M=5.26%). With the TouchGrid subjects only have to tap one TouchGrid cell to complete these trials suggesting that this basic operation is causing subjects some difficulty. With more targets to choose from subjects appear to make more errors with the TouchGrid. The difference in error rates is, however, not significant overall, only for 24 targets, $F(1,5)=31.77$, $p<0.01$.

Overall, error rates with the TouchGrid are high, also for what appear to be relatively simple interactions such as selecting one among eight targets.

Insert figure 5 about here (figures are at the end of the manuscript)

5.4 Subjective satisfaction

Table 2 shows the results of the questionnaires the subjects filled out after using each interaction technique. A Wilcoxon test shows an overall significant difference between interaction techniques, $z=-5.04$, $p<0.001$. Nine of the questions have scales with the most desirable rating at one end and the least desirable rating at the other end (questions 2, 5, 7-13). For these questions, subjects favour the TouchGrid in 33 pairs of ratings; 19 pairs have tied scores; and the touchpad is favoured in only two out of the 6*9=54 pairs. As the sample size is small, this solid satisfaction with the TouchGrid shows itself with significant differences in only two individual questions. The TouchGrid was rated significantly easier for accurate pointing, $z=-2.22$, $p<0.05$, and subjects rated the TouchGrid as causing less shoulder fatigue, $z=-2.12$, $p<0.05$.

Insert table 2 about here (tables are at the end of the manuscript)

On comparing the interaction techniques at the end of the experiment, all subjects preferred the TouchGrid. Two subjects considered the TouchGrid easier to use, and two subjects felt that the TouchGrid was superior especially for small targets, a feeling reflected in the target selection times.

5.5 TouchGrid performance

As suggested above, error rates are high with the TouchGrid. One reason may be that subjects find it difficult to tap in particular touchpad regions. Figure 6 shows the distribution of correct taps (dots) and error taps (asterisks) over the touchpad surface when subjects used the TouchGrid. The data for the left-handed subject was mirrored around the vertical axis before being displayed (similar mirroring is done in all analyses below).

Insert figure 6 about here (figures are at the end of the manuscript)

One characteristic of the subjects' tapping is that they mix up rows more often than columns. As can been seen from the figure errors are often at the border between cells 6 and 9, the border between cells 5 and 8, or the border between cells 1 and 4. On average, 43% of the errors are row errors, 18% are column errors, 10% are diagonal errors (e.g. tapping in cell 7 instead of cell 3), and 29% are other errors.

The figure also shows that taps in the middle column are clearly shifted right. This shift probably reflects that subjects operated the touchpad with their hand approaching from the right, making an approximately 45 degree angle with the laptop. However, the only column error the subjects make frequently is between cells 2 and 3 (7.6% of the total number of tapping errors).

Upon analysing TouchGrid performance for 8-target circles, tapping in cell 4 stands out with a mean time of

856ms (SD=218) and zero errors, probably reflecting the common positioning of fingers over the touchpad with the index finger over cell 4. The cell taking the longest time to select is cell 7, with M=1233ms (SD= 836). In terms of errors, cell 7 is also relatively high (M=8.82%, compared to the 5.26% mean error rate for 8-target circles).

Another determinant of TouchGrid performance is the subdivision of the cursor into cells. With our current algorithm for making the cells some subdivisions of the 24-target circles have three cells (non-corner cells, figure 7a) and others have nine cells (corner cells, figure 7b). It takes longer to select targets in corner cells (M=2000ms; SD=674) compared to targets in non-corner cells (M=1828ms; SD=709), $t=-2.43$, $p<0.05$. Corner cells also have more errors (M=15.5%) than non-corner cells (M=12.6%), especially errors where subjects tap an empty cell. This kind of error accounts for 28% of the errors in trials with the target in corner cells of 24-target circles.

With the 48-target circles, one kind of cell division stands out as particularly slow. The series of taps in figure 7c illustrate the subjects' problems: after selecting for example the top-right corner cell in a circle by tapping cell 9 on the touchpad (leading to figure 7d), and then the lower right cell (cell 3) subjects are faced with the selection problem in figure 7e. If they tap the top-left cell, they have tapped the sequence 9, 3, and 7. The direction change from 3 to 7 crosses the centre of the touchpad and seems hard. Mean selection time for trials that require this kind of change is M=3100ms (SD=1150), while for corner cells without such direction changes (e.g. the cell at the bottom right of figure 7e, selected by tapping the sequence 9, 3, 3) the mean selection time is 2601ms (SD=731), $t=-3.41$, $p<0.001$. Subjects also make more errors in trials involving a direction change (M=13.5%) than in those without such a change (M=8.33%). In the post-experiment comments, one subject explicitly mentioned these changes as particularly challenging.

