USING PSYCHOMOTOR MODELS OF MOVEMENT IN THE ANALYSIS AND DESIGN OF COMPUTER POINTING DEVICES

by

ANANT KARTIK MITHAL

A DISSERTATION

Presented to the Department of Computer and Information Science
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

August 1995

"Using Psychomotor Models of Movement in the Analysis and Design of Computer Pointing Devices," a dissertation prepared by Anant Kartik Mithal in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Political Science. This dissertation has been approved and accepted by:

Sarah Ann Douglas

Dr. Sarah Ann Douglas, Chair of the Examining Committee

August 9, 1995

Date

Committee in charge	Dr. Sarah A. Douglas, Chair Dr. Arthur Farley Dr. Steven Fickas Dr. Steven W. Keele

Accepted by:

Vice Provost and Dean of the Graduate School

An Abstract of the Dissertation of

Anant Kartik Mithal for the degree of Doctor of Philosophy
in the Department of Computer and Information Science
to be taken August 1995

Title: USING PSYCHOMOTOR MODELS OF MOVEMENT IN THE ANALYSIS AND DESIGN OF COMPUTER POINTING DEVICES

Approved: ______________________
Dr. Sarah A. Douglas

This study applies the Stochastic Optimized Submovement Model, a psychomotor model of movement, to computer pointing devices. Modeling such movement is necessary in order to build better pointing devices and to explain differences in performance between existing pointing devices. Current Fitts' law research has shown that isotonic devices perform better than isometric devices, but cannot explain why. The reason for this is that Fitts' law only considers overall movement time and distance, and does not consider what happens during movement. Psychomotor models of movement on the other hand, predict the manner in which movement occurs and are based on an analysis of the microstructure of movement.

Six subjects performed 1800 trials of a pointing task with each of two devices, a mouse and an isometric joystick. The microstructure of movement of these trials was compared between devices, within devices, and within subjects. The comparison showed that the force sensitive isometric joystick picks up tremor causing random changes in cursor velocity that make the device difficult to use. Tremor was not observed in mouse trials. A frequency-domain analysis showed that the joystick trials had higher-frequency components than those seen in mouse movement. This finding is consistent with the presence of tremor. Filters can be built that exploit this difference and reduce the effect of tremor.

The microstructure level analysis of joystick trials revealed a pattern of differences between subjects with high performance and those with low performance. This suggests that some subjects can overcome tremor with practice. The analysis was also used to test the applicability of the Stochastic Optimized Submovement model to the two devices. While the model held for the mouse, it could not be extended to the joystick because the tremor obscured the underlying movement.

The study demonstrates the importance of conducting a microstructure level analysis in addition to the standard Fitts' law analysis techniques that are currently used to study pointing devices.

CURRICULUM VITA

NAME OF AUTHOR: Anant Kartik Mithal

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon
Indian Institute of Technology, New Delhi
St. Stephen's College, University of Delhi

DEGREES AWARDED:

Doctor of Philosophy in Computer and Information Science, 1995, University of Oregon
Master of Science in Computer and Information Science, 1989, University of Oregon
Master of Science in Education, 1987, University of Oregon
Bachelor of Technology in Electrical Engineering, 1983, Indian Institute of Technology

AREAS OF SPECIAL INTEREST:

Human-Computer Interaction
Artificial Intelligence / Natural Language Understanding
Language Understanding

PROFESSIONAL EXPERIENCE:

Research Assistant, College of Education, University of Oregon, Eugene, 1987-88

Research Assistant, Department of Computer and Information Science, University of Oregon, Eugene, 1988-1995

DEDICATION

This thesis is dedicated to my mother, Dr. Shakuntala Mithal, MBBS, Ph. D. and the rest of the nine Ph. D.s in my family, all women, who came before me and were the inspiration.

ACKNOWLEDGMENTS

Generous support in terms of time and equipment by a large number of people made this research possible. I thank:

My subjects, whom regrettably I cannot mention by name. Neff Walker, David E. Meyer and J. B. Smelcer for the code from their experiment. Richard Jagacinski for feedback about early ideas. Gurdeep Pall, formerly a graduate student at the University of Oregon, and his friends at the NT development group at Microsoft who helped with timing and other issues during two enjoyable and highly productive days at Microsoft. Marie Vitulli of the Math Department at the University of Oregon, and Kennon T. Smith (Prof. emeritus) and Adel Faridani of the Math Department at Oregon State University for their help with my numerical analysis. Ted Selker of IBM's Almaden Labs for a TrackPoint II keyboard.

Kesavan Mithal for giving me his own computer so that I could get my research started. Shakuntala and Akhilesh Mithal for financial assistance for the research machine, and years of unwavering support. Jan Saunders, Betty Lockwood and Nagwa Naguib for listening and for smoothing the complications of being a graduate student. Steve Keele for patient support and guidance as I ventured out of my area of expertise into his. Art Farley and Steve Fickas for always having their doors open and making time when I needed to talk to them. Sally (Sarah Ann) Douglas for, well, everything.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Page

CHAPTER I

INTRODUCTION

The advent of Graphical User Interfaces (GUIs) has made the use of a pointing device a vital aspect of interacting with a computer. Once we have used a pointing device, it is almost unacceptable to go back to using a computer without one. The extent to which pointing is used as a means of communicating with the user interface has increased the number of tasks and applications that make use of pointing devices. Many contemporary applications assume the existence of a pointing device and cannot be operated without one.

Pointing is both natural and pervasive in a GUI, and users spend much of their time on computers using a pointing device. This makes the quality of pointing devices an issue, and we would like to have good pointing devices. The devices should be good in the sense that they should be easy to use and learn, quick to work with, have low error rates, should not be fatiguing and should not endanger the health of the users through effects such as like repetitive motion injury.

Because the quality of pointing devices is an issue, researchers and manufacturers try to build new and better pointing devices with the ultimate goal that a superior pointing device will help them capture a larger share of the market. This has prompted research and development in pointing devices to improve existing pointing devices, as well to develop novel pointing devices.

A wide range of pointing devices currently exist , and this might make us believe that new pointing devices are not required. However, novel pointing devices are constantly being developed because as new uses are found for computers, new situations arise which make new kinds of pointing devices necessary. For example, a touch tablet is better for drawing than a mouse. Similarly, when a notebook computer is used on an airplane, a mouse is inconvenient, which lead to the use of trackballs, joysticks and now trackpads in portable computers. Supermarket information consoles, on the other hand, have touch screens rather than mice or trackballs. Thus the circumstances of the task, and the environment within which a computer is used often make one pointing device more appropriate than another.

Unfortunately, the design of pointing devices is an ad hoc process; it has lacked principle, and has not been guided by pointing device research. It has tended to go through cycles of building a device, seeing how users perform, tweaking various parameters and then re-testing the device with users. The cycle is repeated till satisfactory performance has been achieved. The design of isometric joysticks is an example of this approach (Bentley, Custer, & Meyer, 1975; Rutledge & Selker, 1990).

The problem with this technique is that it focuses only on the gross parameters of the pointing task, i.e., the target distance, target width, and the time it takes to complete the movement. It does not focus on the manner in which the movement is made. As a result when the designers make a change in the pointing device, they do not know how the change affects the manner in which their users point with the device, and therefore cannot predict whether or not the change will improve their device.

In addition, they cannot predict how well a novel device will perform without actually building the device and testing. If the device happens to perform poorly, it is often unclear why it performs poorly. This was the case with an isometric joystick that an earlier study showed to be 70% slower than the mouse (Douglas & Mithal, 1993). That study could not provide an explanation for the isometric joystick's relatively poor performance, and was also unable to suggest design changes that could improve the device.

One promising area that has been ignored by pointing device research is research on psychomotor models of movement. This research hypothesizes the manner in which people point, uses the hypotheses to predict the characteristics of movement, and then validates the model by gathering data about how pointer is moved towards the target. While current research on pointing devices looks at gross movement characteristics, this approach looks at the microstructure of movement.

This approach can be very valuable for pointing devices. Knowing the characteristics of movement with a given pointing device will help designers build better pointing devices. This dissertation takes this approach and studies the characteristics of movement of two different pointing devices to see if differences in their performance can be explained by differences in their movement characteristics.

The following sections discuss the current state of research on pointing devices, the goals of this research, and the research questions. The last section provides an overview of the study and summarize the following chapters.

The State of Research on Pointing and Pointing Devices

Research on pointing devices is closely related to research on manual pointing. The two areas are linked together by Fitts' law (Fitts, 1954; MacKenzie, 1992). Fitts' law describes a property of the gross parameters of a movement. It relates the time it takes to move to a target to the target distance and the target width. Chapter II will describe Fitts' law and related research in further detail. What makes Fitts' law important is that empirical research has shown it to hold across a very large range of movements, limbs, devices and tasks. Fitts' law holds not only for pointing actions with the finger, but also for other limbs such as the feet, wrists, and eyes. It also holds for movement with devices such as mice, joysticks and touch tablets, for manual tasks such as stylus tapping, hole pegging, and for computer tasks such as pointing and dragging.

Fitts' law can be used to organize the state of research on pointing as represented by Figure 1. Broadly speaking, research on pointing and pointing devices can divided into two areas, depending on whether it is about manual pointing or pointing on a computer. The categorization depends on whether the research is conducted by psychologists or by researchers in Human Computer Interaction (HCI). Psychologists study manual pointing and may use the computer as a means for presenting the experimental task. Researchers in

HCI study pointing with computer pointing devices. The large number of psychologists using computers blurs the distinction between the two areas.

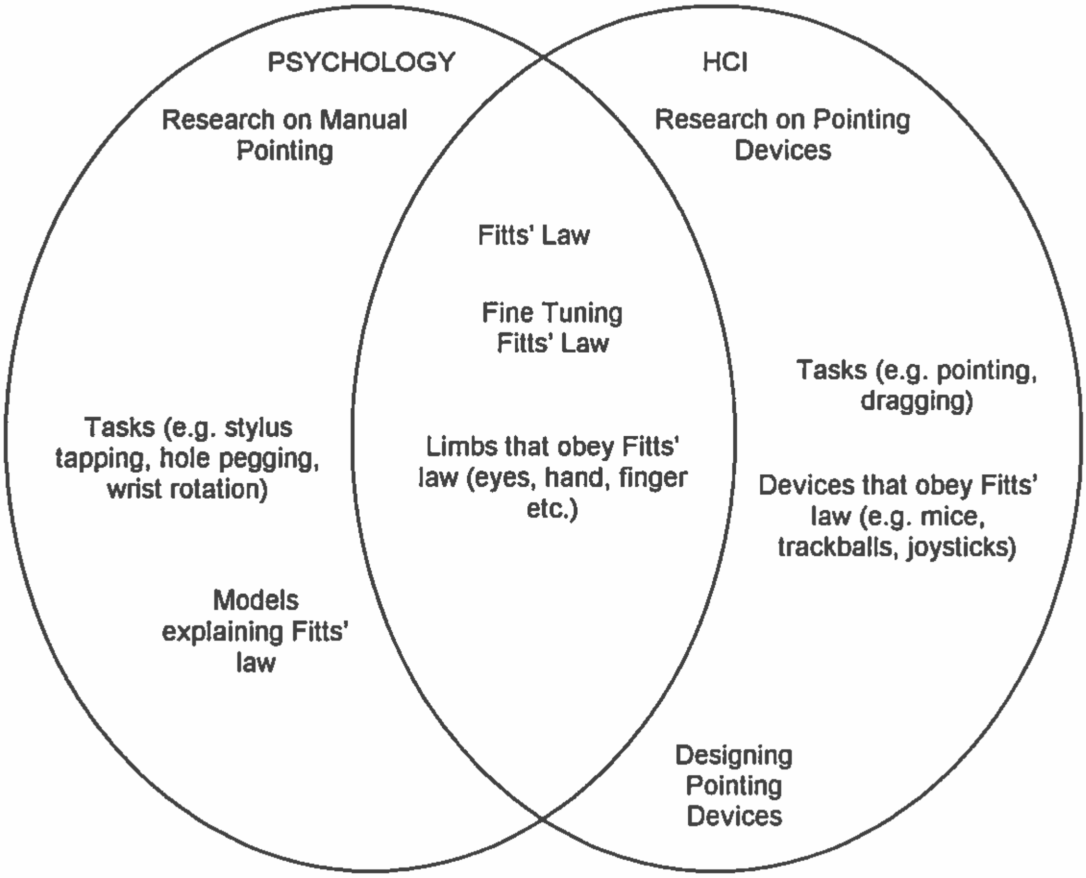

Figure 1. Framework of research on pointing and pointing devices.

Figure 1 shows reassuring parallels and surprising differences between the research done by the two sides. If we call the range of conditions under which Fitts' law holds a space, then both psychology and HCI have tried to map the space in which Fitts' law is applicable. These spaces are somewhat different, yet strikingly similar. One common dimension of the space is the task, where the attempt has been to determine what kinds of tasks conform to Fitts' law. A second common dimension is the limb used in the task. The pointing device used makes for a third dimension in the space for HCI research although the pointing device and the limb used to control it are related. Thus psychology has studied the mix of limbs and tasks that follow Fitts' law, while HCI has examined the limbs, tasks and devices that follow Fitts' law. Limbs, devices and tasks are discussed in detail in Chapter II.

