Two-Handed Input Using a PDA And a Mouse

Brad A. Myers, Kin Pou ("Leo") Lie , and Bo-Chieh ("Jerry") Yang

Human Computer Interaction Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
bam@cs.cmu.edu
http://www.cs.cmu.edu/~pebbles

ABSTRACT

We performed several experiments using a Personal Digital Assistant (PDA) as an input device in the non-dominant hand along with a mouse in the dominant hand. A PDA is a small hand-held palm-size computer like a 3Com Palm Pilot or a Windows CE device. These are becoming widely available and are easily connected to a PC. Results of our experiments indicate that people can accurately and quickly select among a small numbers of buttons on the PDA using the left hand without looking, and that, as predicted, performance does decrease as the number of buttons increases. Homing times to move both hands between the keyboard and devices are only about 10% to 15% slower than times to move a single hand to the mouse, suggesting that acquiring two devices does not cause a large penalty. In an application task, we found that scrolling web pages using buttons or a scroller on the PDA matched the speed of using a mouse with a conventional scroll bar, and beat the best two-handed times reported in an earlier experiment. These results will help make two-handed interactions with computers more widely available and more effective.

Keywords: Personal Digital Assistants (PDAs), Hand-held computers, Palm Pilot, Windows CE, Two-Handed Input, Smart Environments, Ubiquitous Computing.

INTRODUCTION

Many studies of two-handed input for computers have often shown advantages for various tasks [1, 3, 7, 9, 15]. However, people rarely have the option of using more than just a mouse and keyboard because other input devices are relatively expensive, awkward to set up, and few applications can take advantage of them. However, increasing numbers of people now do have a device that they carry around that could serve as an extra input device for the computer: their Personal Digital Assistant (PDA). PDAs, such as 3Com's Palm Pilots and Microsoft's Windows CE devices, are designed to be easily connected to PCs and have a touch-sensitive screen which can be used for input and output. Furthermore, newer PDAs, such as the Palm V and the HP Jornada 420, are rechargeable, so they are *supposed* to be put in their cradles next to a PC when the user is in the office. Therefore, if using a PDA in the non-dominant hand proves useful and effective, it should be increasingly easy and sensible to deploy and configure using hardware devices that users already have.

Submitted for Publication

Another advantage of PDAs over the input devices studied in previous experiments is that they are much more flexible. PDAs have a display on which virtual buttons, knobs and sliders can be displayed, and they can be programmed to respond to a wide variety of behaviors that can be well-matched to particular tasks. However, a disadvantage is that the controls on the PDA screen are virtual, so users cannot find them by feel. Research is therefore needed to assess how well the PDA screen can work as a replacement for other input devices that have been studied for the left hand.

This paper reports on several experiments that measure various aspects of using a PDA as an input device in the non-dominant hand. Two experiments are new and are designed to measure the parameters of using a PDA. One experiment repeats an earlier study but uses a PDA in the non-dominant hand. Since the actual pragmatics of input devices can have a large impact on their effectiveness [2, 8], we wanted to determine whether the results seen in prior experiments would also apply to using PDAs.

In summary, the results are:

- People can quickly and reliably hit large buttons drawn on the PDA with their left hands without looking. 99% of the button taps were correct on buttons that are 1-inch square in a 2x2 arrangement. With a larger number of smaller buttons, the accuracy significantly decreases: 95% were correct for 16 buttons that are ½ inch on a side arranged 4x4. The time from stimulus to button tap was about 0.7 sec for the large buttons and 0.9 seconds for the smaller buttons.

- In a task where the subjects had to move both hands from the keyboard to the PDA and the mouse and then back, we found that it took an average of 0.791 seconds to move both hands from the devices to the keyboard. This was about 13% longer than moving one hand from the mouse to the keyboard (which took 0.701 sec). Moving to a PDA and mouse from the keyboard took an average of 0.838 seconds, which is about 15% longer than moving one hand to the mouse (0.728 seconds).
- In a repeat of the experiment reported in [15], subjects were able to scroll large web pages and then select a link on the page at about the same speed using buttons or a scroller on a PDA compared to using the mouse with a conventional scroll bar. The times we found for scrolling with buttons on the PDA were faster than any of the times in the earlier study, including the 2-handed ones.

RELATED WORK

There have been many studies of using input devices for computers in both hands, but none have tested PDAs in the left hand, and we were unable to find measurements of homing times from the keyboard to devices for two-handed use.

One of the earliest experiments measured the use of two hands in a positioning and scaling task and for scrolling to known parts of a document [3]. This study found that people naturally adopted parallel use of both hands and could scroll faster with the left hand. Theoretical studies show that people naturally assign different tasks to each hand, and that the non-dominant hand can support the task of the dominant hand [6]. This has motivated two-handed interfaces where the non-dominant hand plays a supporting role, such as controlling other drawing tools [9] and adjusting translation and scaling [3, 15]. Other studies have tested two-handed use for 3D interaction [1, 7] and found the extra input to be useful.