Insert figure 7 about here (figures are at the end of the manuscript)

5.6 Predictive performance models

Performance in the touchpad condition may be modelled with Fitts' law (Fitts 1954, MacKenzie 1992). Using the mean of error-free trials for each distance/width combination we get, by linear regression on equation 1:

$$T_{touchpad} = -657 + 469 \times \log_2\left(\frac{D}{W} + 1\right) \text{ ms} \qquad (3)$$

where D is distance and W is width. The regression line slope of 469ms per bit is comparable to other studies of touchpads (Epps 1986, MacKenzie and Oniszczak 1998, Douglas et al. 1999). For this model, r^2 is 0.932.

For modelling TouchGrid performance, we estimate the constants in equation 2 by linear regression on the mean times for each of the 12 tasks. We obtain the equation:

$$T_{TouchGrid} = 212 + 842 \times N \text{ ms} \qquad (4)$$

where N is the number of taps required to select a target. The 842ms per tap corresponds to keying times for 10-key pads such as telephones (Card et al. 1983). For this model, r^2 is 0.986.

An alternative model, using the Hick-Hyman law ($RT = a + b \times \log_2(n)$; Hick 1952, Hyman 1953), can be based on the number of cells users choose among when performing a tap. For example, in figure 2 the user first selects among nine, then three, and finally three cells. Simplifying the model by using the same constant for choices at different recursions, the model is:

$$RT = -1251 + 1946 \times \log_2\left(\prod_{i=1}^{r} \text{cells}_i\right) \text{ ms} \qquad (5)$$

where cells_i refers to the number of cells to choose from at recursion i, and r is the number of recursions. The negative intercept possibly arises because the minimum number of cells to choose from in this experiment is nine. This model also has a good fit with data, $r^2=0.948$. The number of required taps and the base-two logarithm of the number of choices between cells are highly correlated ($r=0.90$) because both numbers are related to the number of screen objects. We prefer equation 4 because it is simpler and offers a straightforward interpretation of the regression coefficients.

6 Conclusion

The TouchGrid maps taps in different touchpad regions to still smaller display regions and thereby provides for object selection by recursively homing in on target objects. Target selection time with the TouchGrid is a linear function of the number of taps required. Thus, the TouchGrid seems particularly suited to situations in which the number of targets is fairly low, for example because the active context is a dialog box or another fraction of the entire display. In such situations target selection can often be accomplished in a few taps, possibly just one.

Compared to standard touchpad operation, the results of our experiment indicate that the TouchGrid was significantly faster for one-tap tasks and marginally faster for two-tap tasks. The TouchGrid was also significantly faster for tasks involving small targets. Conversely, the touchpad was not significantly faster for any of the 12 tasks and only approached significance for a task with 48 large targets, close together. Subjects unanimously preferred the TouchGrid, and they rated it significantly easier for accurate pointing and significantly less straining on their shoulders. In terms of error rates the TouchGrid was, however, inferior to the standard touchpad.

The analysis of the subjects' use of the TouchGrid suggests several possibilities for improvement. One is to reduce the grid to two rows or have cells overlap slightly because many erroneous taps mix up rows or otherwise appear close to cell borders. Another improvement suggested by our analysis is to avoid empty cells, possibly by abandoning rectangular cells in favour of, for example, a tripartition of upper right corner cells into an upper left triangle, a lower right triangle, and a diagonal middle belt. Finally, the TouchGrid should be evaluated with other tasks besides the multi-directional tapping test to exercise the TouchGrid on less regular distributions of targets and analyse how it scales up to environments with large numbers of possible targets. The present evaluation of the TouchGrid suggests a promising interaction technique that evades Fitts' law by being independent of distance to target and target size.

Acknowledgements

Morten Hertzum was supported, in part, by the IT University of Copenhagen. Niels Ebbe Jacobsen provided valuable comments on a draft version of the manuscript. Special thanks are due to the people who participated in the experiment as subjects.

References

Accot, J., and Zhai, S., 2001, Scale effects in steering law tasks. *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 1-8.

Albert, A.E., 1982, The effect of graphic input devices on performance in a cursor positioning task. *Proceedings of the Human Factors Society 26th Annual Meeting* (Santa Monica, CA: Human Factors Society), pp. 54-58.