Both psychology and HCI have looked at refining the Fitts' law equation to get a better fit of the Fitts' law equation to their data, with the research on both sides complementing each other.

While these parallels on the two sides of the research are reassuring, two gaps appear, which are disconcerting. Most of the work on modeling pointing actions has been done by psychologists, while most of the design of pointing devices has been done on the HCI side. At one level, this is not surprising, and is as it should be. Psychology is a science, and one of its goals is to build models to explain phenomena. HCI, on the other hand is somewhat of an engineering discipline, where the goal is the application of science in the creation of artifacts.

Research on models explaining Fitts' law has been prompted in part by the fact that the law holds across a wide range of conditions, which strongly suggests that there is some powerful unifying mechanism that controls movement. HCI, by focusing on the gross characteristics of movement has not been in a position to verify the applicability of psychomotor models to pointing devices. Such an analysis requires detailed time and displacement data as the cursor moves towards the target, data that is typically not gathered in a Fitts' law experiment.

For instance, one psychomotor model, the Stochastic Optimized Submovement model (the SOS model) has been shown to be applicable to pointing movement (Meyer, 1988; Meyer, Smith, & Wright, 1982; Meyer, Smith, Kornblum, Abrams, & Wright, 1990). The model, which is explained in more detail in Chapter II, says that a single pointing action is made up of a sequence of submovements, and predicts the relative size and accuracy of the submovements. It seems that the logical approach to developing novel pointing devices, or to improve existing pointing devices would be to use such models of human movement in the design process, but this has not been done.

There are thus two gaps in the existing knowledge about pointing devices. First, there is little knowledge about the characteristics of movement with different pointing devices. As a result we do not know if differences in the characteristics of movement can account for differences in performance. HCI research on pointing devices has focused on establishing differences in pointing device performance, not on why the differences exist. Second, because we do not have knowledge about the characteristics of movement, we do not know what psychomotor models of movement are applicable to pointing devices. This dissertation aims to fill these gaps.

Goals of this Research

The broad goals of this dissertation are to answer the following questions. Are there differences in the movement characteristics of pointing devices, and can these be used to explain differences in their performance? Can knowledge of the movement characteristics with a pointing device be used to improve its design? Can psychomotor models developed to explain human pointing movement be used to explain movement with computer pointing devices?

In order to achieve these goals the SOS model was used as the framework for analysis (Meyer, 1988; Meyer et al., 1982; Meyer et al., 1990; Walker, Meyer, & Smelcer, 1993). The SOS model relies on an analysis of the microstructure of movement (Jagacinski, Repperger, Moran, Ward, & Glass, 1980a), which is an analysis of the velocity and acceleration over distance as a subject points to a target.

The present study gathered such data for two pointing devices, the mouse and the isometric joystick. These devices were selected for two reasons. First, earlier studies of isometric joysticks have shown that they are much slower than isotonic devices (Card, English, & Burr, 1978; Douglas & Mithal, 1993; Douglas & Mithal, 1994; Epps, 1986), and no one has been able to explain why this difference exists. Second, there are a number of characteristics that make the devices an interesting pair for this study.

The mouse is the ubiquitous pointing device. Movement with a mouse is similar to movement with the hand, with two major differences. The first is that the focus of attention is the computer screen rather than the mouse. When pointing with a finger the focus of attention is the finger. The second is that the mouse moves in two dimensions while natural movement is in three dimensions. The mouse has been studied to examine how well it is described by the SOS model, and the study showed a good fit (Walker et al., 1993) making it suitable as a base line device.

The isometric rate-controlled joystick differs from the mouse in two ways. First it is used to control the velocity of the cursor, while the mouse is used to control the displacement of the cursor. Second, it is isometric, and it is unclear whether models of movement defined for isotonic devices will map onto isometric devices.

Finally, because they are compact, can be mounted on the keyboard and are inexpensive to manufacture, isometric joysticks have become popular in today's notebook computers. As an informal measure of the prevalence of isometric joysticks in notebook computers, we counted the number of pictures of notebook computers in the May 1995 issue of Windows magazine, and categorized them according to whether the pointing device was an isometric joystick or a trackball, and found 23 computers with isometric joysticks, 23 with trackballs, and 5 had some other kind of pointing device showing the extent to which isometric joysticks have come into use. IBM and Toshiba account for a large percentage of the shipments of notebook computers. In addition, some manufacturers such as IBM market add-on keyboards with built-in joysticks with the idea that they would appeal to touch typists who would not have to take their hand off the keyboard in order to point.

Stated simply, the research goals of this study were to use psychomotor models of movement to explain differences in movement with pointing devices.

Research Questions

1. The research questions that arose out of these goals are:
2. Do the devices (mouse and isometric joystick) follow Fitts' law? What are the constants in their Fitts' law equations? How do the indices of performance for these devices compare with one another?
3. In terms of the duration and accuracy of the first and subsequent submovements, what is the microstructure of movement for the isometric joystick? What is it for the mouse?
4. Are there differences in the microstructure of movement of the mouse and isometric joystick, and do these differences explain the difference in performance?
5. Are there changes over time in the microstructure of movement for individual subject's performance?
6. Are there differences in the microstructure of movement between individual subjects?
7. Do models of movement that describe isotonic devices describe isometric devices? In particular, does the SOS model describe the movement of both the mouse and the isometric joystick?

8. If the SOS model does not describe the movement of the isometric joystick, can we modify the model so that it does?
9. Can we use our model for the isometric joystick and the knowledge about its movement microstructure to improve its design?

These questions are discussed in greater detail below.

To begin with, we would like to know whether the devices follow Fitts' law. The SOS model analysis makes this assumption, and therefore it has to be tested. This leads to the first question. Once it has been established that Fitts' law holds, it is important to know the qualitative nature of the submovements is for each of the devices. This leads to question two.

Once the microstructure of movement has been established for each device, differences in the microstructure might provide explanations for differences in performance between the devices, leading to question three. Question four looks at the effect of practice on the microstructure of movement while question five examines performance differences between subjects.

The SOS model has been applied to pointing with the mouse (Walker et al., 1993). This study can provide additional support for that finding. In addition, if it can extend the SOS model to the joystick, then it increases our confidence that the model is universally applicable, which leads to question six. If the SOS model does not hold for the isometric joystick, the microstructure level analysis can be used to develop an alternative model, leading to question seven.

The motivation for modeling pointing devices is to enable designers to make better design decisions, and the eighth question examines whether the knowledge gained through this study can be used to improve the design of the isometric joystick.

While this approach was applied to an isometric joystick, a similar approach can be applied to any pointing device. This approach can also be applied to other machine mediated movement, such as prosthetic devices, remote control of space-craft, and movement in virtual reality.

Summary

To summarize, this dissertation arose out of the need for a better understanding of pointing devices, out of the need to know how people point with them. The next chapter describes existing research on pointing and pointing devices. It also describes how psychologists have developed models of how people point. Applying the same ideas to pointing devices is the goal of this study.

In order do this, an experiment was conducted to gather pointing data from the mouse and finger-operated rate-controlled isometric joystick. The methodology of this experiment is described in Chapter III. Chapter IV describes the results of a pilot study and the main study. The microstructure level analysis conducted on the data from the isometric joystick revealed that it picks up tremor which makes it hard to control. Chapter V discusses this and other results. Chapter VI describes the conclusions and opportunities for further research.

CHAPTER II

LITERATURE REVIEW

Many physical skills require a person to rapidly change the position of a limb from one location to another. Tasks such as reaching for a pen or an elevator button, hitting the reset button on a computer, using a mouse to select a screen icon and fitting a bolt into its hole, are examples of actions that can be collected under the rubric of pointing actions. Pointing can be defined as the act of moving a limb from a starting position to some target location with some tolerance in the final resting position of the limb. Of interest are the width of the target location, the distance from start location to the center of the target, and the time it takes from the start to the end of the movement. This is illustrated in Figure 2.

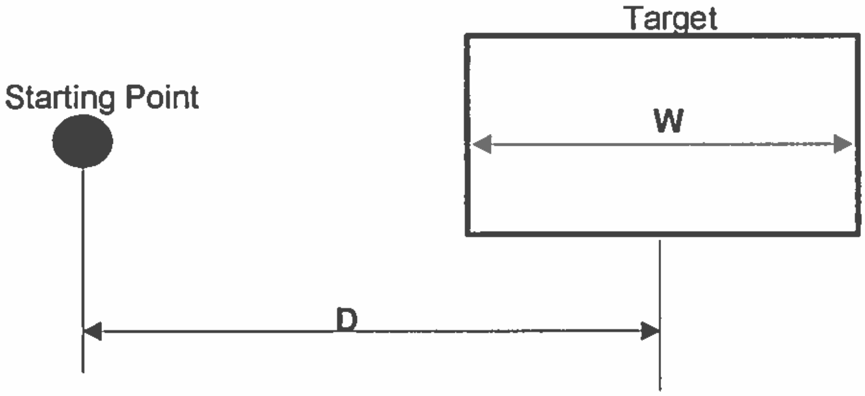

Figure 2. The parameters of interest in a pointing action.

Pointing actions are fundamental to the understanding of mechanical skills in humans (Keele, 1986; Meyer et al., 1990). The analysis of pointing actions has therefore been a major concern in the field of motor performance, with notable research being published as early as 1899 (Meyer et al., 1990). Pointing is also an important part of interacting with GUIs, as described in Chapter I, and research on computer pointing is an important aspect of Human Computer Interaction (HCI).

Research into pointing movements came to the fore in psychology at a time when worker efficiency was of major concern. The early experiments performed with pointing movements reflected the kinds of tasks that might be performed by workers on a production line. This led to the establishment of Fitts' law which predicts the time it takes to move to a target using the relationship between the target width and target distance (Fitts, 1954). Fitts' law provides industrial engineers with predictive tools to minimize task time on a production line. Subsequent research was able to establish that Fitts' law holds in a wide range of conditions, and that its predictions are very accurate.

In the late 70's it was established that Fitts' law holds for pointing with a mouse, and subsequent research shows that the law extends to computer pointing devices such as mice, joysticks and touch tablets. The accuracy of the predictions of Fitts' law are so unique and so valuable for research in HCI, that it has been upheld as a vital principle in the field (Card, Moran, & Newell, 1983). HCI focuses on human behavior which, generally speaking, varies widely. In this light, it is both amazing and reassuring that once we make some simple empirical measurements, we can predict human performance with a pointing device very accurately.

Chapter I described a framework to categorize research on pointing. It differentiated between research on manual pointing and research on pointing with pointing devices. Fitts' law predicts movement time in both these areas. With that in mind, this chapter starts by describing the original Fitts experiment, followed by a description of research on manual pointing by psychologists. This section also looks at psychomotor models of movement. The third section looks at research on the evaluation and design of computer pointing devices. The concluding section discusses the characteristics of pointing devices that could lead to differences in the way they are used. Due to the wide variety of pointing devices, it is possible that existing psychomotor models of movement may not be applicable to all pointing devices.

Fitts' Law and a Description of Fitts' Paradigm

Any discussion of psychomotor research on pointing has to include Fitts' original work which established a mathematical relationship between the time it takes to move to a target, the movement distance and target width (Fitts, 1954). This is an excellent place to start not only because it was the first research to establish Fitts' law, but also for a number of other reasons. First, strictly speaking, the parameters and variables in the Fitts' law equation must be interpreted according to Fitts' original experiment. Second, this research has been used as the research paradigm for a number of other studies (Fitts & Peterson, 1964; MacKenzie, Sellen, & Buxton, 1991). It is therefore important to describe this experiment in some detail.

Fitts' experiment was an attempt to extend information theory to human motor systems. He started with Shannon's Theorem 17 (Fitts, 1954; MacKenzie, 1992):

$$C = B \log_2(S/N + 1) \tag{1}$$

Where

C = effective information capacity of a communication channel

B = bandwidth of the channel

S = signal power

N = noise power

Fitts then suggested that the distance of movement (A for amplitude) was analogous to the signal power, and the width of the target (W) was analogous to allowable noise (Figure 2). He noted that it takes longer to hit a target further away, and it is easier to hit a larger target. With this starting point, he derived the following equation which is now known as Fitts' law:

$$MT = a + b \log_2(2A / W) \tag{2}$$

Where

MT = Movement Time

a, b = Empirically determined constants. 'a' is sometimes considered to include a ballistic time such as when a mouse button is depressed.

A = Target amplitude (distance of center of target from starting location).

W = Width of target.

The term $\log_2(2A/W)$ is described as the Index of Difficulty (ID), acknowledging that it forms a measure of the difficulty of the pointing task. The ID is measured in terms of bits, a term that owes its name to its information processing heritage. Fitts also describes an Index of Performance (IP), which is analogous to the channel capacity (C) from Shannon's theorem. We can say:

$$IP = ID / MT \tag{3}$$

IP has units of bits per second. In an idealized Fitts' law task, the constant a is zero, so IP is taken as $1/b$.

