A STUDY IN TWO-HANDED INPUT

William Buxton and Brad A. Myers

Computer Systems Research Institute
University of Toronto
Toronto, Ontario
Canada
M5S 1A4

(416)-978-6320
willy@toronto.CSNET

ABSTRACT

Two experiments were run to investigate two-handed input. The experimental tasks were representative of those found in CAD and office information systems.

Experiment one involved the performance of a compound *selection/positioning* task . The two sub-tasks were performed by different hands using separate transducers. Without prompting, novice subjects adopted strategies that involved performing the two sub-tasks simultaneously. We interpret this as a demonstration that, in the appropriate context, users are capable of simultaneously providing continuous data from two hands without significant overhead. The results also show that the speed of performing the task was strongly correlated to the degree of parallelism employed.

Experiment two involved the performance of a compound *navigation/selection* task. It compared a one-handed versus two-handed method for finding and selecting words in a document. The two-handed method significantly outperformed the commonly used one-handed method by a number of measures. Unlike experiment one, only two subjects adopted strategies that used both hands simultaneously. The benefits of the two-handed technique, therefore, are interpreted as being due to efficiency of hand motion. However, the two subjects who did use parallel strategies had the two fastest times of all subjects.

CR Categories and Subject Descriptors: B.4.2 **[Input/Output and Data Communications]**: Input/Output Devices; D.2.2 **[Software Engineering]**: Tools and Techniques - User Interfaces; I.3.1 **[Computer Graphics]**: Hardware Architectures - Input Devices; I.3.6 **[Computer Graphics]**: Methodologies and Techniques - Interaction Techniques, Ergonomics.

General Terms: Experimentation, Human Factors

Additional Keywords and Phrases: Two-Handed Input, Parallel Input, Compound Tasks.

INTRODUCTION

A researcher turns a page of a book while taking notes. A driver changes gears while steering a car. A recording engineer fades out the drums while bringing in the strings.

What each of these tasks has in common is that the human operator is assigning a separate continuous task to each hand. What is clear is that we all perform this type of task every day. What is less clear is why hardly any user interfaces allow us to utilize this demonstrated ability in communicating with a computer.

From our experience in building systems for music and graphics, we were convinced that tapping this human ability could result in improvements in human performance for both experts and novices. Especially with the trend towards direct manipulation systems (Shneiderman, 1983), we were further convinced that such tasks were applicable beyond specialized applications such as process control and music.

In order to test our hypotheses, we designed and ran two experiements. The first has its roots in computer aided design and involves what we call a *positioning/scaling* task. The task for the second experiment is drawn from word processing, and involves selecting specified words from within a document. In the first experiment, all subjects (all of whom were novices) used two hands. In the second experiment, experts and novices were used to compare one and two handed techniques. The one-handed method was based on the scrolling mechanism used by the MacWrite word processor (Apple, 1984). The two-handed method was of our own design.

EXPERIMENT 1: POSITIONING/SCALING

Introduction

In our first experiment we had subjects perform a compound task where they positioned a graphical object with one hand and scaled its size with the other. The task was designed so that it could be perfectly solved serially by first positioning the object, and then scaling it. In addition, in our instructions and training, we did everything to bias users to perform it in a sequential manner.

Our hypothesis was that when the positioning and scaling sub-tasks were split over two hands and two devices, users would gravitate towards performing them in parallel. Despite the tendency towards sequential task performance assumed by most computer systems, our belief was that, for the positioning/scaling task, parallel performance of the sub-tasks was more "natural".

We also believed that the motor skills required to perform the task were either already existent or easily acquired.

The Task
Subjects were presented two squares on a CRT (shown in Figure 1). One square, known as the target, is positioned randomly on the screen and scaled to a random size. The other square is known at the tracker. The task is for the subject make the tracker box match the position and size of the target box. The position of the tracker square is controlled by the subject's right hand using a graphics tablet with a puck. The size of the tracker is controlled by the subject's left hand using a treadmill-like slider.

The squares were easily distinguished. The tracker was drawn with solid lines. The target was represented by bold-face corners. The centre of each square was marked by an identical fixed-size cross.