There has been prior work on using PDAs at the same time as regular computers for various tasks including meeting support [11], sharing information [12], and to help individuals at their desks [10], but we found no prior work on measuring performance of non-dominant hand use of PDAs.

EXPERIMENTAL DESIGN

Two new studies were performed. In the first, the subjects did five tasks in a row. The first task was a typing test to see how fast the subjects could type. Next, they performed a button size task to measure the error rates and speeds when tapping on different size buttons on the PDA. Next, the subjects performed a homing speed task where we measured the how quickly the subjects moved among the keyboard and the devices. Finally, they performed a scrolling task using a variety of devices, which is a repeat of an earlier experiment [15]. The subjects reported a number of problems with the scrolling devices on the PDA in the last task, so we redesigned the scrolling devices, and in a second study with new subjects, we evaluated the performance of the new scrollers on the same task. Each of these experiments is described below.

Apparatus

Subjects sat at a normal PC desktop computer that was running Windows NT. On the right of the keyboard was a mouse on a mouse pad. On the left was an IBM WorkPad 8602-30X PDA (which is the same as a Palm IIIx). In the first study, we put the WorkPad in its cradle. Subjects complained that the WorkPad was wobbly it its cradle, so for the second study, the new subjects used a WorkPad resting on a book and connected by a serial cable to the PC. There were no further comments about the positioning.

The WorkPad has a 3¼-inch diagonal LCD display screen (about 2 ¼ inches on a side) which is touch sensitive. It is 160x160 pixels. Figure 1 shows a picture of the full WorkPad.

The software running on the WorkPad was the Shortcutter program [10] that allows panels of controls to be created so that each control sends specified events to the PC. The software on the PC consisted of various applications specifically created for this experiment (except in the scrolling task, which used the Netscape browser running a custom JavaScript program to collect the data).

Typing Test

We used a computerized typing test called "Speed Typing Test 1.0" [14]. The subjects were asked to type a paragraph displayed on the screen as quickly as possible.

Button Size Task

In this task, the PDA displayed between 4 and 16 buttons in eight different arrangements: 2 rows by 2 columns, 2x3, 3x2, 2x4, 4x2, 3x4, 4x3, and 4x4 (see Figure 1). To control for ordering effects, half of the subjects used the order shown above, and the other half used the reverse order (4x4 first down to 2x2 last). In the 2x2 condition, the buttons were about one inch square, and in the 4x4, they were about ½ inch square.

At the beginning of each condition, a picture was displayed on the PC screen showing the corresponding layout of the buttons (with the same size as the PDA). Then one of the buttons was shaded black (see Figure 2). The subjects were asked to tap on the corresponding button on the WorkPad as quickly and accurately as possible with a finger on their left hand. The stimulus button was then cleared on the PC and the next stimulus button was shaded 500 milliseconds later. The stimuli appeared in random order. A total of 48 stimuli were used in each condition. Every button appeared the same number of times. For example, for the layout of 2 rows by 2 columns, each button appeared 12 times, while for the layout of 3 rows by 4 columns, each button appeared 4 times. There was a break after each condition. Our hypotheses were that people could more accurately select among fewer, larger buttons, and that people could make selections without looking at the WorkPad.

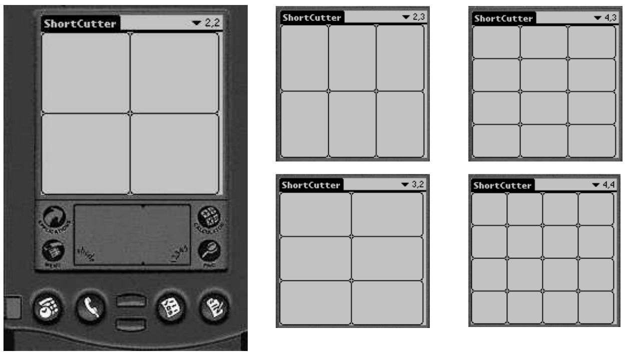

Figure 1. Left: a picture of a Palm Pilot (the WorkPad is similar) showing the 2x2 layout of buttons. Right: the screens for 3x2, 2x3, 4x3 and 4x4. The other layouts are similar.

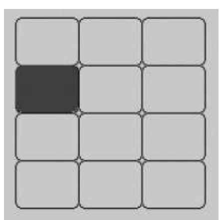

Figure 2. Part of the PC screen showing the stimulus during the 4x3 condition of the button task.