The Original Fitts Experiment

In order to test out this hypothesis, Fitts conducted an experiment using the apparatus shown in Figure 3. Subjects were asked to alternately tap a metal stylus in the central shaded areas. The instruments collected electrical impulses indicating a hit in the target area (shaded with horizontal lines) or outside (shaded with diagonal lines). The separation of the target centers was called the amplitude (A). The other dimension of interest was the width (W). The target height was 6 inches. The distance between the plates was varied between 2, 4, 8 & 16 inches, and the widths were varied between 0.25, 0.50, 1.00 and 2.00 inches. These amplitudes and widths were chosen to provide many different combinations of A/W. Note that the amplitudes and target widths are of the form

$n2^0$, $n2^1$, $n2^2$, $n2^3$, ... , $n2^m$, where n and m are positive integers. They were selected in this fashion because of the logarithmic term in ID. This pattern of increasing width and distance by powers of 2 has been followed by most subsequent studies. Two styli were used, a 1 oz stylus, and a 1 lb. stylus.

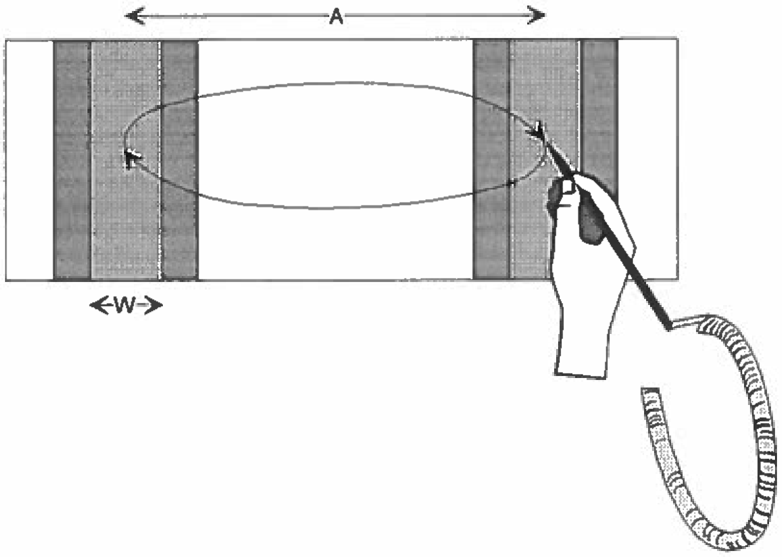

Figure 3. The original Fitts' law experiment.

Fitts found that the data matched Fitts' law (Equation 2) closely, with a correlation coefficient of 0.97. While such a high coefficient is unusual in other human factors studies, studies of pointing and pointing devices commonly have correlation coefficients higher than 0.95. The task had a performance index of ~10.5 bits/second. The weight of the stylus did not have an effect on IP.

This study was performed at the period in history when there was emphasis on time-and-motion studies aimed at workers on production lines. The kinds of tasks tested here are similar to the kinds of tasks that factory workers might face. It is also interesting to note that the experiment studied repetitive movements.

Fitts' experiment and the Fitts' law equation highlight the factors that are important in a pointing task, namely pointing speed (or the time it takes to point), target size, target distance and accuracy. Fitts' law gives us a means for comparing tasks, limbs and devices both in manual as well as in computer pointing. For example, if two pointing devices have to be compared, an experiment based on a Fitts' law paradigm (i.e. the task is similar to Fitts' task) can be used to determine their Indices of Performance (IPs). The device with the higher IP is the faster device *on the average.*

At the same time note that Fitts' law does not in and of itself give us any means of predicting the performance of a limb or device, and therefore no means of design. It is a post hoc test, but does not tell us what the performance will be without conducting an empirical study. Note also, that Fitts' law equation does not consider what happens during movement, it only looks at the aggregate movement time. In doing so, it hides what happens during movement. As we shall see, what happens during movement is key to this dissertation.

Fitts' Law and Manual Pointing

Fitts' law is an important result both because pointing is fundamental to human movement, and because it is a very accurate predictor of movement times. Psychologists have therefore tried to determine the extent of the space within which Fitts' law holds. This search space has two dimensions, limbs and tasks, and the attempt has been to determine the universality of Fitts' law within this space.

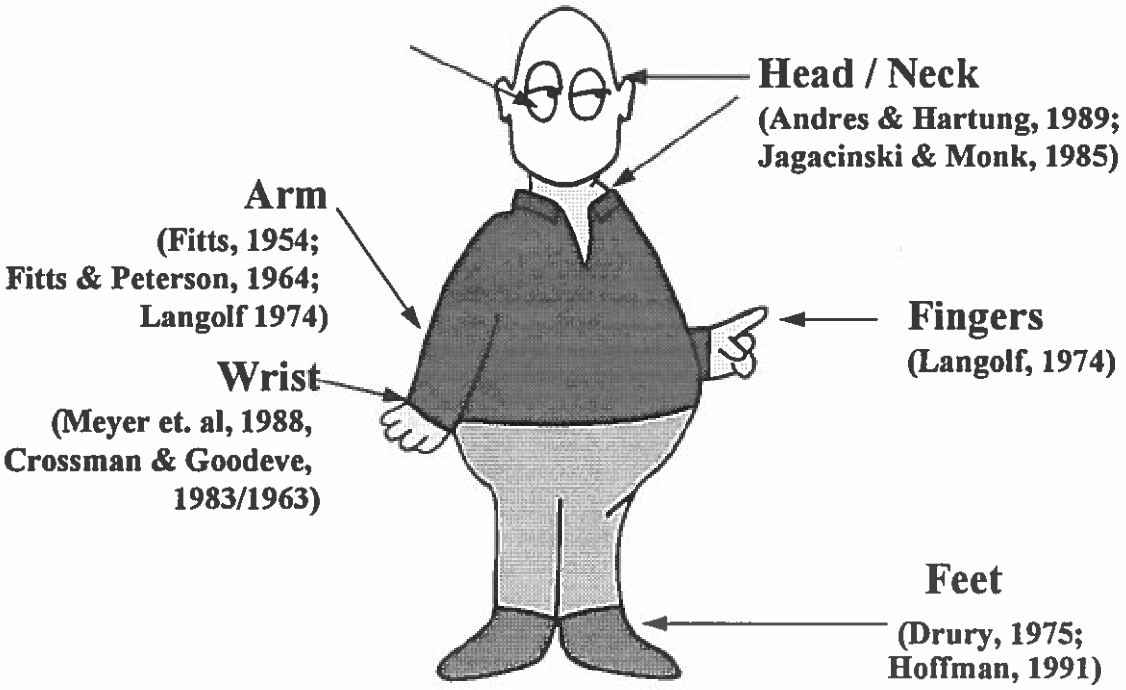

Figure 4. Limbs that have been shown to follow Fitts' law.

The Applicability of Fitts' Law to Different Limbs

Over a number of years, research has shown that Fitts' law holds for a number of limbs, as shown in Figure 4. Fitts' experiments (Fitts, 1954; Fitts & Peterson, 1964) indicated that movements of the arm followed Fitts' law. Fitts speculated whether the equation would be applicable to other limbs, and whether different limb segments might show different indices of performance. This idea was tested by Langolf who studied Fitts' law under various amplitude ranges (Langolf, Chaffin, & Foulke, 1976). Langolf's experimental setup included operations under a microscope, and extended to operations requiring movement by the whole arm. He found that

> If the reciprocal of the slope, 1/b, in the relationship MT = a + b ID infers the information processing capacity of the motor system, then the fingers showed an information processing rate of about 38 bits/sec; when the hand flexed and extended about the wrist the rate was 23 bits/sec; when the hand flexed and extended around the wrist, the rate was 23 bits/sec; and the rate for the arm was 10 bits/sec. (Langolf et al., 1976, page 118)

Langolf's studies thus extended the scope of Fitts' law to fingers, wrist, and elbow, and suggested that the IP of the smaller limbs (fingers) is higher than the IP of larger limbs (the arms). It should be noted that there is a concern with Langolf's work that it has not been replicated, and he used only three subjects.

The mobility of the head is an issue for physically handicapped persons, and a number of researchers studied head movements and showed them to follow Fitts' law (Andres & Hartung, 1989; Jagacinski & Monk, 1985; Lin, Radwin, & Vanderheiden, 1992).

In many situations, workers make extensive use of their hands and do not use their feet. This led to speculation that feet could also be used as controllers. Studies by Drury and Hoffmann showed that Fitts' law holds for feet (Drury, 1975; Hoffmann, 1991) and computer input devices called 'moles' were designed by Pearson (Pearson & Weiser, 1986; Pearson & Weiser, 1988).

A number of researchers, such as Meyer, Crossman and Goodeve, Schmidt etc., became interested in wrist rotation because they noted that a) such movement had a single degree of freedom, making it easier to model; b) because the ratio of torque to inertia is high, wrist movement would avoid filtering by the characteristics of bones and muscles involved in the movement; and c) the movement had to be stopped primarily by antagonistic muscle action, thus reducing the effect of impact from an object such as a stylus. It is therefore useful to study wrist movement in order to validate models of movement, and in the course of conducting such studies, wrist rotations have been shown to follow Fitts' law (Crossman & Goodeve, 1983; Meyer, 1988).

The kinds of actions studied in these experiments are referred to as rapid aimed movements, because the movements are rapid and aimed at a target. Research has shown that in rapid aimed movement involving the hands, the eyes start moving towards the target before the hand starts moving, and arrive at the target before the hand gets there (Abrams, Meyer, & Kornblum, 1990). In addition, the eye is in some sense the "natural pointing device" because it indicates the user's focus of attention. Researchers in HCI became interested in the eye as a pointing device, and Ware and Mikaelian showed the eye to follow Fitts' law (Ware & Mikaelian, 1987).

Tasks in the Fitts' Law Space

As explained above, researchers have extended the applicability of Fitts' law to various limbs. The other dimension along which Fitts' law has been explored is in the task being carried out. This has lead to the use of the term 'Fitts' law tasks'.

Broadly speaking, there are three kinds of movements, called Type A, B & C (Keele, 1968). Type A movements are movements that are stopped by impact with an object, such as hitting a nail with a hammer. These are the fastest movements. Type B and C movements are stopped by antagonistic muscle action, and therefore take longer. In Type B movements, the exact location where the movements stop is not important, making them faster than Type C movements. The upstroke of a hammer is an example of a Type B movement. Type C movements, which require precision, are the slowest of the three. All manual pointing movements are of Type C. The speed-accuracy tradeoff described by Fitts' law is applicable to movements of Type C.

The Fitts' law equation is a speed-accuracy tradeoff. There is an inherent tradeoff between speed and accuracy in a pointing task. The more accurate a person's movements need to be (the smaller the target), the slower they tend to become. The exact form that this tradeoff takes could be logarithmic, polynomial or linear, depending on the task. A linear speed-accuracy tradeoff is described later in this chapter.

Fitts' original experiment explored the applicability of Fitts' law to different tasks (Fitts, 1954). In addition to the pointing (stylus tapping) task, Fitts also conducted two additional experiments. One, which was similar to the Towers of Brahma problem (sometimes referred to as the Towers of Hanoi problem), had two posts with disks on them. Subjects moved the disks from one post to the other as rapidly as possible. The second experiment had two rows of holes with pins, and subjects had to move the pins from one set of holes to the other. Both tasks closely conformed to the law. Langolf, in replicating part of Fitts' work, also used a peg transfer task that conformed to Fitts' law (Langolf et al., 1976).

Fitts differentiated between discrete and continuous tasks. The task studied in his firs experiment was a continuous task. He later conducted a second experiment where the plate to be struck was indicated by a stimulus light. The subject waited with the stylus on a mark midway between the plates, and then struck the plate indicated by the light (Fitts & Peterson, 1964). There were therefore distinct pauses between each movement. This is a discrete task and it also followed Fitts' law, and had an IP greater than that from the earlier experiment, and Fitts and other researchers have speculated that the continuous experiment includes a 'dwell' time on each plate (Kantowitz & Knight, 1976).

Jagacinski studied tasks with two kinds of endpoint conditions, one where the endpoint of a trial was reached when the subject stopped on the target, and the other where the endpoint was reached when the subject held the pointer steady over the target for a specified amount of time (Jagacinski et al., 1980a). Both movements conformed to Fitts' law.

Fitts' studies were one-dimensional to the extent that the movement was along a line. Jagacinski extended Fitts' law to two dimensions by placing targets in a circle (Jagacinski & Monk, 1985). Fitts' original experiment is considered to be in one-dimension because the targets were placed in a straight line. On the other hand, the stylus moves through a 3-dimensional arc while carrying out the movement, an important distinction that will be discussed later in this chapter.

These studies combine to show that there is a wide range of combinations of tasks and limbs where Fitts' law holds.