Albinsson, P.-A., and Zhai, S., 2003, High precision touch screen interaction. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 105-112.

Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P., Bederson, B., and Zierlinger, A., 2003, Drag-and-pop and drag-and-pick: techniques for accessing remote screen content on touch- and pen-operated systems. *Proceedings of the IFIP TC13 INTERACT 2003 Conference on Human-Computer Interaction* (Amsterdam: IOS Press), pp. 57-64.

Card, S.K., English, W.K., and Burr, B.J., 1978, Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. *Ergonomics*, 21(8), 601-613.

Card, S.K., Moran, T.P., and Newell, A., 1983, *The Psychology of Human-Computer Interaction* (Hillsdale, NJ: Erlbaum).

Cohen, O., Meyer, S., and Nilsen, E., 1993, Studying the movement of high-tech. rodentia: pointing and dragging. *Proceedings of the INTERCHI'93 Conference on Human Factors in Computing Systems: Adjunct Proceedings* (New York: ACM Press), pp. 135-136.

Douglas, S.A., Kirkpatrick, A.E., and MacKenzie, I.S., 1999, Testing pointing device performance and user assessment with the ISO 9241, part 9 standard. *Proceedings of the CHI 99 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 215-222.

Epps, B.W., 1986, Comparison of six cursor control devices based on Fitts' law models. *Proceedings of the Human Factors Society 30th Annual Meeting* (Santa Monica, CA: Human Factors Society), pp. 327-331.

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.

Graham, E., and MacKenzie, C.L., 1995, Pointing on a computer display. *Proceedings of the CHI 95 Conference on Human Factors in Computing Systems: Conference Companion* (New York: ACM Press), pp. 314-315.

Hick, W.E., 1952, On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, 4, 11-26.

Hornbæk, K., Bederson, B.B., and Plaisant, C., 2002, Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Transactions of Computer-Human Interaction*, 9(4), 362-389.

Hyman, R., 1953, Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, 45, 188-196.

ISO 9241-9, 2000, *Ergonomic Requirements for Office Work with Visual Display Terminals (VDT's) – Part 9: Requirements for Non-Keyboard Input Devices* (Geneva: International Organization for Standardisation).

Jellinek, H.D., and Card, S.K., 1990, Powermice and user performance. *Proceedings of the CHI 90 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 213-220.

Kabbash, P., and Buxton, W., 1995, The "Prince" technique: Fitts' law and selection using area cursors. *Proceedings of the CHI 95 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 273-279.

Karat, J., McDonald, J.E., and Anderson, M., 1986, A comparison of menu selection techniques: touch panel, mouse and keyboard. *International Journal of Man-Machine Studies*, 25(1), 73-88.

MacKenzie, I.S., 1992, Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction*, 7(1), 91-139.

MacKenzie, I.S., Kauppinen, T., and Silfverberg, M., 2001, Accuracy measures for evaluating computer pointing devices. *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 9-16.

MacKenzie, I.S., and Oniszczak, A., 1998, A comparison of three selection techniques for touchpads. *Proceedings of the CHI 98 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 336-343.

Perlin, K., and Fox, D., 1993, Pad: an alternative approach to the computer interface. *Proceedings of the SIGGRAPH '93 Conference* (New York: ACM Press), pp. 57–64.

Robbins, D.C., Cutrell, E., Sarin, R., and Horvitz, E., 2004, ZoneZoom: map navigation for smartphones with recursive view segmentation. *Proceedings of the AVI '04 Working Conference on Advanced Visual Interfaces* (New York: ACM Press), pp. 231-234.

Zhai, S., Buxton, W., and Milgram, P., 1994, The "Silk Cursor": investigating transparency for 3D target acquisition. *Proceedings of the CHI 94 Conference on Human Factors in Computing Systems* (New York: ACM Press), pp. 459-464.

	n (%)	*Added time*
Touchpad (non-target clicks)	88 (6.43%)	+2449ms (+114%)
TouchGrid (any error)	151 (11.04%)	+2577ms (+139%)

Table 1. Errors, their frequency, and their impact on target selection time (over average error-free time). Each technique has been used for 1368 trials.