Homing Speed Task

The purpose of this task was to measure the times to move the hands back and forth from the keyboard to the mouse and WorkPad as the subjects switch between two-handed selection operation and two-handed typing. We compared moving a single hand to and from the keyboard to moving both hands.

There were three conditions with three trials in each. In each trial, 14 textboxes were shown on the screen with a label in front of each. The conditions were that the subjects had to first select a text box by either clicking in the field with the mouse in the usual way, tapping on a full-screen button on the WorkPad (which therefore worked like a "TAB" key and caused the cursor to jump to the next field), or tap on the WorkPad and click the mouse at the same time. In other words, the selection operation in this last condition was like a "Shift-Click" operation in which the button on the WorkPad was treated as a Shift key. After the textbox was selected, the subjects typed the word indicated on the left of the textbox. The word was either "apple" or "peach" (in alternating order). These words were chosen because they are easy to type and remember, and they start and end with keys that are under the opposite hands. The user could not exit the field until the word was typed correctly. After typing the word correctly into the textbox, the subject then continued to perform the same selection-typing operation in the next textbox. The trial ended when all 14 textboxes on the screen were filled in. There was a break after each trial. We measured the time from the mouse and WorkPad click to the first character typed, and from the last character typed to the first movement of the mouse or tap on the WorkPad. We did not count the time spent actually typing, and we eliminated the times for the first and last words, because they were biased by start-up and transients.

We hypothesized that moving to the WorkPad and the mouse would not take much longer than moving one hand since people would move both hands at the same time. We were also interested in the actual numbers for the time measurements. These might be used with future models of human performance for two-handed homing tasks.

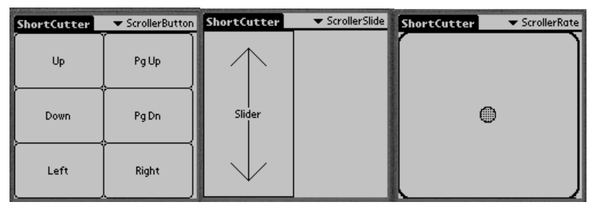

(a) (b) (c)

Figure 3. (a) The button scroller on the WorkPad used in the first experiment. (b) The Slide Scroller and (c) Rate scroller used in both experiments.

Scrolling Task

For this task, we were able to replicate the conditions of a previous experiment [15] exactly.[1] The purpose of this task was to evaluate and compare subjects' performance in scrolling web pages in a standard browser using different scrolling methods. The web pages contain text from an IBM computing terminology dictionary, and each page is about 12 screen-fulls of text. In each web page a hyperlink with the word "Next" is embedded at an unpredictable location. The subjects were asked to find the target hyperlink by scrolling the web page using the different scrolling mechanisms. Once the link was visible, they used the mouse in the conventional way to click on it. Clicking on the hyperlink brought the subject to the next web page. For each condition, the subjects first performed a practice run of 10 pages, during which they were asked to try out the scrolling method without being timed. Then, the subjects did two consecutive trials of 10 pages each as fast as they could.

1 Thanks very much to Shumin Zhai of IBM for supplying the experimental material from the earlier study.

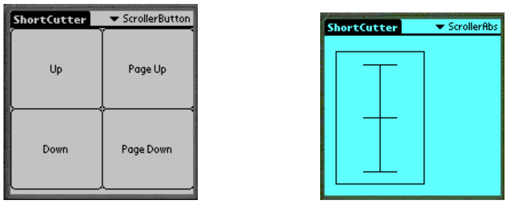

(a) (b)

Figure 4. (a) The revised button scroller on the WorkPad used in the second experiment. (b) The "absolute scroller".

The condition with the fastest time in the previous experiment used a "pointing stick" joystick to scroll, but we were not able to reproduce this condition.[2] The conditions we used in our first experiment were:

- Scrolling using the mouse and the regular scroll bar.
- Scrolling using a "scroll wheel" mounted in the center of the mouse (a Microsoft "IntelliMouse"). We were careful to explain to the subjects the three different ways the wheel can be used, including rolling the wheel, or tapping or holding the wheel down to go into "scroll mode" where the further you move the mouse from the tap point, the faster the text scrolls. The subjects could choose which methods to use.
- Scrolling using buttons on the WorkPad (see Figure 3a). There were 6 buttons that scrolled up and down a line, up and down a page, and left and right (which were not needed for this experiment). The buttons auto-repeated if held down.
- Scrolling using a "slider" on the WorkPad (see Figure 3b). Putting a finger on the slider and moving up or down moved the text the corresponding amount. If you reach the edge of the slider, then you need to lift your finger and re-stroke. Tapping on the slider has no effect since only relative movements are used.
- Scrolling using a "rate scroller," which acted like a rate-controlled joystick with three speeds (see Figure 3c). Putting a finger on the WorkPad and moving up or down started the text moving in that direction, and moving the finger further from the start point scrolled faster.