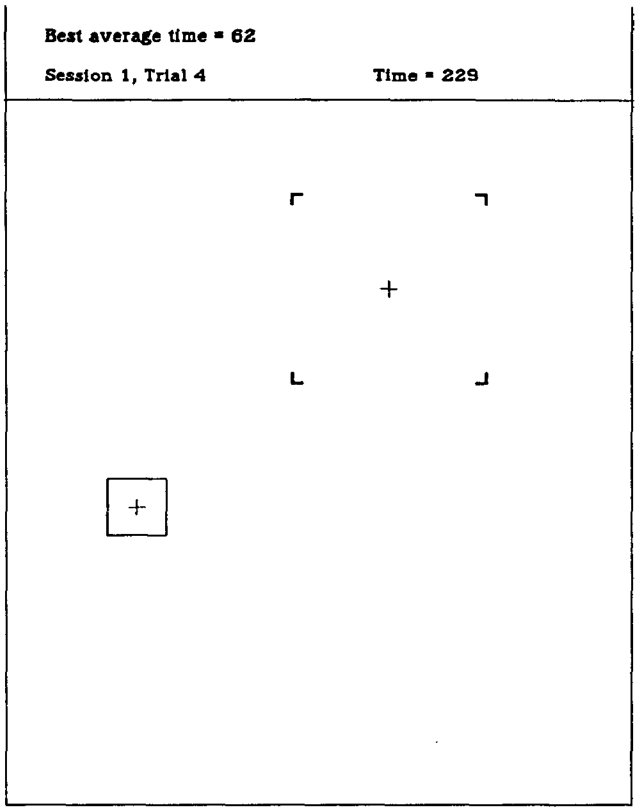

Figure 1: Experiment 1 Trial.

The target square is defined by bold corners. The tracker is the square defined by the solid lines. The goal is to position the tracker over the target and scale its size to match.

The tracker square changed size symetrically about its centre. Therefore, the two sub-tasks could be performed sequentially by aligning the centre cross of the tracker square with that of the target, and then adjusting the tracker's size.

Trials began by the subject depressing a button on the tablet's puck. A trial automatically finished when the scaling and positioning were within a system-defined degree of tolerance. The end of each trial was signalled to the user by an audio beep from the computer. The final position of the tracker for trial n became its initial position for trial n + 1. At the start of each trial, the target jumps to a new random position and assumes a new random size. Subjects could either hold the puck button down during a trial or click and release.

After training, subjects ran five sets of ten trials each. Sets were timed. At the end of each set, subjects were told their average time over that set as well as their own best average time. Subjects were instructed to try to beat their best average time. The total time taken by a subject to complete the experiment, including training and filling out a questionnaire, was about seventeen minutes.

The Environment
The experimental environment is shown in Figure 2. A PERQ I workstation from PERQ Systems Corp was used. It features a high-resolution (1024 x 768) non-interlaced bit-mapped display. The aspect ratio of the CRT was rectangular, and it was vertically oriented in portrait style.

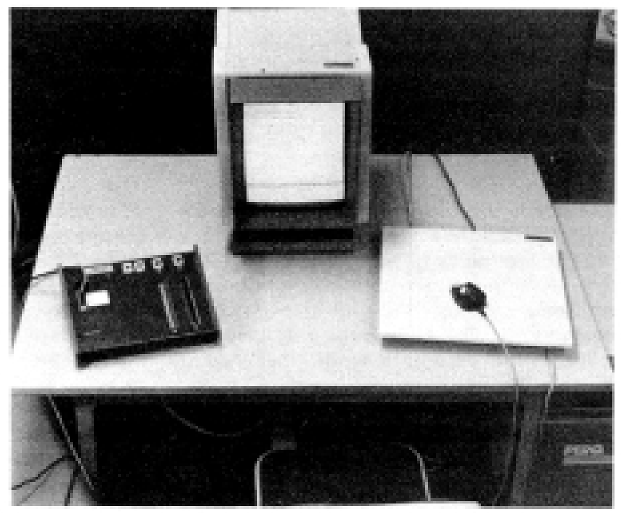

Figure 2:The Experimental Environment

The graphics tablet used was a Bit Pad-1 with a 4-button puck manufactured by Summagraphics Ltd. The tablet controlled the tracker in absolute mode so that there was a direct mapping of the position of the puck on the tablet to the position of the tracker on the screen.

The slider box was made at our Institute using a treadmill-like device developed by Allison Research of Nashville, Tennesse. The slider is, in effect, a 1-D mouse, providing relative information proportional to the amount of motion up or down. The slider is about 13 cm by 2 cm. A cut-away schematic of the slider is shown in Figure 3.

The workstation was in an area isolated by office partitions. All subjects used the same configuration with the workstation keyboard removed, the sliders on the left and the tablet on the right.