Manual Tasks that Do Not Follow Fitts' law

A number of interesting studies have produced cases where Fitts' law does not hold, or where a variation of Fitts' law holds, such as the study by Kerr. Keele describes a study by Kerr which studied movement under water in a tapping task (Keele, 1986). Kerr analyzed the effects of target width and distance separately (Kerr, 1978), obtaining the functions:

$$MT(\text{land}) = 34\text{msec} + 111\log_2 D + 109\log_2 \frac{1}{W} \tag{4}$$

$$MT(\text{water}) = 145\text{msec} + 155\log_2 D + 115\log_2 \frac{1}{W} \tag{5}$$

Keele notes:

> Thus, while distance and width have equivalent effects for land-based movement, underwater distance has a larger effect than width, presumably because viscosity of water slows down the fast distance-covering portion of movement. (Keele, 1986, page 24)

This suggests that there are two parts to a pointing action, one fast distance covering portion, and one slower error correction portion, and the fast distance covering portion is slowed by the water. Note that distance and width play the same roles in increasing and decreasing movement time as they do in the Fitts' law equation.

Jagacinski studied Fitts' law for moving targets (Jagacinski, Repperger, Ward, & Moran, 1980b). He studied the effect of movement in the target on target acquisition time. He found that Fitts' law was not a good predictor of movement time with moving targets, and empirically derived the equation:

$$CT = c + dA + e(V+1)(\frac{1}{W} - 1) \tag{6}$$

Where

CT is the capture time in seconds,

A, V & W are amplitude, velocity and target width, and

c, d & e are the regression coefficients.

The Fitts' law equation describes a logarithmic speed-accuracy tradeoff. Under some conditions, this gives way to a linear speed-accuracy tradeoff. Schmidt et al. performed an experiment where the subjects had to make single aimed tapping movements whose distances and durations were both supposed to match specified values (Schmidt, 1979; Schmidt, Zelaznik, & Frank, 1978). They found that the speed-accuracy tradeoff could be characterized in terms of the equation:

$$S = A + B\frac{D}{T} \tag{7}$$

where

S is the standard deviation (variable error)

D is the mean movement distance, and

T is the mean movement duration.

This equation holds not only for stylus tapping, but also for wrist rotations (Meyer, 1988). Meyer notes that this equation might have the same functional status as Fitts' law when the movements have to match a temporal goal (Meyer et al., 1990).

Variations to the Fitts' Law Equation

It was pointed out in Chapter I that one area where efforts in HCI and psychology work together is in the attempt to fine tune Fitts' law. In Fitts' original experiment, he used $\log_2(2A/W)$ as the ID, primarily to ensure that the log term did not become negative for any of the movement amplitude and target width conditions he had. Since then researchers have proposed other forms of ID in order to provide a better fit to the experimental data. Three forms of Fitts' ID have been used by various researchers. Fitts' original study produced the first form (Fitts, 1954):

$$MT = a + b \log_2 (2A / W)$$

$$ID = \log_2 (2A / W) \quad (8)$$

This form was subsequently justified by Fitts saying that in attempting to hit a target of width W, a subject can make an error of at most W/2, giving a log term of A/(W/2)=2A/W as in the equation above (Fitts & Peterson, 1964). In 1960, Welford proposed the following variation (Welford, 1968):

$$MT = a + b \log_2 (A / W + 0.5)$$

$$ID = \log_2 (A / W + 0.5) \quad (9)$$

This variation was used by Card et al. (1978) in their study of pointing devices. While Fitts picked his variation for ID primarily to ensure that the value of ID was always positive for his target conditions, Welford proposed the second variation *based in part on the observation that this definition reduces the numerical value of the first constant,* a, *giving theoretical predictions of MT near zero for an ID of zero* (Fitts & Peterson, 1964). In 1992 MacKenzie proposed the Shannon variation (MacKenzie, 1992):

$$MT = a + b \log_2(A/W + 1)$$

$$ID = \log_2(A/W + 1) \quad (10)$$

Note that the equations differ only in the log term, i.e., they differ only in ID. In all forms, increasing the movement amplitude increases movement time, and increasing target width decreases movement time.

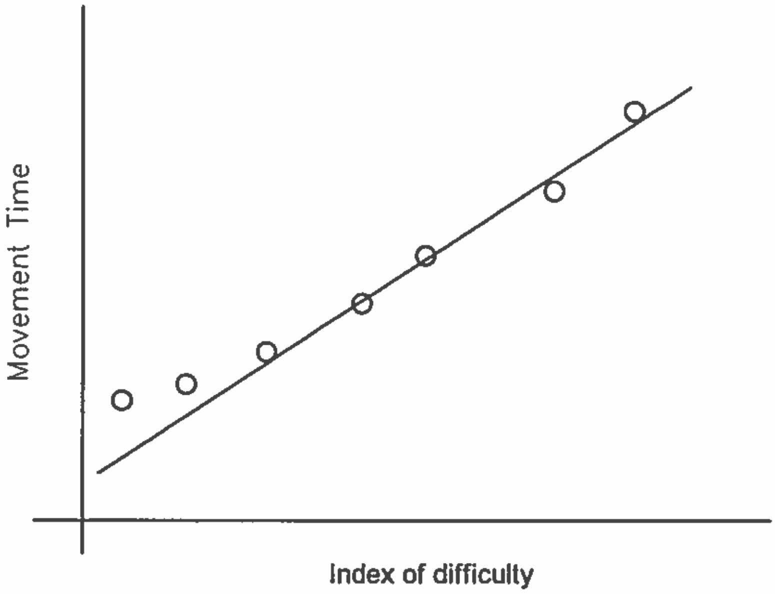

Figure 5. Curvature of data when compared to the Fitts' law regression line.

The variation MacKenzie proposed was intended to more closely match Shannon's theorem (Equation 1). He suggested that this formulation is more faithful to the information theoretic origins of Fitts' original derivation. He showed that this formulation also had a better fit to the data from a number of experiments, including Fitts' original study (Fitts, 1954) and the study by Card et al. (Card et al., 1978). In addition, this formulation has the advantage that the ID can never be negative, a condition that was possible in both the original Fitts formulation as well as in the Welford variation. The effect of Welford and MacKenzie's variations is to better fit the data at the bottom of the curve, where ID is low. The graph in Figure 5 represents how the data appears with the original formulation of ID.

Unfortunately, MacKenzie adopts the information theoretic nature of Fitts' law too literally (MacKenzie, 1992). While Shannon's work on information theory was a good analogy to apply to the problems of human movement, there is no evidence at all that the theory applies to pointing tasks. In fact, other researchers have proposed different forms of the speed-accuracy tradeoff.

For instance (as described earlier), Jagacinski and others investigating moving targets found a better fit using the equation

$$MT = c + dA + e(V+1)(\frac{1}{W} - 1) \tag{11}$$

where *c*, *d* and *e* are constants, and *V* is the velocity of movement of the target (Jagacinski et al., 1980b). If *V* is zero, the equation becomes

$$MT = c + dA + e(\frac{1}{W} - 1) \tag{12}$$

and using this equation, Jagacinski got a better fit to the data for stationary targets in their task than the fit using the Fitts' law equation (Jagacinski et al., 1980b).

Kvalseth suggested a relationship of the form

$$MT = aA^bW^c \tag{13}$$

(Kvalseth, 1981) which takes the form

$$MT = a(A/W)^b \tag{14}$$

when $b = -c$. Meyer notes that values of p ranging between 0.25 and 0.5 often have a better fit to target acquisition data than Fitts' equation (Meyer et al., 1990). These three contrasting models were studied by Epps, and he found Kvalseth's model to have the best r^2, though the Fitts and Jagacinski models had better mean squared error and predictive sum of squares (Epps, 1986).

Attempts to fine tune the Fitts' law IP are perhaps premature because we are still unsure what factors affect the speed accuracy tradeoff. As Meyer notes, it is likely that differing tasks will have a better fit to one or other form of the Fitts' law equation, so finding a form that has a better fit to some subset of data is likely to yield results that are not generalizable to other data (Meyer et al., 1990).

Models Underlying Fitts' Law

There are a large number of mathematical models that predict Fitts' law, and without anchoring the variables to psychomotor models, there is no way of determining the validity of any one. Psychologists have therefore developed models that describe the manner in which people plan and execute movements. The framework in Chapter I noted that all the research on modeling pointing behavior has been done by psychologists, and researchers in HCI have neither contributed to this area of research, nor have they applied psychomotor models to designing pointing devices. In this sub-section I discusses various theories of movement that have been suggested, starting with work from the previous century by Woodworth, and concluding with current work by Meyer and others.

Psychologists have long understood the importance of a model describing the kinematics of movement. This problem has received the attention of modern psychology as far back as the 19th century. Once investigators started examining the basis for Fitts' law, the information theoretic reasoning that Fitts used was rejected as a model of movement, and various other theories were developed.

Initial Adjustment and Current Control

Some of the earliest work on models of movement was done by Woodworth in 1899 (Meyer et al., 1990). Woodworth thought that rapid movement had two phases, which he called initial adjustment and current control. The initial adjustment phase (also called ballistic movement) transports a selected part of the body quickly towards a target location. The current control phase (also called feedback control) subsequently corrects any errors made along the way, using sensory feedback to reach the target accurately (Meyer et al., 1990). Based on experiments that compared movement accuracy of subjects with closed eyes against subjects with open eyes, Woodworth estimated that the initial adjustment takes place in the first 300 ms. of a movement, and the bulk of time required for movement is taken up by the current control phase.

Fitts' work with Shannon's theorem from information theory was the next major model of rapid aimed movements to emerge. Although his empirical results were easy to replicate, information theory has not been well received as a theoretical framework of movement control. Critics felt that the information theory hypothesis was strained at best, and totally wrong at worst, and this triggered a search for other ways of explaining the logarithmic speed-accuracy tradeoff (Meyer et al., 1990).

Deterministic Iterative Corrections Model

In work originally presented in 1963, and subsequently re-published in 1983, Crossman & Goodeve raised doubts about Fitts' idea that there was noise in the initial adjustment phase of rapid movements (Crossman & Goodeve, 1983; Meyer et al., 1990). They note that there is an empirical difficulty of establishing the existence of the postulated "noise" or initial uncertainty. ... Thus the supposed "noise" is apparently not present in the effector system. (Crossman & Goodeve, 1983 page 253).

They proposed an alternative model which Meyer calls the deterministic iterative-corrections model (Meyer et al., 1990). Movements under this model are depicted in Figure 6. The horizontal axis represents movement distance, and the vertical axis represents movement velocity. The curves correspond to successive hypothetical submovements between an initial home position (distance = 0) and a target region bounded by vertical lines whose center is A units from the home position and whose width is W units (Meyer, 1988 page 343).

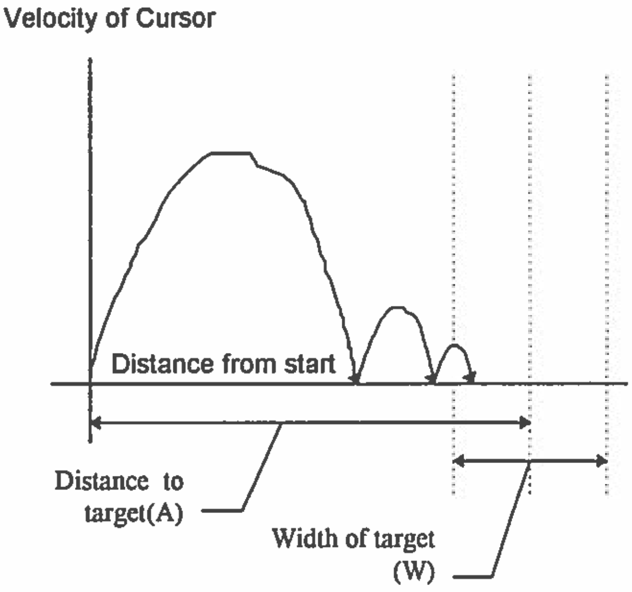

Figure 6. Movement characteristics of the deterministic iterative-corrections model.

The model assumes that a movement to a target is made up of a number of submovements, and that each submovement covers a constant portion (p) of the distance between the starting location and the target. For instance, the first submovement would travel a distance pA, the second would travel a distance pA(1-p), the third would travel a distance $pA(1-p)^2$, and so on. The model's second assumption is that each of these submovements takes the same amount of time. The third assumption is that the submovements are guided by sensory feedback, either visual or kinesthetic. Movement ends when the distance to the target is less than W/2, where W is the width of the target. This gives a function representing movement time in terms of p, A and W, which has the same form as Fitts' law. Keele suggested a similar model (Keele, 1968). This model was rejected later in studies.