The order of the conditions was varied across subjects.

Revised Scrolling Task

We received a number of complaints and suggestions about the scrollers on the WorkPad in the first session, so we redesigned some of them and repeated the scrolling task in a second study with new subjects. In this study, we only used four buttons for the button scroller (since the left and right buttons were not needed—see Figure 4a). We also tried to improve the rate scroller, by adjusting the scroll speeds and the areas where each speed was in affect. Finally, we added a new (sixth) condition:

- Scrolling using an "absolute scroller," where the length of the scroller represented the entire document, so putting a finger at the top jumped to the top of the document, and the bottom represented the bottom (see Figure 4b). The user could also drag up and down to scroll continuously. Therefore, it was as if the scroll bar's indicator was attached to the finger. The motivation for this scroller was that we noticed that most people in the mouse condition of the first session dragged the indicator of the scroll bar up and down, and we wanted to provide an equivalent WorkPad scroller.

Subjects

There were 12 subjects in the first study, which took about an hour and they were paid $15 for participating. 12 different subjects did the second study, which took about ½ hour and they were paid $10. All subjects were Carnegie Mellon University students, faculty, or staff. 25% (6 out of 24) were female, and the age ranged from 19 to 46 with a median of 26. All were moderately to highly proficient with computers, and half had used PDAs. The data from some extra subjects were eliminated due to technical difficulties. The measures from two subjects who were left-handed are not included in the data, but informally, their numbers did not look different.

RESULTS

General

Pearson product-moment correlation coefficient between typing speed and tap speed in the button size task (namely the mean tap speed across all 8 layouts) was .60, which means the faster typists were somewhat faster at tapping. The correlation coefficient between typing speed and the speed for moving one hand to the keyboard in the homing task was .79 which means, as expected, subjects who were better typists could put their hands in the home position more quickly. There was little correlation of typing speed to the other measures in the homing task. The correlation coefficient between typing speed and scrolling speed (in the revised scrolling task) across all 6 conditions and both trials was 0.32, which means there was little correlation for the scrolling task.

Age and gender did not affect the measures.

Button Size Task

Figure 5 shows the times to tap on the button measured from the time the stimulus appeared on the PC monitor. These numbers only include correct taps. There were two orders for the trials, so each condition was seen by some subjects early in the experiment, and by other subjects later. The chart presents the data for the early and late cases along with the average of both.

[2] We did not have a pointing stick to test, and anyway, it would have been difficult to connect one to the computers we had, which illustrates one of the claims of this paper—it can be difficult to connect multiple conventional input devices to today's computers. Since the experimental set up was identical to the original experiment [15], it should be valid to compare our times with the times reported there.

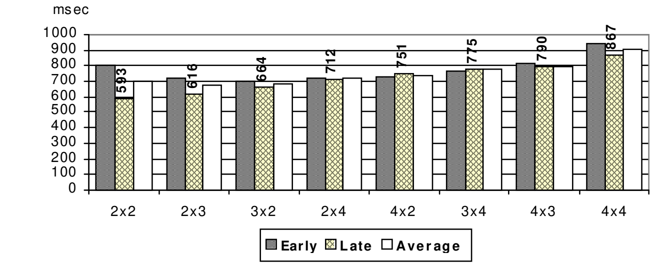

Figure 5. Times to tap each button depend on the size. The times are shown for the subjects who saw each condition later.

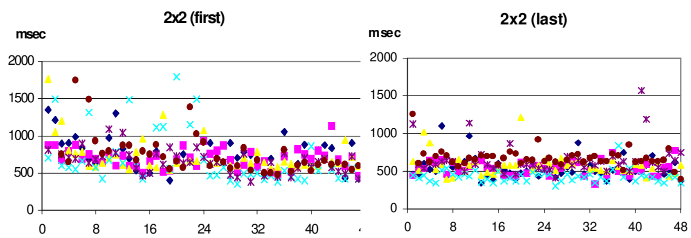

Figure 6. Plot of all times for the 2x2 layout shows (on the left) learning happening for those subjects who saw this condition first, but not (on the right) for those who saw it last.

Figure 6 shows the times to tap on a button in the 2x2 trial for each of the buttons for each of the subjects. The left graph is of those subjects who saw the 2x2 condition first, and roughly matches the power law of practice. However, for those subjects who did the 2x2 condition last, there was no apparent learning during that trial, and the times are flat. Therefore, we feel it is more valid to use the times from only the subjects who saw the condition later. The average time for just the second set is 593msec.

As shown in Figure 5, and predicted by Fitts's law [5, p. 55], the time to tap on a button is inversely proportional to the size of the button, ranging from 593 msec in the 2x2 condition to 867 msec in the 4x4 (for those the subjects who saw each condition later).