Question	Touchpad	TouchGrid
1. The force required for actuation was (1:too low – 5:too high)	3.33	3.67
2. Smoothness during operation was (1:very rough – 5:very smooth)	2.83	3.50
3. The mental effort required for operation was (1:too low – 5:too high)	3.17	2.80
4. The physical effort required for operation was (1:too low – 5:too high)	3.83	3.17
5. Accurate pointing was (1:easy – 5:difficult)	4.17	2.17 *
6. Operation speed was (1:too fast – 5:too slow)	3.17	2.67
7. Finger fatigue (1:none – 5:very high)	3.33	1.83
8. Wrist fatigue (1:none – 5:very high)	3.33	2.17
9. Arm fatigue (1:none – 5:very high)	2.33	2.00
10. Shoulder fatigue (1:none – 5:very high)	2.83	1.67 *
11. Neck fatigue (1:none – 5:very high)	3.00	2.00
12. General comfort (1:very uncomfortable – 5:very comfortable)	2.50	3.17
13. Overall, the input device was (1:very difficult to use – 5:very easy to use)	2.67	3.83

Table 2. Satisfaction with the interaction techniques (questionnaire from Douglas et al. 1999). Overall, subjects rate the TouchGrid significantly better; significant differences for individual questions are marked with an asterisk.

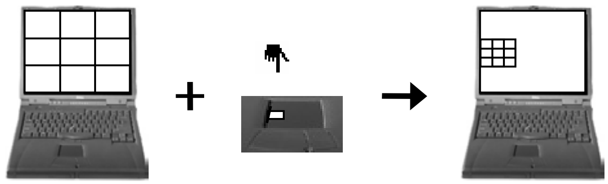

Figure 1. The TouchGrid overlays the active window with a 3x3 grid and maps the grid cells to regions of the touchpad. A tap in, for example, the middle left of the touchpad (the region coloured white above) specifies the middle left cell in the grid. If this cell contains more than one object the grid is recursively applied to the cell, and the user taps once more. If the cell contains one object, this object is selected.

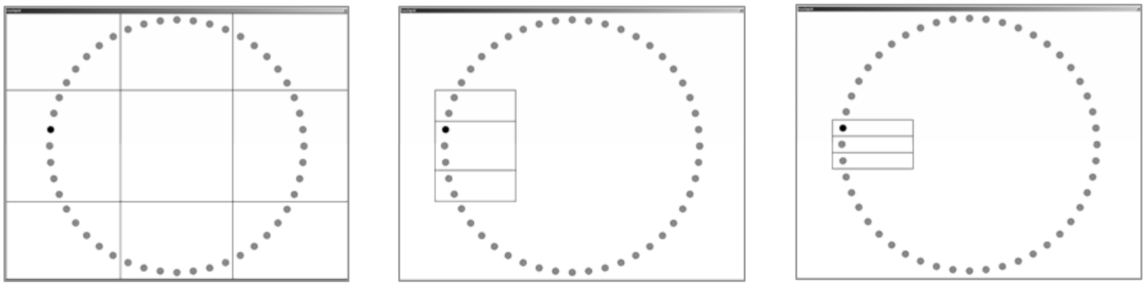

(a) (b) (c)

Figure 2. TouchGrid pointing. Selecting the dark target object involves three taps with concomitant cursor changes: (a) initial appearance of the cursor, (b) after the first tap, and (c) after the second tap. The final selection of the target consists of a third tap, in the top third of the touchpad. Note that the grid lines have been made thicker in the figure to increase their visibility.

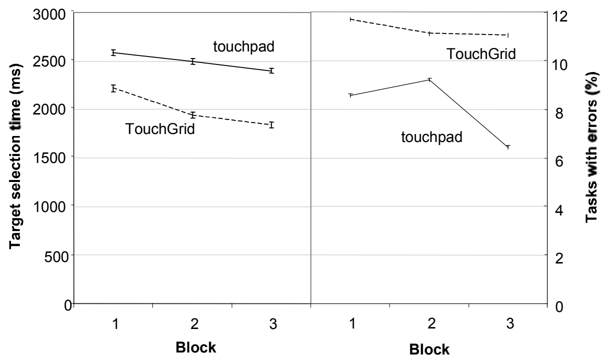

Figure 3. Learning effects. Error bars show standard error of the mean. Diagram to the left summarizes n=7817 error-free trials; diagram to the right summarizes n=8208 trials.