Second-Order Under-damped Function

In 1976 Langolf suggested using control system theory as an alternative to information theory as a means of explaining Fitts' law (Langolf et al., 1976). He noted that a second-order under-damped system with a constant damping ratio ξ will settle according to the equation

$$MT = \frac{1}{\xi \varpi_n} \ln\left(\frac{2A}{W} \frac{1}{\sqrt{1-\xi^2}}\right) \tag{15}$$

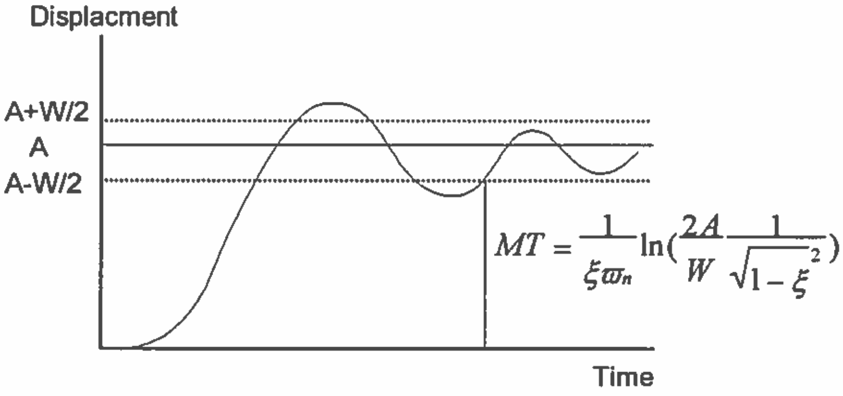

Figure 7. Movement characteristics predicted by an under-damped second order system. From (Langolf et al., 1976 page 115).

He pointed out that such a system shows an oscillatory behavior, frequently characteristic of human movements. This function has a time versus displacement function as shown in Figure 7. The figure shows how a system with damping ratio ξ will settle towards a signal of amplitude A. MT is the time after which the oscillations in the signal are within W/2 displacement units of the target distance A. This theory has not gained much popularity because it does not appear to describe the movement trajectories gathered in various studies, such as the one by Jagacinski described below. Langolf's model suggests that all movement will be characterized by a large initial movement, followed by a movement back to the start, followed by a movement towards the target, oscillating in this fashion till the oscillations are completely contained in the target. These characteristics were not found when movement data was gathered.

Dichotomous Movement Models

Jagacinski et al. studied the microstructure of movement for joysticks employing position control and velocity control (Jagacinski et al., 1980a). Position and velocity control are discussed later in this chapter. Briefly, a position control joystick controls the position of the cursor, while a velocity control joystick controls the velocity of the cursor. Their experiment was aimed at determining which of two classes of models describing movement was valid. One class of movement models suggests a "unitary convergence pattern," i.e., that all the submovements that form a movement towards a target are similar in form, (Crossman & Goodeve, 1983; Keele, 1968). These are the models described earlier as deterministic iterative corrections models and the second-order under-damped function. For example, in Crossman & Goodeve's model, the movements are of constant duration and accuracy. The other class of movement models suggests that there is an initial ballistic movement, initiated by a motor program, and subsequent movements are made under feedback control (also called current control). These are the kinds of models suggested by Woodworth.

In order to determine which of the two classes of models was accurate, Jagacinski varied the ending criterion. One required subjects to maintain the cursor over the target for a specified time (350 ms.), and the other required subjects to hold position *as well as* keep the velocity of the cursor below a specified level (.5 degrees/sec). Jagacinski hypothesized that if the unitary movement models were correct, there would be a significant change in the structure of the entire movement, while if the dichotomous model was correct, there would be a change only in the end of the movement.

To examine the structure of the movements, plots of cursor velocity against cursor position were made. The results showed that the trajectories did not resemble straight line, parabolic or spiral shapes that would be expected from linear first or second order continuous systems. The movements usually consisted of a series of irregular episodes of acceleration and deceleration, which Jagacinski referred to as sub-movements (Jagacinski et al., 1980a).

In order to be able to measure and quantify these submovements, Jagacinski used a computer program to parse the input data to determine the beginnings and endings of submovement durations. A submovement was considered terminated when (a) the velocity changed sign, or (b) the acceleration changed sign. The submovements thus parsed were tabulated in terms of duration and accuracy.

He defined submovement accuracy as

$$accuracy = \left| \frac{dist_{final}}{dist_{start}} \right| \tag{16}$$

where $dist_{final}$ is the distance to the target at the end of the movement, and $dist_{start}$ is the distance to the target at the beginning of the movement.

The results clearly indicate that there is a marked difference between the first submovement and the rest of the submovements. The first submovements were both slower and more accurate than the rest of the submovements. This dichotomy is more marked for the velocity control system than for the position control system.

The results also showed that the main effect of the position plus steadiness criterion was to increase the number of submovements that the subjects made at the end of the movement, which is illustrated by Table 1 below.

Table 1. Number of submovements required for different control mechanisms and termination conditions.

	Number of Submovements	
	Position Criterion	Position Plus Steadiness Criterion
Position Control	2.1	3.1
Velocity Control	2.5	3.0

The result that the first submovement was quite different from the subsequent movements rejects the deterministic iterative corrections models (Crossman & Goodeve, 1983; Keele, 1968). The finding also rejects Langolf's second order under-damped system (Langolf et al., 1976).

Stochastic Optimized Submovement Model

The stochastic optimized submovement model is the most elaborate model of movement currently available. By making a few simple assumptions, it attempts to predict a number of observed behaviors. The model has been proposed by Meyer et al., who argued that the deterministic iterative corrections model could not account for a number of experimental results (Meyer, 1988; Meyer et al., 1982; Meyer et al., 1990), in addition to those mentioned above from the Langolf and Jagacinski studies.

The deterministic iterative corrections model postulates multiple submovements till the subject reaches the target, but a number of studies (Jagacinski et al., 1980a; Langolf et al., 1976; Meyer, 1988) found many occasions when there are only one or two submovements, even when the ID is very large. Submovements measured in these experiments have not been found to have constant duration. In addition, submovements do not travel a constant portion of the distance to the target, and the initial submovement tended to be much longer and much more accurate than the subsequent movements (Jagacinski et al., 1980a).

The deterministic iterative corrections model does not account for errors, but errors do occur (Fitts, 1954). A number of researchers have found that the error rate increases with increasing ID (Fitts & Peterson, 1964; Meyer, 1988, Douglas, 1993). Finally, Meyer notes that there are a number of cases when the speed-accuracy tradeoff is not logarithmic such as in the study by Jagacinski with moving targets (Jagacinski et al., 1980b).

In order to explain these differences, Meyer et al. (Meyer, 1988; Meyer et al., 1990) have suggested the stochastic optimized submovement model (SOS model), which is illustrated in Figure 8. The model brings back the notion of noise in the motor system, which was originally introduced by Fitts (1954), then rejected in subsequent studies (Crossman & Goodeve, 1983). The model is influenced by the impulse-variability model (Schmidt, 1979; Schmidt et al., 1978) which was originally described for time matching movements such as making movements in time with some signal such as a metronome (Meyer et al., 1990).

The model makes a number of assumptions. These are listed below, based on there being at most two submovements. The more general model is an extension of this model, and does not place a limit on the number of submovements.

Primary submovement. A movement begins with a primary submovement which is programmed to hit the center of the target region (illustrated by the middle trajectory of Figure 8). If the submovement is successful, the action terminates.

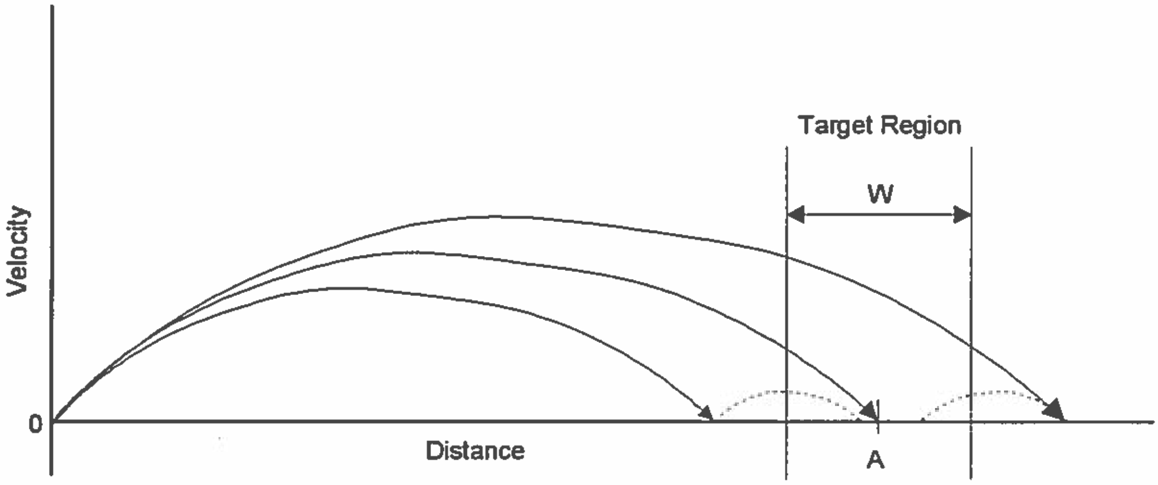

Figure 8: Movement characteristics predicted by the Stochastic Optimized Submovement Model.

Motor noise and secondary submovements. Noise in the motor system may cause the primary submovement to miss the target. This might cause a secondary corrective submovement to be made, based on sensory feedback. The secondary movement will in most cases hit the target, but on occasion, it will miss. In this limited case of the model, movement will cease, although as many as five submovements have been noted (Crossman & Goodeve, 1983). In most cases, there are only 2 submovements as shown by Table 1.

Effects of noise on submovement endpoints. The model assumes that errors due to motor noise will increase with increasing velocity. In particular, primary submovements are assumed to have endpoints whose standard deviation (S_1) in space is proportional to the average primary submovement velocity (V_1):

$$S_1 = KV_1 = K\frac{D_1}{T_1} \tag{17}$$

where D_1 is the average distance traveled by the primary submovement, and T_1 is the average time of the primary submovement. This assumption is based on studies that have shown that the variable error is a function of average movement velocity in a time-matching movement task (Schmidt, 1979).

Similarly, for the secondary submovement, we have:

$$S_2 = KV_2 = K\frac{D_2}{T_2} \tag{18}$$

Minimization of movement durations. The last key assumption is that the average velocities of the primary and secondary submovement are programmed to minimize the average total movement duration *(T)*. Meyer notes that this assumption stems from the demands of the typical time-minimization movement task, where subjects are presumed to try and reach the target region as quickly and with as few errors as possible.

The model makes a number of predictions. It predicts that the average total movement duration *(T)* is a quasi square-root function of the target distance - width ratio (Meyer, 1988):

$$T = 2K\frac{2\theta\sqrt{\frac{D}{W}} - \sqrt{\frac{W}{D}}}{\theta\sqrt{\theta - \frac{W}{D}}} \tag{19}$$

This function is very closely approximated by the function (Meyer et al., 1990)

$$T = A + B\sqrt{\frac{D}{W}} \tag{20}$$

Meyer et al. note that they expect Equation 20 to come closer to the actual movement durations than a logarithmic equation. Kvalseth noted a better fit of the data to this equation (Kvalseth, 1981). Epps noted that the equation fit better than Fitts for some cases, but not for others (Epps, 1986).

The model predicts that the primary submovement duration will also follow an equation like Equation 19, and can be approximated by an equation similar to Equation 20. It predicts that the proportion of trials on which there are secondary submovements will increase with D/W.

> When target distance increases or width decreases, the primary submovements should slow down somewhat, but the decrease in their velocity would not be sufficient to maintain a constant rate of target hits. Instead, given the nature of the optimization, the primary submovements should miss the target more often, thereby requiring more corrections through secondary submovements. (Meyer et al., 1990 page 207)

The model also makes a prediction that error rates should increase monotonically with D/W.

Meyer et al., have studied the SOS model for wrist rotations, and found its predictions to be remarkably accurate. Thus, the model is possibly the best theory for explaining Fitts' speed accuracy tradeoff, that is available today. It is important to note, however, that the SOS model was developed for isotonic movement, and even among its authors, there is disagreement as to whether it will work for isometric movement.[1]

---

[1] Personal communication.

Thus we have seen that pointing movement has been extensively studied by psychologists, and the regularity described by Fitts' law has been observed in a wide range of movement types, and with a number of limbs. Psychologists have also worked on fine tuning Fitts' law, as well as on developing models to explain Fitts' law.

This completes the half of the framework described in Figure 1, that describes the work done by psychologists. The next section examines the other half of the framework, examining the work done on pointing with a computer pointing device.

Fitts' Law and Computer Pointing Devices

The framework developed in Chapter 1 suggested that the work done in HCI on pointing devices can also be thought of as an attempt to explore the space within which Fitts' law holds. The dimensions of this space are device and task, and research attempts to determine which combination of tasks and devices follow Fitts' law. I now discuss the results of this research.

Note the similarities and differences of this space and the space within which psychologists have explored Fitts' law. While the psychologists studied tasks and limbs, HCI researchers have studied tasks and devices. Some of the tasks that can be studied on a computer have direct counterparts in the real world, while others do not. For example, while pointing exists in both areas, clicking and dragging are computer based tasks, and do not have ready equivalents in the real world. Computers can mediate the relationship between the input variable (device movement) and the output variable (cursor movement) in manners that are not possible with manual movement. The kinds of tasks performed with computer pointing devices include serial and discrete tasks, one-dimensional tasks, two dimensional tasks, pointing, dragging, drawing etc.