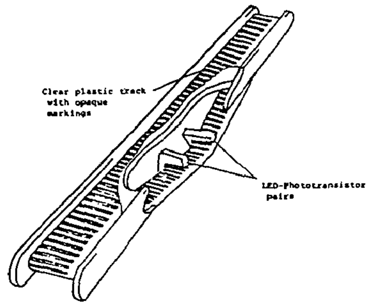

Figure 3: Cutaway view of an Allison Research Slider

Subjects
Fourteen subjects were used in the experiment. All were graduate students or staff associated with the Computer Systems Research Institute. Subjects were respondents to a call for volunteers posted in our building. Subjects were not paid, and none obtained course credit for their participation. Virtually all subjects were computer literate, some holding advanced degrees. However, all of the subjects in this experiment were novices in the use of computer pointing devices.

Training
Subjects were trained for the experiment in two stages, corresponding to the positioning and scaling sub-tasks, respectively. It was our intent during the training not to do anything (beside provide a device for each hand) to bias toward using parallel strategies in the experimental task. For consistency, all instructions were provided in written form on-line.

The first training session involved a task identical to that used in the experiment, except that the tracker and target were the same size. Hence, there was no scaling involved. After reading the instructions, subjects performed the task in sets of 10 trials until they reached a specified standard of proficiency.

The second part of the training was to develop familiarity with the slider and the scaling task. In this case, both squares were centred on the screen. In sessions of 45 seconds, the target square continuously grew and shrank. The subject was instructed to continuously match the target's size with the tracker square using the slider. If a specified degree of proficiency was not reached after the first session, additional practice sessions were presented.

Following completion of the two stages of training (which typically took on the order of 5 minutes), instructions for the experimental task were presented. Of utmost importance is to note that at no time was a subject informed that both devices could be used at the same time. Furthermore, the sequencing of the two training sessions follows the sequence in which the task can be performed perfectly without any parallel activity.

Results
The most important result was that all but one subject used both hands simultaneously in performing the task. Of the 14 subjects, six used both hands simultaneously *from the very first trial*. (The one subject who did not use both hands simultaneously reported that he thought that it was not allowed.) Averaged over all trials, subjects were engaged in parallel activity 40.9% of the time. If we look at only the best session for each subject, the figure becomes 45.1%.

In order to see how they correlated, we plotted speed of task performance against percentage of time engaged in parallel activity. This data is shown in Figure 4.

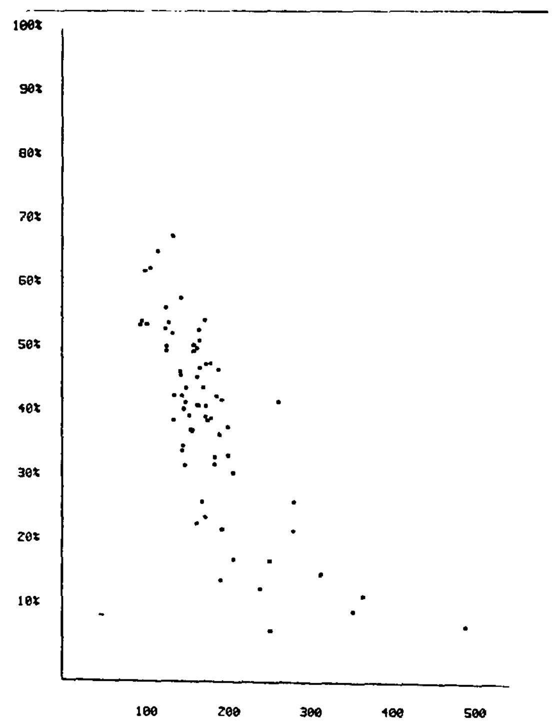

Figure 4: Time vs Parallel Activity

The horizontal axis represents time to complete the task (in 60ths of a second). The vertical axis represents percentage of time engaged in parallel activity. Average data for each session of each subject is plotted (5 x 14 = 70 points).

Interpretation
Subjects clearly have no difficulty in performing the task. The high incidence of parallel activity suggests that neither of the two sub-tasks presented a significant load on the cognitive or motor systems. The experiment shows that the efficiency of subjects' performance correlates positively to the degree of parallelism used. Perhaps most important, we believe that the experiment demonstrates that such behaviour is natural, at least for the task presented. This we support by the subjects' unprompted adoption of parallel strategies.