The times to tap differ significantly among different numbers of buttons ($F_{4,40}$=16.8, p<.001). There is significant interaction between button number and order of conditions (2x2->4x4 or 4x4->2x2) ($F_{4,40}$=6.0, p=.001), but the learning effect is most prominent among layouts with small number of buttons. The Tukey Test at .05 significant level indicates that there is no significant difference between the 4-button condition and the 6-button condition, between the 6-button and 8-button, or between the 8-button and 12-button. However, the 12-button condition is faster than the 16-button condition by a statistically significant margin.

The times for different layouts of the same number of buttons is not statistically significant, however: the Tukey Test at .05 significant level indicates that times for the 2×3 are not statistically different from 3×2, 2×4 compares to 4×2, and 3×4 compares to 4×3.

Figure 7 shows the error rates for the various configurations, which varies from 1.04% to 4.17% for the subjects who saw each condition later. The error rates do not differ significantly among different layouts ($F_{7,70}$=1.6, p=.14) nor among different numbers of buttons ($F_{4,40}$=2.1, p=.07). For the 4×4 layout, 45% of the errors were in the wrong row, 48% were in the wrong column, and 7% were wrong in both (on the diagonal from the correct button). There was no consistent pattern of where the problematic buttons were located (see Figure 8).

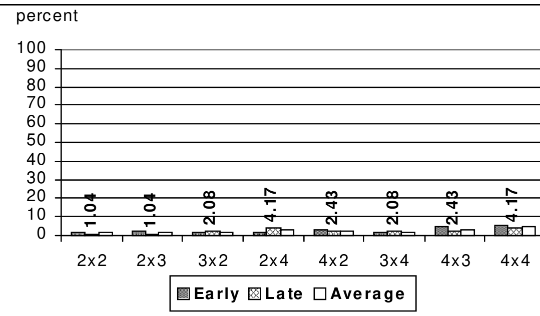

Figure 7. Error rates for each condition of the button task. Numbers shown are for the subjects who saw each later.

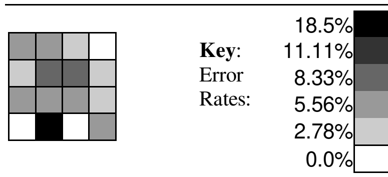

Figure 8. Percent of the taps in each button that were in error in the 4x4 layout.

Homing Speed Task

Figure 9 shows the times for moving each hand in the various conditions of the homing speed task. When moving only one hand at a time (top 4 rows), the subjects took 728 msec to move to the mouse and 701 msec to move back to the keyboard from the mouse. The times to move to the PDA were 744 msec to the PDA and 639 back.

When required to move both hands, the subjects took only slightly longer, requiring about 15% more time to acquire both the PDA and the mouse (838msec), and about 12% more time to acquire the keyboard (791 msec).

1H Keyboard->Mouse	728	
1H Keyboard->PDA	744	
1H Mouse->Keyboard	701	
1H PDA->Keyboard	639	
Keyboard -> Mouse&PDA	838	15.1%
Mouse&PDA -> Keyboard	791	12.8%

Figure 9. Times in milliseconds to move hands. "1H" means when only one hand is moving. The third column shows the percent slowdown of moving both hands compared to the corresponding one-handed mouse time.

Scrolling Task

As in the study we reproduced [15], the time for the first trial with each input device was for practice, so Figure 10 shows the times for the second and third trials.

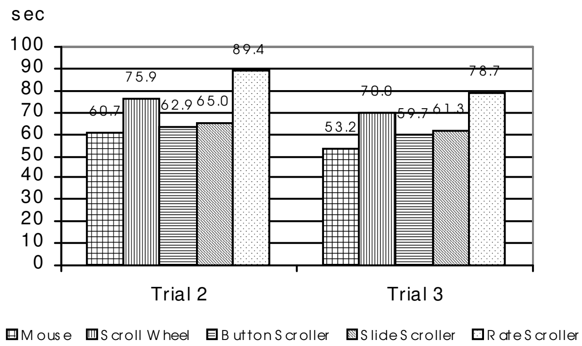

Figure 10. Times in seconds to scroll through 10 pages in trials 2 and 3 of the first version of the web page scrolling task using different input devices.

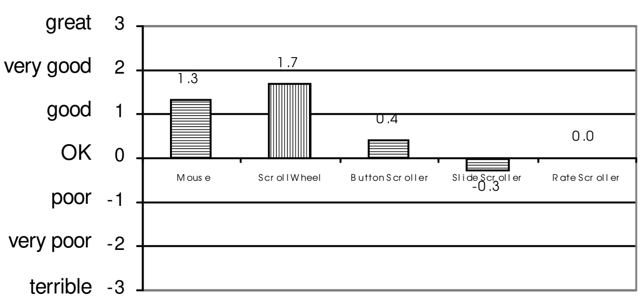

Figure 11. Ratings of the various input methods by the subjects in the first version of the scrolling experiment. We used the same scale as [15].