Computer Pointing Tasks in the Fitts' law Space

To begin exploring the space of where pointing devices follow Fitts' law, let us start with tasks that follow Fitts' law. Pointing tasks on a computer vary along three dimensions:

Whether the task is discrete or continuous.

What task is performed (e.g. pointing or dragging).

What is the number of dimensions involved?

Discrete vs. Continuous Tasks

As described earlier, pointing tasks can be discrete or continuous. Continuous tasks are performed one after another, without pauses between pointing actions, like Fitts' first experiment (Fitts, 1954). The task in Fitts' second experiment was discrete, with pauses between each pointing action (Fitts & Peterson, 1964).

Actual computer use spans the spectrum from discrete to continuous pointing actions, including what might be described as semi-continuous actions. When a user of a word processor moves the insertion point from one location to another, he or she has performed a discrete pointing action. Most pointing tasks, however, are made up of a series of pointing actions. For example, selecting an item from a menu requires two pointing actions, one to pull down the menu, and one to select an option from it. If this action was preceded by an action which selected an object, then the command takes three continuous pointing movements. In this manner, pointing with a computer typically is made up of semi-continuous pointing actions. The use of arcade-type computer games typically requires continuous pointing movements, and so computer use runs the gamut from discrete pointing movements, to continuous pointing movements.

Fitts showed that his equation was applicable to both continuous and discrete tasks (Fitts, 1954; Fitts & Peterson, 1964). Both kinds of movements when performed on a computer have been shown to follow Fitts' law. Discrete tasks have been the subject of a number of studies, and (Card et al., 1978; Douglas & Mithal, 1993; Jagacinski & Monk, 1985; MacKenzie & Buxton, 1992) are representative. Continuous tasks have not commonly been studied, an exception being a study by MacKenzie et al. who replicated Fitts' original experiment on a computer (MacKenzie et al., 1991).

Tasks Differentiated by GUI Primitives

In addition to discrete and continuous tasks, the interaction primitives of GUIs also create categories of tasks such as pointing, clicking, dragging and drawing. The most common action done with a GUI is pointing which in its purest sense is the action of moving the cursor from one location on the screen to another. Users point to an object by placing the cursor 'hot spot' on top of it. In most GUIs, it is common to indicate the end of the pointing action by clicking a button. This button is often, but not always, located on the pointing device, such as the buttons on a mouse or trackball. Most of the studies involved a pointing task that was terminated by a click, and have shown such tasks to follow Fitts' law (Card et al., 1978; Epps, 1986; MacKenzie et al., 1991).

Not all GUIs require the button to be clicked in order to indicate the object selected. For example, in the X Windows interface, the active window is the one beneath the cursor, and users do not have to click to point at a window. Jagacinski et. al showed that such movements followed Fitts' law (Jagacinski et al., 1980a). In addition, if we can assume, as is done in the Keystroke Level Model, that the button click takes a constant time (Card, Moran, & Newell, 1980; Card et al., 1983), then, it seems reasonable that the pure pointing action also follows Fitts' law.

When the pointing device is moved with the selection button held down, the action is called dragging. Dragging is a basic action in GUIs and is used to indicate operations such as moving an object to another location, and selecting a group of objects. MacKenzie showed dragging to follow Fitts' law (MacKenzie et al., 1991). One of the devices he studied was the trackball, and he noted that it performed poorly for dragging apparently because it was awkward to hold the button down and move the ball at the same time. This is probably the reason why some trackballs have a "drag lock" feature. This requires two button clicks, one to initiate dragging, and one to terminate it, thus leaving the fingers free to move the ball.

Meyer and Walker did a study of the applicability of Meyer's Stochastic Optimized Submovement Model to mouse movement (Walker et al., 1993). The pointing task that they studied had the button down for the duration of movement, making it a dragging task, although the user did not get the visual feedback usually associated with a dragging action. The task had a high correlation to Fitts' law, and thus they indirectly showed that dragging is a Fitts' law task.

In addition to dragging, there are other cases where the pointer is moved with one or more buttons pressed. For example, "option-drag" is used on the Macintosh interface to copy a file. Option-drag is a dragging action performed with the "Option" key on the keyboard held down. There are numerous such "chording" actions, none of which have been studied, though it is reasonable to assume that they all follow Fitts' law.

In pointing and dragging, the path through which the cursor moves is not important. When that path becomes important, we have cases of drawing, such as drawing with a stylus on a touch tablet. Drawing has not been the subject of a Fitts' law analysis, but it has been suggested that it would be similar to a time-constrained movement that has a linear relationship between time and distance (Meyer, 1988; Schmidt, 1979; Schmidt et al., 1978).

Pointing Devices and Task Dimensions

The third way of categorizing computer pointing tasks is the number of dimensions in which the task is performed. MacKenzie replicated the one-dimensional study by Fitts, showing that the one-dimensional task conformed to Fitts' law for a number of devices (MacKenzie et al., 1991). Much of Jagacinski's work establishing Fitts' law for joysticks has been in one dimension (Jagacinski, Hartzell, Ward, & Bishop, 1978; Jagacinski et al., 1980a; Jagacinski et al., 1980b).

Manual pointing actions are inherently three-dimensional. When a pointing device such as a mouse or touch tablet are used, the hand controlling the device moves in three-dimensional space. When this is translated onto the screen for a GUI, the movement occurs in two dimensions. A number of studies that established Fitts' law for pointing devices were conducted in two dimensions (Card et al., 1978; Gillan, Holden, Adam, Rudisill, & Magee, 1990; Gillan, Holden, Adam, Rudisill, & Magee, 1992). These studies show that Fitts' law holds in a two-dimensional task.

At least four studies explicitly looked at establishing Fitts' law in two dimensions (Boritz, Booth, & Cowan, 1991, Douglas, 1993 #1612; Jagacinski & Monk, 1985; MacKenzie & Buxton, 1992). The studies by Jagacinski and Boritz established that the Fitts' IP appears to vary from angle to angle. For example, Boritz found that movements with the right hand moving from left to right were the fastest, and movements moving in towards the body were the slowest. Both Jagacinski and Douglas study found variations in IP, but their results did not match those by Boritz, and while it appears that IP varies with angle, the exact nature of this variation is unclear.

MacKenzie's work had a different focus, to determine the exact nature of the target width in the Fitts' law equation when the task takes place in two dimensions(MacKenzie & Buxton, 1992). His experiment had subjects point to rectangular targets, and then used one of three widths in the Fitts' law analysis. These were the "canonical" width, which is the width of the target in the horizontal direction, the width of the target in the direction of movement, and the greater of the width and height. He found that when he used the canonical width as the width in the Fitts' law equation, he got a poor fit to the data. However, using either of the other two formulations, Fitts' law was a good predictor of movement time.

While current GUIs are two dimensional, three dimensional interface constructs are beginning to appear in a number of ways. First, virtual reality exists in three dimensions, and Fitts' law in those three dimensions takes on meaning. In addition, many computer games such as flight simulators and the ubiquitous shoot-em-up games exist in three dimensions. The other way in which three dimensions are beginning to appear is in the manner of layered interfaces, where a document can have layers, and there is a need to move in the *z* axis. While a number of input devices (mostly joysticks with a *z* wheel) exist, there is no published research that establishes Fitts' law in three dimensions. A possibility is a thesis in progress at the University of Toronto, titled Input Techniques for HCI in 3D Environments, by Shumin Zhai. It is about the performance evaluation of isometric, isotonic and elastic input devices in 6 degree-of-freedom docking and tracking tasks. Early results indicate that Fitts' law holds in three dimensions, but the thesis has not yet been completed.

Pointing Devices in the Fitts' law Space

Human factors researchers have tried to establish the space within which Fitts' law holds for pointing devices. As described above, one of the dimensions of this space is the kind of task performed. This subsection describes the other dimension, the kinds of devices used.

Figure 9 represents a system block diagram of a generic pointing device. The user manipulates the device, and a transducer in the device converts the user's input signal into a device output signal, which in turn is an input to the device's software driver whose output is the position of the cursor on the screen.

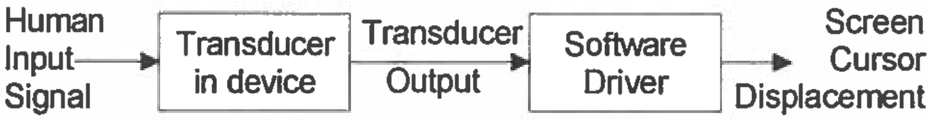

Figure 9. Generic block diagram of a pointing device.

Discrete and Continuous Pointing Devices

Card et al. (1978), in the earliest HCI study showing the importance of Fitts' law for computers, divided pointing devices into two categories, discrete and continuous.

Discrete devices allow step movements of the cursor, such as cursor control keys and step keys. Continuous devices like the mouse, joystick and touch tablet, allow continuous movement of the cursor. Card found that Fitts' law was applicable to movement with continuous devices, but not with discrete devices.

The Card et al., study also showed that the error rates were the lowest and the rates of learning were highest for the continuous devices, with the mouse performing better than the joystick. This finding was one of the factors that lead to the wide acceptance of the mouse in contemporary GUIs.

Mouse

Among continuous devices, the mouse is probably the most highly studied, with a large body of literature establishing that its movement time follows Fitts' law (Boritz et al., 1991; Card et al., 1978; Douglas & Mithal, 1993; Douglas & Mithal, 1994; Epps, 1986; MacKenzie, 1992; MacKenzie et al., 1991; Walker et al., 1993).

Joystick

Joysticks are used in military aircraft and in video games, and they have also been well studied. They are unusual among pointing devices in that there are many types of joysticks, as illustrated in Figure 10.

Joysticks can be categorized in three ways, depending on whether they are isometric or isotonic, position controlled or rate controlled, and spring loaded or free.

Isometric joysticks do not move when pressure is applied to them, and use transducers (see Figure 9) such as strain gauges or varistors to sense the amount of force applied to them. The force sensed is turned into an output signal. Isotonic joysticks move when pressure is applied to them. Isotonic joysticks can be spring loaded or free depending on whether a spring is utilized to return the joystick to the central position. Isotonic spring loaded joysticks can measure either displacement or force, while free joysticks measure displacement.

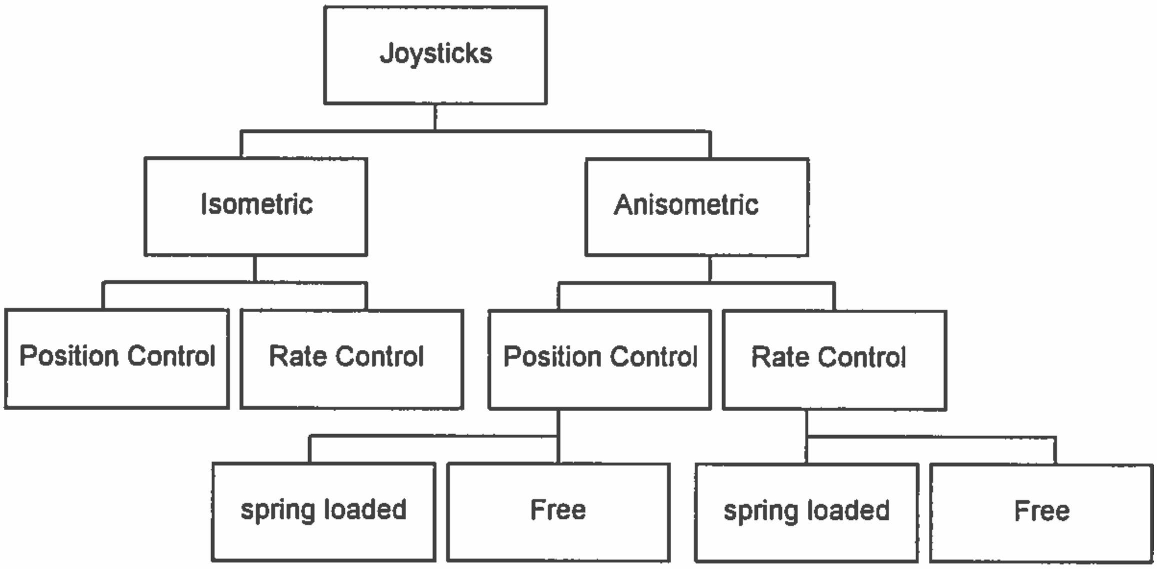

Figure 10. The different types of joysticks.

If the device driver maps the transducer output signal into a displacement on the screen, the joystick is position controlled. The device driver can also map the transducer output signal into a velocity of movement of the cursor, in which case the joystick is rate or velocity controlled.

The following types of joysticks have been shown to follow Fitts' law:

1. Isometric rate control (Card et al., 1978; Douglas & Mithal, 1993; Douglas & Mithal, 1994; Epps, 1986; Kantowitz & Elvers, 1988; Rutledge & Selker, 1990)
2. Isometric position control (Kantowitz & Elvers, 1988)
3. Isotonic position control (Epps, 1986; Jagacinski et al., 1980b)
4. Isotonic rate control (Epps, 1986; Jagacinski et al., 1980b)

The literature does not consistently report whether the isotonic joysticks used were spring loaded or not. While not all kinds of joysticks have been tested, it appears likely that Fitts' law will hold for all the kinds of joysticks listed above.