EXPERIMENT 2: NAVIGATION/SELECTION

Introduction
Having established subjects' ability to utlize two hand effectively, we were then interested in determining if there were common transactions where a two-handed approach would result in significant improvements in performance when compared to common one-handed techniques. We chose a task from word processing for the experiment. The task was to select specified

words in a stylized document. The experiment was designed so that the subject had to *navigate* (scroll or jump) to the appropriate part of the document before *selection* could take place.

To establish a known frame of reference, we modelled the one-handed technique on the scroll arrows and scroll bar of the MacWrite word processor (Apple, 1984). This we compared to a two-handed technique of our own design. MacWrite was chosen since it is representative of the current state-of-the-art. It also gave us access to a population of expert subjects.

Our hypothesis was that a well-designed strategy that partitioned the navigation/selection task between two hands would be easier to learn and use than the popular one-handed technique tested. Based on our previous experience, our belief was that complete novices using the two-handed technique would come close to matching the performance of experts using one hand.

The Task

The screen was partitioned into two halves. (See Figure 5). In the top half of the screen was a window 80 characters wide and 24 lines long. Part of a document was displayed within the window. In the bottom half of the screen, a one-line instruction was presented to the subject. Instructions were always to select a particular word on a particular line. Selection was always done using a puck on a graphics tablet. However, the specified word was never visible in the window at the time the instruction was given, so the user would have to navigate to the appropriate part of the document before selection could take place.

Subjects were divided equally into two groups. The group using one hand used the tablet and their right hand for selection and navigation (using the MacWrite-like scroll bar and arrows). The group using two hands used the tablet and their right hand for selection and two touch-sensitive strips and their left hand for navigation.

The document consisted of 60 numbered lines double-spaced. Lines were numbered at both the left and right margins. Each line contained three words: Left, Middle, and Right. The words were placed in three columns at the left, middle, and right of the lines.

We chose this stylized document to better approximate the case where one is navigating within a familiar document. Subjects performed three sessions of 21 trials (resulting in 20 transitions per session). To better focus on operational issues, the same questions were presented in each of the three sessions.

With both the one-handed and two-handed versions, there were two strategies that one could use to navigate. One was to smooth scroll, the other was to jump. With the one-handed version the scroll arrows were used to smooth scroll and the scroll bar to jump. With the two handed version, one touch-sensitive strip was used to scroll the document up or down (by sliding the finger). The second strip caused a jump to the same relative position in the document as the position on the strip that was touched (top->beginning, bottom->end, ...).

Subjects were timed and presented their average time and best time at the end of each session. They were instructed to try to beat their best time. The time required for subjects to complete their participation, including training and filling out a questionnaire was about twenty minutes.

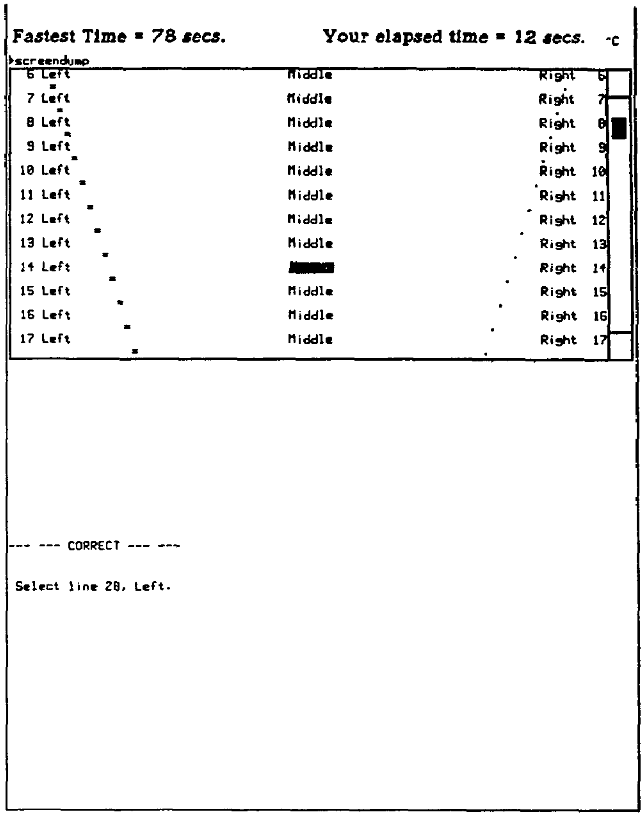

Figure 5: Sample Trial for Two-Handed Version.