A repeated measure variance analysis shows that subjects' completion time was significantly affected by input method ($F_{4,44}$ = 13.3, p < .001). Trial 3 was significantly faster than Trial 2 ($F_{1,11}$=17.2, p<.001), showing a learning effect. However, this improvement did not alter the relative performance pattern of the input method ($F_{4,44}$=.5, p=.73).

Taking the Mouse condition as the reference and averaging over both trials, the scroll wheel, the Slide Scroller, and the Rate Scroller conditions were 28, 11, and 48 percent slower. The Tukey Test at .05 significant level indicates that the difference between mouse and scroll wheel conditions, between the mouse and button scroller, and between the mouse and slide scroller conditions were not significant, while the difference between mouse and rate scroller conditions was significant.

Figure 11 shows the subjects' ratings of the various scrollers using a rating scale from Zhai *et. al.* [15]. Contrary to the results of that previous study, the Tukey Test at .05 significant level indicates that the difference between ratings of mouse and scroll wheel was not significant. Subjects gave the mouse a significantly higher rating than the slide scroller, while the difference between ratings of mouse and button scroller and the difference between ratings of between mouse and rate scroller were not significant.. Subjects gave the scroll wheel a significantly higher rating than slide scroller and rate scroller, while the difference between ratings of scroll wheel and button scroller was not significant. The differences of ratings among the three Pebbles scrollers were not significant.

Revised Scrolling Task

We were not happy with the performance of the scrollers on the PDA, and the subjects provided useful feedback on ways to improve them. Therefore, we performed iterative design on the software, and tried the scrolling task again with 12 new subjects. Figure 12 shows that we were able to improve the performance of the new versions of the button scroller, but the rate scroller may be worse. The new absolute scroller was quite fast. The ratings of the new versions are shown in Figure 13 and parallel the performance.

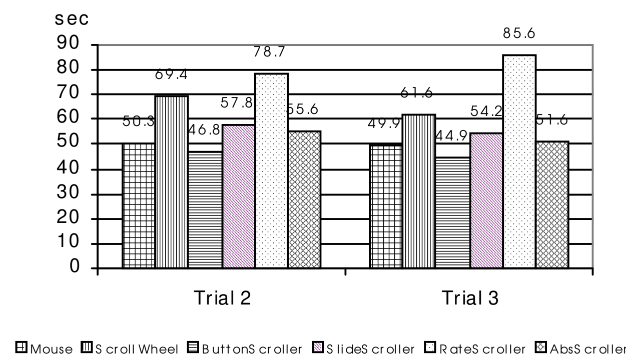

Figure 12. Times in seconds to scroll through 10 pages in trials 2 and 3 of the second version of the web page scrolling task.

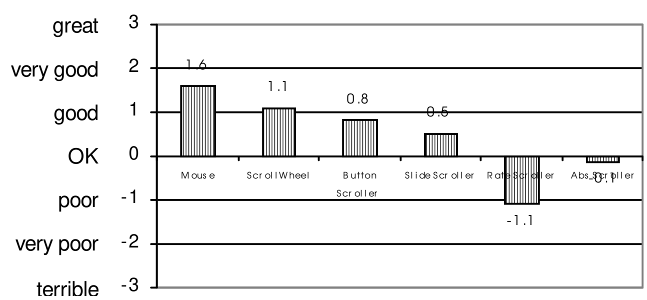

Figure 13. Ratings of the various input methods by the subjects in the second version of the scrolling experiment.

A repeated measure variance analysis showed that subjects completion time was significantly affected by input method ($F_{5,55}$ = 29.3, p < .001). Taking the Mouse condition as the reference and averaging over both trials, the button scroller was 8 percent faster but the Tukey Test at .05 significant level indicates that such difference is not significant. The scroll wheel, the absolute scroller, the slide scroller, and the rate scroller conditions were 31, 7, 12, and 64 percent slower than the standard mouse condition. The Tukey Test at .05 significant level indicates that the difference between mouse and scroll wheel conditions and the difference between mouse and rate scroller conditions were significant, while the difference between mouse and absolute scroller conditions and between Mouese and slide scroller conditions were not significant.

DISCUSSION

Button Size

The subjects were able to hit buttons quite accurately with their left hand, especially for small numbers of buttons. The predicted decrease in performance with decreased button size was observed. There seems to be a threshold of about 12 buttons before there is any affect due to the size.