Jagacinski, performed a series of studies on joysticks, contrasting conditions such as position control versus velocity control, endpoint conditions, and moving vs. stationary targets. Based on the studies by Jagacinski, Epps and Kantowitz, we can make two following generalizations. First, isotonic joysticks have a higher IP than isometric joysticks (Epps, 1986). Second, position control joysticks have a higher IP than velocity control joysticks (Jagacinski, 1989; Jagacinski et al., 1980a; Jagacinski et al., 1980b; Kantowitz & Elvers, 1988).

Trackball

Trackballs are used in air traffic control applications, medical equipment, CAD applications and in laptop and notebook computers. Studies by Epps and MacKenzie have shown them to follow Fitts' law (Epps, 1986; MacKenzie et al., 1991).

Touch Tablet

Epps and MacKenzie et al., studied the touch tablet as an input device and showed it to follow Fitts' law (Epps, 1986; MacKenzie et al., 1991). The touch tablet can be operated in two modes, relative and absolute. In relative mode, the user puts the pen down on the tablet and then moves it in the direction that the cursor should move. The cursor is displaced on the screen an amount that is proportional to the displacement of the pen. This mapping is very similar to the mapping of the displacement of a mouse on a mouse pad to the displacement of the cursor on the screen. Relative mode touch tablets have started appearing in notebook computers since the middle of 1994 when they were introduced by Apple in their PowerBook computers.

In an absolute mode touch tablet, there is a one-to-one mapping of the surface of the touch tablet to the screen. For example, the top left corner of the touch tablet corresponds to the top left corner of the screen. Placing the pen in this corner immediately moves the cursor to that position.

The Epps study included both a relative mode tablet, as well as absolute mode tablet, and showed the relative mode tablet to have a higher IP than the absolute mode tablet. This might have been because of the tablet technology available at the time. It is possible that for modern tablets with cordless pens, absolute mode tablets will be faster than relative mode tablets.

Other Continuous Devices

Three kinds of continuous devices, light pen, light gun, and digitizer pucks were studied to determine their pointing speed, but a Fitts' law analysis was not done on them (Goodwin, 1975; Rosenberg & Martin, 1988). However, it seems likely that Fitts' law holds for them because pointing with them is continuous, and is similar to pointing with a finger.

This body of literature demonstrates that the Fitts' law is applicable to a wide range of computer pointing devices. We now discuss the manner in which the computer affects the experimental condition in a Fitts' law task.

The Relationship Between Computers and the Study of Pointing

The computer provides researchers with an unusual medium for studying human behavior, because it can be used not only to present the experimental task, but also gather the data, and even to analyze it. In studying pointing movement, the match is incredibly good, because pointing devices closely mimic human movement in the real world and the computer becomes the experimental medium. I believe that this match is the reason that so many computers have played a role in Fitts' law studies, but it is not without problems. This section examines the relationship between computers and pointing, and discusses the effect of computer mediation.

The computer is only a representation of experiments such as Fitts', which in turn is a simplification of real movement. Any results are only as good as the representation is faithful. Computer based experiments, even those that closely mimic Fitts' study, such as those by MacKenzie et al., differ from studies of real movement on two significant points (MacKenzie et al., 1991).

First, Fitts' experiment took place in three-dimensional space even though the placement of the target varied only in one dimension. Computer based experiments take place in a two-dimensional space, the computer screen. Movements with a device such as a mouse on a mousepad take place on a two-dimensional plane, while the stylus used by Fitts' subjects described arcs in three-dimensional space.

Second, Fitts' experiment was, for want of a better word, immediate, i.e., subjects were moving the stylus in order to make the tip of the stylus connect with a target plate. In computer-based experiments, the object that subjects are moving, e.g., the mouse, is not the focus of their attention (see Figure 9, where the mouse would be used to generate the input signal). The focus of attention is the cursor on the screen (the screen cursor displacement in Figure 9).

There is also no direct mapping between the action performed on the pointing device and the movement of the cursor on the screen (indicated by the transducer and software driver in Figure 9). This might be as simple as a linear mapping from displacement to displacement, as is normal for most mice, but it need not be so. In fact, devices such as power mice use non-linear mappings between mouse displacement and cursor displacement in an attempt to increase performance (Jellinek & Card, 1990), while devices such as isometric joysticks map force into velocity.

The importance of this is that while computers have played a valuable role in understanding Fitts' law and movement, the differences between computer-based movement (virtual movement) and real-world movement (real movement) should be kept in mind.

Fitts' IP Does Not Map Across Studies

One difference between Fitts' law studies on manual pointing and computer pointing is that generally speaking, two experiments comparing manual pointing with the same limb are comparable. The same is not true of computer pointing tasks. For example, an earlier study that compared an isometric joystick with a mouse found that joystick was approximately 50% slower than the mouse (Douglas & Mithal, 1993; Douglas & Mithal, 1994). This was surprising because in the Card et al. study, the joystick was only 22% slower than the mouse (Card et al., 1978). MacKenzie provides a framework to understand this. He lists seven factors can affect the average movement time and IP obtained in a study. These are the task, the selection technique, the range of conditions, choice of model, the approach angle and target width, the device, error handling and learning effects (MacKenzie, 1992).

The tasks that are studied in Fitts' law experiments have an effect on the IP. As discussed earlier, tasks can be serial or discrete, and discrete tasks are faster. The selection technique used to indicate the end of a trial also makes a difference. In some studies, the end of the trial was indicated by a button press, e.g. the space bar (Card et al., 1978), and the mouse button (Douglas & Mithal, 1993). In one joystick study (Jagacinski et al., 1980a), the termination criteria required the subjects to hold the cursor steady over the target.

There is some evidence that two handed input might speed performance (Buxton & Myers, 1986). Some of the tasks have been two handed, e.g. the key joystick, while other have been one-handed, e.g. most mice (Douglas & Mithal, 1993).

MacKenzie suggested that while it is not possible to compare devices across experimental conditions, the ratio of IPs within a study can be compared across studies (MacKenzie, 1992). Based on this, Douglas and Mithal (1993) compared their data to data reported in other studies that compared mice and isometric joysticks (Card et al., 1978; Douglas & Mithal, 1994; Epps, 1986; MacKenzie, 1992). The results are summarized in Table 2.

Table 2. The IPs for mouse and joystick from three studies, and their ratios (Douglas & Mithal, 1994).

	Card et al.	Epps	Douglas & Mithal
IP mouse	10.4	2.6	4.15
IP joystick	4.5	1.2	1.97
mouse:joy	2.31	2.16	2.10

Thus, the ratio of IP for mouse and key joystick are comparable to those for devices studied by other researchers, indicating that while it might not be possible to compare IPs across studies, the ratios of the IPs of devices across studies can be used to compare them.

Pragmatic Considerations for Pointing Devices

MacKenzie noted that differences in devices are one of the factors that contribute to differences in IP (MacKenzie, 1992). These differences are important, because they might indicate that a psychomotor model of movement that is applicable to one device might not be applicable to other devices. There are a number of factors that can differ between devices. For example, some joysticks control position, while others control velocity. This sub-section, discusses differences between pointing devices.

Buxton, suggested that there are lexical and pragmatic factors to be considered regarding input devices (Buxton, 1983). He suggested that devices be categorized along three dimensions: the property sensed, the number of degrees of freedom, and a third dimension which captured whether the device moved, or was touch sensitive. For example, a touch tablet senses position in two dimensions by touch, while a light pen senses position in two dimensions by non-touch sensitive means (electronic in this case). A mouse senses motion in two dimensions by measuring mechanical movement.

More recently, Mackinlay, in work inspired by a language of graphical primitives developed by a cartographer, developed a representation for input devices that could be used to categorize all existing input devices, as well as predict novel input devices (Card, Mackinlay, & Robertson, 1990; Mackinlay, Card, & Robertson, 1990). He also used three dimensions in his taxonomy. The first dimension was the property sensed, which was either force, difference in force, position, or difference in position. The second dimension was the number of dimensions of the property that were sensed, and the third dimension was whether a rotary or linear movement was detected. For instance, a trackball measures a difference in position in two spherical coordinates.

While these taxonomies are useful for categorizing devices, none is particularly useful in helping us understand pointing devices at a psychomotor level. I believe that this is because not enough attention has been paid to the pragmatics of the device, in the sense used by Buxton.

The mechanisms involved in pointing are illustrated in Figure 11. Each one of these parts plays a role in defining an overall transfer function between the input at 'A' and the displacement of the cursor on the screen at 'E'. The factors are:

1. A — The property being controlled by the user
2. B — The transfer function of the transducer
3. C — The property being sensed by the software driver
4. D & E — The transfer function of the software driver, and therefore the parameter in the cursor being controlled (e.g. displacement or velocity)
5. The overall transfer function between the input parameter and the output parameter
6. Whether or not the device is isometric.

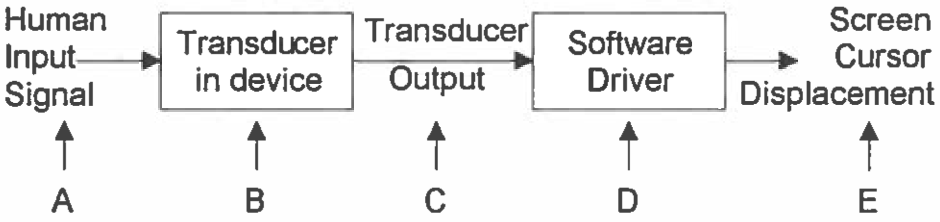

Figure 11: The important parts in a computer pointing system

All these factors except the last factor are represented in Figure 11. These factors are discussed below.

The Transducer

When controlling a pointing device, a user controls some parameter, which is sensed by the transducer and translated into another parameter for the software driver. These parameters might be the same, as in the case of a touch sensitive tablet, when the user moves a stylus, thereby changing its position. This is illustrated in Figure 12.

Figure 12. Block diagram of a touch pad.

In other devices, the mapping between the input parameter and the intermediate parameter is so direct. For example, in a mechanical mouse, the user changes the physical location of the mouse, and this is sensed by rollers moved by the mouse ball. This is illustrated in Figure 13.

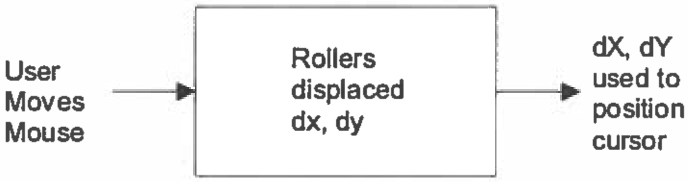

Figure 13. Block diagram of a mouse indicating how the mouse's movement turns rollers.

Because there is no direct mapping of user input (movement) into the intermediate parameter for the software driver, users can pick up the mouse and put it down in another position without moving the cursor, i.e., users can move the mouse without affecting its sensor. If the mapping were direct, i.e., if the mouse could directly measure displacement, then repositioning the mouse would not work. There are in fact 3D mice, such as one by Apple Computer Inc., which sense the position of the mouse in 3-D space, where such a trick cannot be used.

The mapping gets even more indirect in the case of isometric joysticks (Douglas & Mithal, 1993), as illustrated in Figure 14. In this device, users vary the amount of force they applied on a joystick, thereby changing the resistance of varistors attached to the joystick. A similar device can be built using strain gauges in place of the varistors. The change in resistance is used to control the velocity of the cursor.

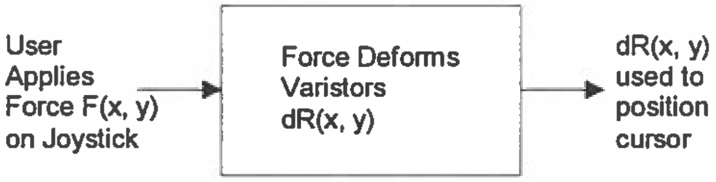

Figure 14. Block diagram of the Home Row joystick.

Thus a number of parameters in the pointing device are important when considering psychomotor models applicable to the device. For example, it might require different strategies to control the position of the cursor with an isometric joystick and a mouse.

Software Driver Transfer function (D), Cursor Parameter (E)

Once the device feeds its signals into the computer, other factors come into play. One such factor is the (cursor display) output parameter that the software driver controls. For instance, in the mouse, a displacement in the mouse is mapped onto a displacement in the cursor. On the other hand, in most isometric joysticks, it is the *velocity* of the cursor that is controlled. Based on the output parameters, we can call devices (absolute) position control, displacement (relative position) control, velocity control and acceleration control (and any arbitrary time derivative of position that we choose to take).