The subject has just correctly selected line 14, Middle. The program has responded by instructing "Line 28, Left" to be selected. The current relative position in the document is indicated by the black bar in the scroll bar in the right margin of the window. In the two-handed version, the graphic scroll bar is for output only.

The Environment

The environment used was the same as that described for Experiment 1. The only difference was that the slider box was replaced by a touch-sensitive tablet. The touch-sensitive surface and its controller were manufactured by Elographics Corp. The power supply and housing were of our own manufacture. The touch-tablet's surface was partitioned into two vertical strips by using a cardboard template. Each exposed strip measured about 4.5 cm by 2 cm. A photograph of the touch-tablet is shown in Figure 6.

Subjects

Twenty-four subjects ran the experiment. Twelve were experts in the use of a mouse and twelve were novices. Half of each group ran the one-handed version of the experiment, the other half ran the two-handed version. Hence, there were four groups of six subjects in a two-by-two comparison.

Subject expertise was determined using a questionnaire. The data generated in the experiment strongly verifies our grouping of subjects. Subjects were staff or students (graduate and undergraduate) associated with the Department of Computer Science. Subjects were respondents to either posted or verbal calls for volunteers. No subjects were paid, and none obtained course credit for their participation.

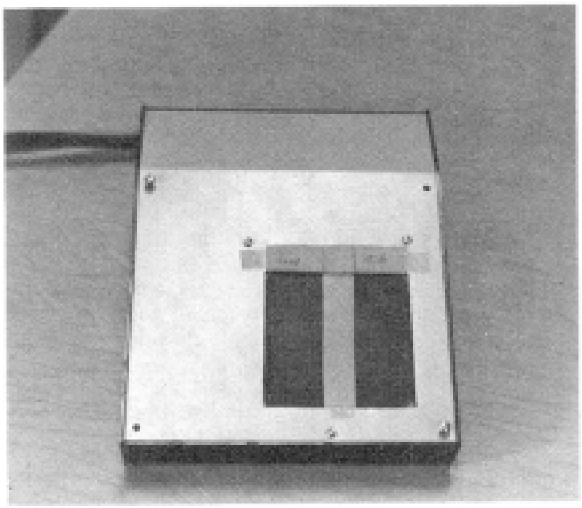

Figure 6: The Touch-Sensitive Tablet

The tablet surface is partitioned by a template into two virtual devices. The left one is a position-sensitive strip used to jump to specific locations in the document. The right one is a 1-D relative device used to smooth scroll the document in the window. See Buxton, Hill and Rowley (1985) for additional information on the use of touch-tablets.

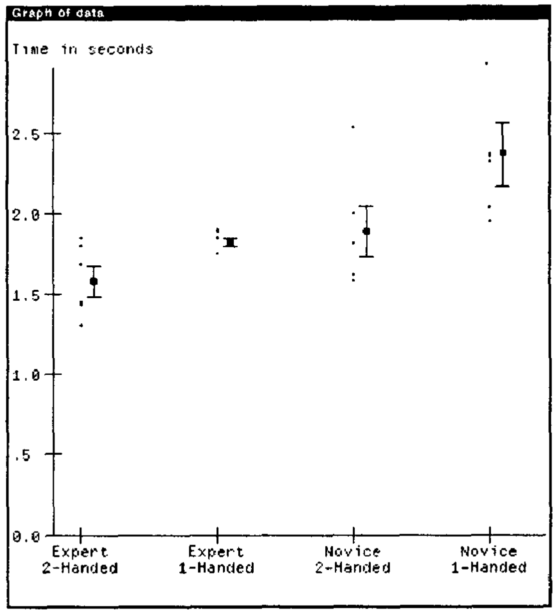

Figure 7: Subjects' Performance by Group

The average time per trial for the best set of each subject organized by group. Group median and standard deviation are also shown.

Training
To maintain consistency, all training was done on-line. Subjects were presented the document and given instructions on how to use the particular navigational tools assigned to them. Different instructions were obviously provided to the two-handed and one-handed groups.

Results
The first result showed that the two-handed approach resulted in better performance by experts and novices alike.

1. Experts: the two-handed group outperformed the one handed group by 15%.
2. Novices: the two-handed group outperformed the one handed group by 25%.

Using the two-handed technique greatly reduced the gap between expert and novice users. If we look at the average times taken from the *first* set of trials, we see the following:

1. One-handed: experts outperformed novices by 85% ($p = 0.05$).
2. Two-handed: experts outperformed novices by only 32% ($p = 0.02$)
3. Experts using one hand outperformed the novices using two hands by only 12%, and this difference has no statistical significance.