We believe that we achieved expert performance (the learning curve flattened out) by the end of the experiment, so we tried using the times in models of expert human performance. One candidate is Fitts's law, but we do not know exactly where the subjects' fingers were when they started to move to tap. Assuming a movement of about 2 inches and a target size of 1 inch (in the 2x2 case), Fitts's law as formulated in [5, p. 55] predicts a movement time of about 150msec, compared to our measurement of 593msec. In our task, however, there is also perception and thinking time. For the smaller buttons (½ inch in the 16x16 case), Fitts's law predicts an increase in time of about 100msec, but we saw an increase of about 275msec. We observed that subjects looked back and forth from the monitor to the PDA, at an increasing rate depending on the number of buttons to choose from. Therefore, we believe the performance cannot be modeled simply as a Fitts's law task, but we were unable to find an appropriate alternative model.

Our results showing that users can tap up to 12 buttons accurately and quickly with the left hand is relevant since there are a number of applications where having several buttons on the PDA would be useful. Examples include scrolling with buttons (Figure 3a and Figure 4a), and panels created with the Shortcutter tool for controlling a compiler, playing music on the PC, reading mail, etc. [10].

Homing Times

Our one-handed homing time to move from the mouse to the keyboard (701 msec – see Figure 9) is longer than the time to move from the PDA to the keyboard (639 msec). This may be because the physical distance to the mouse from the home position on the keyboard is longer (14 inches compared to 7 inches) due to the number-pad and arrow keys sections of the PC keyboard. In the other direction, the increased time to acquire the PDA may be due to the unfamiliarity of homing to this kind of device.

In the classic study of text selection devices [5, p. 237], the homing time to move from the space bar to the mouse was measured as 0.36 seconds. This was measured from videotapes of subjects moving. An average homing time of 0.4 seconds was incorporated into the Keystroke Level Model [4]. However, we measured one-handed homing times of around 0.7 seconds, which is substantially longer. Our time was measured from the time of the mouse click to the time that the first keystroke was recognized. Our typing test shows that the average time per keystroke was 0.3 seconds, so this might be subtracted from our measured time to get the predicted 0.4 seconds.

An important observation is that, as predicted, subjects moved both hands simultaneously, and this did not penalize the movement time much. The sum of the one-handed times to move from mouse and PDA to the keyboard is 1340msec (701+639). This is much larger than the time to move from both mouse and PDA to the keyboard in the two handed case which is 791 msec (1340msec is 69% larger). A similar relationship holds for the movement from the keyboard to the PDA and mouse (728+7443=1473 > 838; 76% larger).

Overall, it takes only about 15% longer to acquire both the mouse and the PDA than just to acquire the mouse, and it takes only about 13% longer to get back to the keyboard from both devices than from just the mouse.

We were not able to find any prior studies of the time to acquire two devices at the same time. Most studies of two-handed use of input devices (including our button-size and scrolling tasks) allow the subjects to stay homed on the devices. We found that moving both hands slowed down each hand a little, but there was substantial parallel movement. Realistic tasks are likely to include a mix of keyboard and other input device use, so homing issues may be important.

Scrolling

Our measured times for scrolling the web pages with the mouse (about 60 seconds) is a little faster than the time reported in [15], and in the revised web task, the time for scrolling with the button scroller is 45.9 sec (average of trial 2 and trial 3) which is faster than the time for scrolling with the in-keyboard isometric joystick (around 50 sec). This shows that using the PDA can match or beat the speed of other non-dominant hand devices.

An interesting comparison is between their joystick, our Rate Scroller (Figure 3c) and the scroll wheel used in its most popular manner as a rate-controlled scroller. All provide the same rate-controlled style of scrolling, but they have significantly different performances and ratings by users. Our attempt to improve the rate scroller obviously did not help, showing that further work is needed to make this scrolling method effective. We observed that the fast speed was much too fast, but the medium speed was too slow. The popularity of the scroll wheel and the success of the pointing stick give us reason to keep trying. Furthermore, IBM did significant experimentation and adjustments before the pointing stick had acceptable performance [13]. Therefore, an important conclusion from the scrolling experiment is that the specific design and pragmatics of the input methods has a very important influence on the performance.

Another interesting result is that our subjects quite liked the scroll wheel (average rating of 1.7 ≈ very good), whereas in the earlier study it was rated much worse (-1 ≈

poor) [15]. This may be due to the increased experience people have with a scroll wheel (many of our subjects have a scroll wheel on their own mouse), and because most of our subjects used it in its rate-controlled joystick mode, whereas most of the earlier study's subjects used the rolling mode.

An interesting observation about this Web scrolling task in general is that it primarily tests scrolling while searching for information, so the scrolling must go slow enough so the subjects can see the content go by. This is why the methods that provided the best control over the speed are preferred. The low rating of the rate scroller on the PDA is because the fastest speed was much too fast to see the text go by, and the medium and slow speeds were rated as too slow. However, other scrolling tasks, such as those tested by [3], require the user to go to a known place in the document, and then a method that can move long distances very quickly may be desirable.