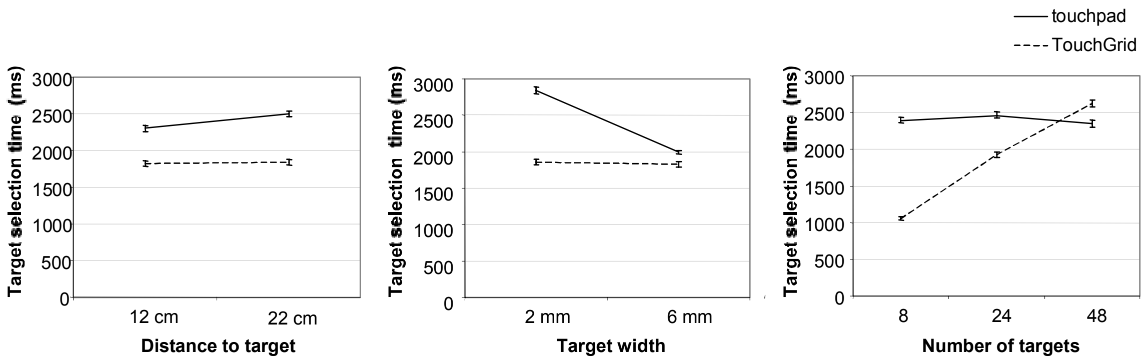

Figure 4. Target selection time. Error bars show the standard error of the mean. Each diagram shows n=2497 error-free trials.

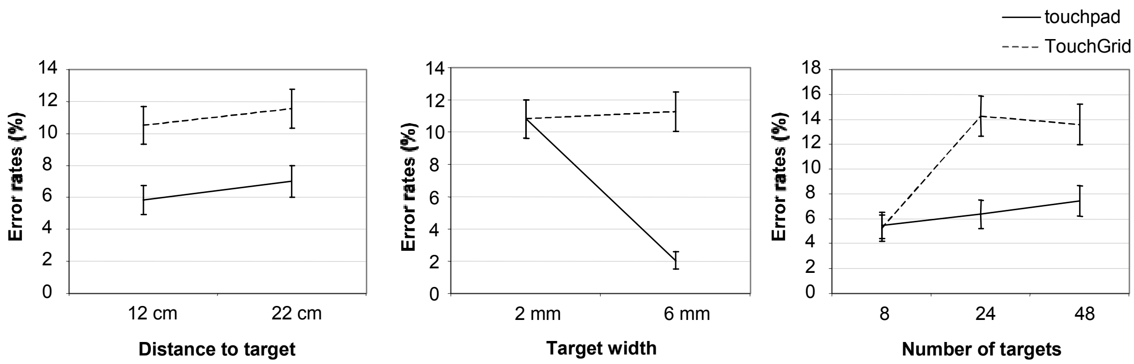

Figure 5. Error rates. Error bars show the standard error of the mean. Each diagram shows n=2736 trials.

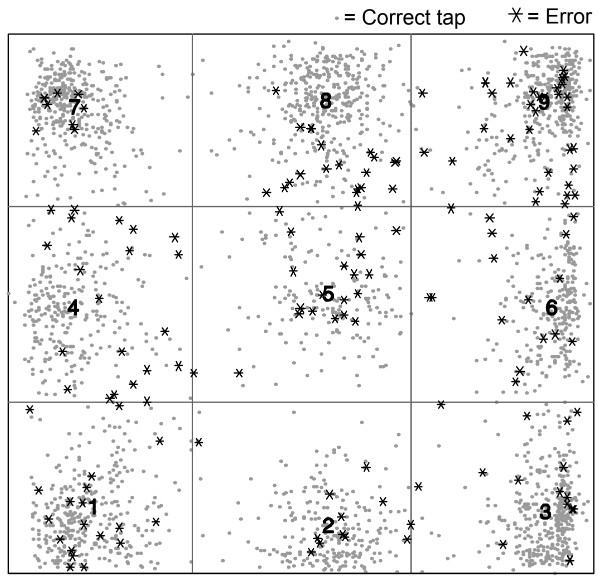

Figure 6. Distribution of taps over the touchpad surface during use of the TouchGrid (all subjects, all trials). The lines indicate the division of the touchpad surface into cells; the numbers are placed at the mass-midpoint of subjects' taps. Note that as a result of pilot testing, we made the middle column and row slightly larger, as subjects appeared to have difficulty hitting them.

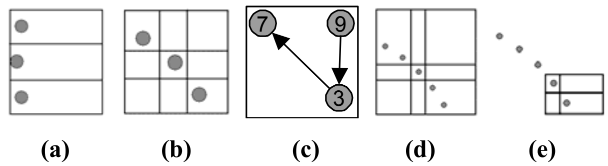

Figure 7. Cell divisions with the TouchGrid. From left to right is shown (a) left side of 24-target circle, (b) top-right corner of 24-target circle, (c) a series of particularly challenging taps on the touchpad, (d) top-right corner of 48-target circle, and (e) selection of lower right cell in corner of 48-target circle.