If we look at average times for subjects' *best* of three sets, we also see that the two-handed technique resulted in superior performance. The top six times of all subjects were obtained by those using the two-handed technique. A comparison of the performance of the subjects by group is summarized in Figure 7.

Based on our results from Experiment 1, we expected to observe subjects using both hands in parallel (for example, moving the tablet puck from one side to the other while the target was being scrolled into view). However, with the majority of users this was not the case. Only two subjects employed parallel strategies. Significantly, these two subjects had the two best times in the experiment.

Regarding strategies, all subjects jumped significantly more than smooth scrolled during searches, although the effect was more pronounced with experts. The data also shows that in their best set, experts jumped far more when using the one-handed version than when using the two handed version (93% vs 74% of the time, respectively).

Interpretation
The results show that the partitioning of the navigation/selection task between the two hands results in improved performance for experts and novices. The first order benefit cannot, however, be attributed to the two hands being used at once. Rather, the improvement is interpreted as being due to the increased efficiency of hand motion in the two-handed technique. In the one-handed approach, significant time is consumed in moving the pointer between the document's text and the navigational tools. In the two-handed version, the hands are always in home position for each of the two task, so no such time is consumed.

If this interpretation is correct, we would expect to see the greatest improvement in performance in transitions where there is the greatest distance between the target and the navigational tools. This situation occurs in the one-handed technique where two selections occur in sequence on the left side of the display (since the scroll-bar and scroll-arrows are along the right margin). This expectation was confirmed by the data. With such transitions, the two-handed technique resulted in performance improving by 30% . However, with transitions that minimized the movement between target and navigational tools (two targets appearing in sequence on the right side of the display), the two-handed technique still resulted in an improvement of 15%.

Unlike the experimental condition, in many real-world tasks, time is lost to homing with the two-handed technique as well. An example would be where the hands frequently move back-and-forth from the keyboard. This may make the benefits of the technique of less practical significance overall. Note, however, that in such contexts, time is equally lost in homing using the one-handed technique.

Finally, we must address the question of why more simultaneous use of two hands was not observed. We can only conclude that the task in Experiment 2 was more difficult than that in Experiment 1. It is important to remember, however, that despite the fact that the entire experiment took subjects less than twenty minutes, two did adopt a parallel strategy, and these two subjects obtained the best times overall. Consequently, while more difficult, the skill can be easily learned and performance benefits can accrue when it is.

CONCLUSIONS

The data generated makes a strong case for improving performance by splitting the sub-tasks of compound continuous tasks between the two hands. Experiment 1 shows that even novice users have the requisite manual skills, and Experiment 2 shows that significant improvements can be made over one-handed techniques which are the current practice.

Experiment 2 shows that performance improvements can occur with two handed input even where the tasks are performed sequentially. Furthermore, by splitting the tasks between two hands, the foundation is laid for further improvement by the ability to support parallel task performance by more skilled users.

To date, very few computer systems easily lend themselves to experimentation with the types of interaction described in this paper. This may be largely due to the serial nature of existing programming languages and processors. Technological biases nonwithstanding, we feel that the results reported here warrent increased attention being paid to multi-handed and parallel input structures.

ACKNOWLEDGEMENTS

We are indebted to a number of people for help in running this experiment. Dorothy Philips and Stu Card gave a great deal of help in the design of the positioning/scaling experiment. Christine Warchol wrote the first implementation of that experiment, although that version was never run due to a computer death. Guy Fedorkow built the slider box used in Experiment 1. Jan Venus constructed the touch-sensitive unit used in Experiment 2. Ralph Hill and the CHI'86 referees provided many useful comments. Liz Russ helped type and proof-read the manuscript. Finally, we are indebted to the volunteers from the University of Toronto who participated as subjects.

This research has been funded by the Natural Sciences and Engineering Research Council of Canada. This support is gratefully acknowledged.

REFERENCES

Apple (1984). *MacWrite User's Manual*. Apple Computer Inc.

Buxton, W. (1982). Lexical and Pragmatic Considerations of Input Structures. *Computer Graphics* **17(1)**, 31 - 37.

Buxton, W., Hill, R. & Rowley, P. (1985). Issues in Touch-Sentitive Tablet Input, *Computer Graphics* **20(3)**, 215-224.

Shneiderman, B. (1983). Direct Manipulation: A Step Beyond Programming Languages. *IEEE Computer*, **16(8)**, 57 - 69.