CONCLUSIONS

Many studies have shown the effectiveness of two-handed input to computers in certain tasks. One hindrance to two-handed applications has been that there may be only a few tasks in which using both hands is beneficial, and the benefits are relatively minor. Another problem is that although it is interesting to study custom devices for use by the non-dominant hand, in order for there to be wide-scale use, it is better to provide mechanisms that users can easily get and configure. Since increasing numbers of people have PDAs that are easy to connect to PCs, it makes sense to see if PDAs can be used effectively in the non-dominant hand. The research presented here shows that PDAs can be used as buttons and scrollers, and that the time to home to two devices is only slightly longer than for one. Our study of one application shows that at least for the scrolling task, a PDA can match or beat other 1-handed and 2-handed techniques. Because there is no incremental cost for the PDA since users already own it, and since the PDA is connected to the PC anyway, even small efficiencies many be sufficient to motivate its use as a device for the non-dominant hand. Our studies and many others have emphasized the importance of the pragmatics and the exact behavior of controls. Because the PDA can be programmed with a variety of controls with various properties, further research is required to determine the most effective ways that a PDA can be used to control the PC in both the dominant and non-dominant hand.

ACKNOWLEDGMENTS

For help with this paper, we would like to thank Rob Miller, Bernita Myers and Shumin Zhai.

The research reported here is supported by grants from DARPA, Microsoft, IBM and 3Com. This research was performed in part in connection with Contract number DAAD17-99-C-0061 with the U.S. Army Research Laboratory. The views and conclusions contained in this document are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory or the U.S. Government unless so designated by other authorized documents. Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

REFERENCES

1. Balakrishnan, R. and Kurtenbach, G. "Exploring Bimanual Camera Control and Object Manipulation in 3D Graphics Interfaces," in *Proceedings SIGCHI'99: Human Factors in Computing Systems.* 1999. Pittsburgh, PA: pp. 56-63.

2. Buxton, W., "Lexical and Pragmatic Considerations of Input Structures." *Computer Graphics*, 1983. **17**(1) pp. 31-37.

3. Buxton, W. and Myers, B. "A Study in Two-Handed Input," in *Proceedings SIGCHI'86: Human Factors in Computing Systems.* 1986. Boston, MA: pp. 321-326.

4. Card, S.K., Moran, T.P., and Newell, A., "The Keystroke-Level Model for User Performance Time with Interactive Systems." *Communications of the ACM*, 1980. **23**(7) pp. 396-410. July.

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

6. Guiard, Y., "Asymmetric Divison of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model." *Journal of Motor Behavior*, 1987. **19**(4) pp. 486-517.

7. Hinckley, K., *et al.*, "Two-Handed Virtual Manipulation." *ACM Transactions on Computer Human Interaction*, 1998. **5**(3) pp. 260-302. Sept.

8. Jacob, R.J.K., *et al.*, "Integrality and Separability of Input Devices." *ACM Transactions on Computer-Human Interaction*, 1994. **1**(1) pp. 3-26.

9. Kurtenbach, G., *et al.* "The Design of a GUI Paradigm based on Tablets, Two-hands, and Transparency," in *Proceedings, CHI'97: Human Factors in Computing Systems.* 1997. Atlanta, GA: ACM. pp. 35-42.

10. Myers, B.A., *et al.*, "Individual Use of Hand-Held and Desktop Computers Simultaneously," 1999. Submitted for Publication.

11. Myers, B.A., Stiel, H., and Gargiulo, R. "Collaboration Using Multiple PDAs Connected to a PC," in *Proceedings CSCW'98: ACM Conference on Computer-Supported Cooperative Work.* 1998. Seattle, WA: pp. 285-294.

12. Rekimoto, J. "A Multiple Device Approach for Supporting Whiteboard-based Interactions," in *Proceedings SIGCHI'98: Human Factors in Computing Systems.* 1998. Los Angeles, CA: pp. 344-351.

13. Rutledge, J. and Selker, T. "In-Keyboard Analog Pointing Device: A Case for the Pointing Stick," in *Technical Video Program of the CHI'90 Conference.* 1990. SIGGRAPH Video Review, Issue 55, No. 1.

14. TestedOK Software, "Speed Typing Test 1.0. Available from http://hometown.aol.com/tokfiles/typetest.html. TestedOK@aol.com," 1999.

15. Zhai, S., Smith, B.A., and Selker, T. "Improving Browsing Performance: A Study of Four Input Devices for Scrolling and Pointing," in *Proceedings of Interact97: The Sixth IFIP Conference on Human-Computer Interaction.* 1997. Sydney, Australia: pp. 286-292.