It is important to note that the sensed input parameter can be mapped onto any of these output parameters simply by changing the software driver. Take the example of the mouse, which can be used to control displacement, velocity or acceleration. We could re-write the mouse's device driver so that its displacement from a central location controlled the velocity of movement of the cursor. For example, we could have an algorithm that mapped one mickey (smallest unit of mouse movement) away from the central location, to a movement of one pixel per second away from the center. So, if (0, 0) were at the center of the screen, and the mouse were at (2, -3) mickeys, it would have a velocity of (2, -3) pixels per second away from (0, 0).

Similarly, a joystick, is most often configured so that it controls the velocity of the cursor, but it could be configured so that it controls the absolute position of the cursor away from the center. Jagacinski studied such a joystick (Jagacinski et al., 1980b).

Pointing movements are sometimes referred to as target acquisition tasks because their goal is to acquire a target. Jagacinski's studies on target acquisition with moving targets suggested that there is an issue of stimulus response compatibility for pointing devices, so that position control devices would be better for stationary targets, and velocity control devices for targets moving at a constant speed (which makes the target acquisition problem a tracking problem). By extrapolation, it might be reasonable to suggest that acceleration control devices might be best for capturing accelerating targets.

For example, automobile accelerators can be roughly called velocity control devices: the pressure on the accelerator pedal is converted into a forward velocity in the car. The harder the driver pushes on the pedal, the faster the car travels. This is stimulus-response compatible with driving, when all the cars are moving, and drivers attempt to move into spaces between cars that are themselves moving. However, when parking a car, the driver tries to put a car into a stationary target, and at such a time, a position control device might make it easier to park. Similarly, velocity controlled devices might be more stimulus response compatible with video games and for air traffic controllers.

The Overall Transfer Function

The relationship between the input and output parameters has been used by some researchers to categorize some kinds of input devices. For example, a mouse maps a displacement (of the mouse on the mouse pad) to a displacement (of the mouse on to the screen). This is referred to as a zero order device. An isotonic joystick that maps joystick displacement to cursor velocity maps a displacement to a velocity and is referred to as a first order device. If it mapped the displacement to an acceleration, it would be called a second order device.

This categorization is somewhat simplistic. For example, it is not clear that a force-joystick that controls velocity is a first order device. Instead, it is better to say that the relationship between the input parameter and the output parameter affects device performance without qualifying that statement.

Isometric Versus Isotonic Devices

A factor that does not fit into the block diagram of a pointing device shown in Figure 11, but that I believe is important, is whether or not the device is isometric. The difference between an isometric and an isotonic device is that the isotonic device provides users with two kinds of feedback, visual and kinesthetic, while an isometric device does not provide kinesthetic feedback. One of the aspects of the experiment described in Chapter III is to determine if this factor makes a difference in the way in which users control a pointing device.

Other Pragmatic Considerations

There are a few other pragmatic aspects of input devices that must be considered, but which are more difficult to categorize. For example, a non-isometric joystick can be spring loaded or not. If it is spring loaded, it will always return to a central rest position when released. When it is not spring loaded, it will remain in the last position it was moved to. Such factors can in turn affect other factors. For example, the spring loaded device will more likely be used as a velocity control device, while the other will more likely be used as a position control device. In addition, the spring loaded device could sense either force or position as the input parameter.

Such pragmatic considerations can sometimes be carried across to other devices. For example, a spring loaded mouse is conceivable, which always returns to a central location when released. Note that this return to the central location can either be performed explicitly by springs, or tacitly by software, which could return the cursor to the center of the screen after a few seconds of inactivity.

Designing and Modeling Pointing Devices

The pragmatics of pointing devices shows that there are wide variations in the kinds of pointing devices available, and the literature shows that a large number of these devices, used in a variety of different tasks, all conform to Fitts' law. This is a very powerful result. It almost appears that whatever we do to the task, and however we mediate the manner in which the output variable is controlled, humans manage to make everything the same, i.e. their movement patterns match Fitts' law.

Which gives rise to a number of questions. Is the problem of controlling a pointing device a problem of cognition? Is there some kind of control mechanism in the brain that makes everything the same? The problem is that we cannot answer these questions with today's level of knowledge, and the goal of this dissertation is to add to our base of knowledge, so that these questions can ultimately be answered.

Chapter I indicated that all the work on designing pointing devices has been done by HCI researchers. In this sub-section, examines some of this research.

Absolute devices such as touch tablets, and displacement devices such as mice and trackballs are somewhat uninteresting from the point of view of design, because they have a very simple overall transfer function. Design efforts for these devices have focused on improving the number of mickeys per inch that mice, trackballs and touch tablets can sense. Some additional efforts have been made to improve pointing speed through the use of accelerated drivers for mice and trackballs. Accelerated drivers have a non-linear mapping between input displacement and output displacement. The faster the mouse moves, the greater the gain. A study of such *power mice* concluded that accelerated drivers had no effect on pointing speed, and that any subjective gains were probably as a result of the reduced amount of desk space that the mice required to move on (Jellinek & Card, 1990).

Another aspect of pointing devices that has been studied is the control-display gain that is most appropriate for device control, and research has found that there is a range within which the device is usable, and that within that range, the control-display gain does not appear to have much effect (Kantowitz & Elvers, 1988; Lin et al., 1992; MacKenzie, 1991).

The most interesting work on pointing device design has been done in the area of isometric velocity control joysticks (Bentley et al., 1975; Douglas & Mithal, 1993; Rutledge & Selker, 1990).

Bentley et al. designed the joystick used by Card et al. (1978). They used the following transfer function to determine the position of the cursor on the screen based on the force being applied to the joystick:

$$Cx_k = a_1Cx_{k-1}+a_2Cx_{k-2}+a_3Cx_{k-3}+$$
$$b_0\,(Vx_k + b_1Vx_{k-1} + b_2Vx_{k-2} + b_3Vx_{k-3}) \qquad (21)$$

Where

Cx, Cy = cursor coordinates

Vx, Vy = transducer voltages

$k-3$, $k-2$, $k-1$ = discrete sequential sampling times

a_i = constants

There are two design problems with this approach. First, there are many different forms that the transfer function can take, and designers have to pick one. Second, having picked a transfer function, designers do not have a good means of picking the values for the constants in the transfer function, such as $k-3$, $k-2$, $k-1$ and a_i in this case.

As a result, all these design efforts are marked by a cycle of picking a transfer function, setting the control parameters, testing the device with users, changing the parameters and re-testing, and going through this cycle until acceptable performance is achieved. Such an approach has the problem that any solution might not be optimal. In addition, manufacturers claims of "improved drivers" are difficult to substantiate because few if any perform any experiments to show that their drivers improve pointing speed, and therefore it is not clear that any changes they made to their drivers are in fact "better."

Thus pointing device design has not made use of models of pointing, which is not surprising to some extent, because there has been very little work on modeling pointing devices. A recent example is a study by that examined the validity of Meyer's Stochastic Optimized Submovements Model for mouse movement, and concluded that model accounted for some, but not all of the characteristics of mouse movement (Walker et al., 1993).

Summary

To summarize, there are a lot of similarities between manual pointing and computer pointing, and the two are tied together by Fitts' law. In fact, the two are so similar, that it is difficult to differentiate between research done by psychologists, and that done by HCI researchers.

While a single model might be able to explain manual pointing with different limbs, the same might not be true for pointing devices, which vary widely. This implies that we cannot simply take models developed by psychologists and apply them to pointing devices, they need to be tested on different devices.

On the other hand, the fact that Fitts' law holds for such a wide range of limbs, tasks and devices, suggests that there is a single underlying mechanism that explains human movement. If this is so, we need to capture that commonalty. This knowledge will come in part by modeling computer pointing behavior, where enough work has not been done.

Once we are able to model pointing behavior on a computer, we will be in position where we will be able to make better designs of computer pointing devices.

These issues are addressed in the experiment described in the next chapter.

CHAPTER III

METHOD

This chapter describes the experiment that was designed to answer the research questions. The questions fell into three categories, about the applicability of Fitts' law to the two devices, about the microstructure of movement for the devices, and about applying the SOS model to them.

There were two main guiding principles in the design of the experiment. It was designed to be a simple pointing task that allowed easy comparison of differences between the two pointing devices, the mouse and the isometric joystick, as well as comparisons between subjects and changes over time for individual subjects. It was also designed to give subjects enough practice so that the data from the end of the study would represent practiced performance.

Subjects

Six volunteers participated in the experiment. All were students or employees at the University of Oregon. Four subjects were from the Department of Computer Science, and two were from the College of Business. One each of the subjects from each of these groups (a total of two) were women, and there were four male subjects.

Of the four subjects from the Department of Computer Science, all were expert mouse users, and used the mouse on various computers throughout a normal day. Two of these subjects were graduate students, and the other two were employees of the Computer Science Department. None of the subjects had any experience with the TrackPoint or other isometric joysticks apart from trying out various isometric joysticks that have been available at our labs from time to time. One of the subjects, a graduate student, had played a lot of video games using a joystick.

Of the two subjects from the College of Business, one was a graduating senior, and one was a doctoral student. Both of these subjects owned IBM ThinkPads, which have embedded TrackPoint joysticks. One subject also had extensive experience with a mouse, while one had only minimal experience with a mouse, but extensive experience with the TrackPoint.

All subjects were right handed.

Subjects were promised a home cooked Indian meal for participating in the experiment.

Experimental Equipment

The hardware used for the experiment was a Gateway 2000 P5-90 XL computer with an Intel Pentium processor running at 90 Mhz. Tests showed that the processor did not have the floating point division error identified with earlier versions of the Pentium. The computer had 16 MB of memory and 1 GB of hard disk. The display adapter was an

ATI Mach 64 accelerated video card with 2MB of video ram, displaying 65,536 colors at a resolution of 1024x768 pixels. This was attached to a 17" NEC MultiSync XE17 monitor. The monitor displayed 1024 pixels in 310 mm. The operating system was Microsoft Windows NT version 3.5.

The mouse used was a PS/2 compatible mouse from Microsoft, FCC ID C3K5400PS2, Part # 31660. The isometric joystick was a TrackPoint II keyboard from IBM. This keyboard, represented in Figure 15, has an isometric joystick embedded between the 'G', 'H' and 'B' keys of a standard keyboard. The joystick being isometric does not move, but senses the magnitude and direction of the force applied to it, and moves the cursor with a velocity proportional to the force, and in the direction that the force is applied. The mapping of force to velocity is non-linear as described in (Rutledge & Selker, 1990).

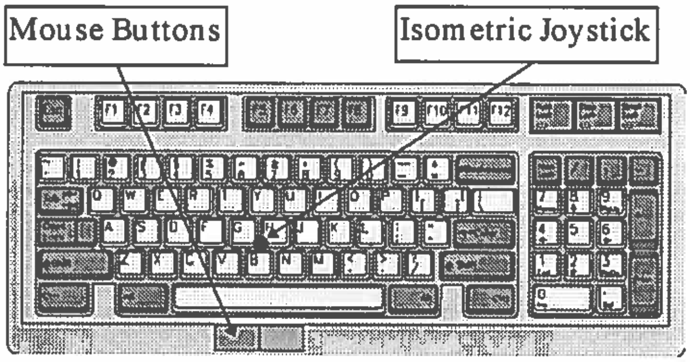

Figure 15. The IBM TrackPoint II keyboard.

The buttons below the space bar correspond to the buttons on a standard two-button mouse, the left button mapping on to the left button on the mouse, and the right button mapping onto the right button on the mouse. In the standard touch typing position, the joystick is equidistant from the index fingers of both hands, and the mouse buttons fall under the thumbs. Some subjects chose to click the left mouse button with the thumb of the right hand, while others used the index finger of the left hand (the right button was not used for this experiment). They were instructed to use whatever technique they preferred.

The manner in which the mouse and TrackPoint keyboard are connected to a computer is a little unusual, and it is explained below. Understanding these connections are critical for interpreting the results that are presented in later chapters.

A pointing device can be connected to an MS-DOS computer in one of two ways, either through a serial (RS-232C) port, or through a PS/2 mouse port (i8042), as shown in Figure 16.

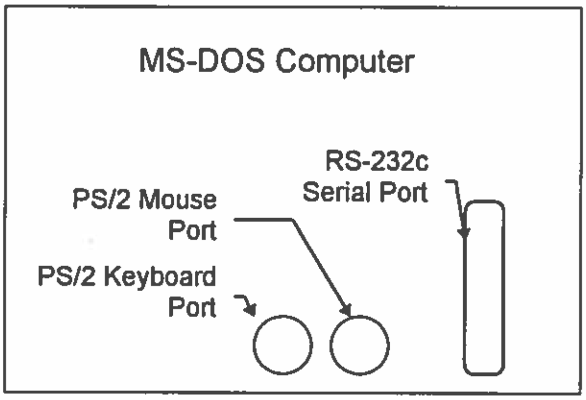

Figure 16. Input device connections for an MS-DOS computer.

The TrackPoint keyboard has two cables that come out of it (Figure 17), one going to the PS/2 keyboard port, and the other going to the PS/2 mouse port. In addition, there is a PS/2 mouse port on the keyboard that allows a user to attach a mouse to the keyboard for use in addition to the TrackPoint joystick.

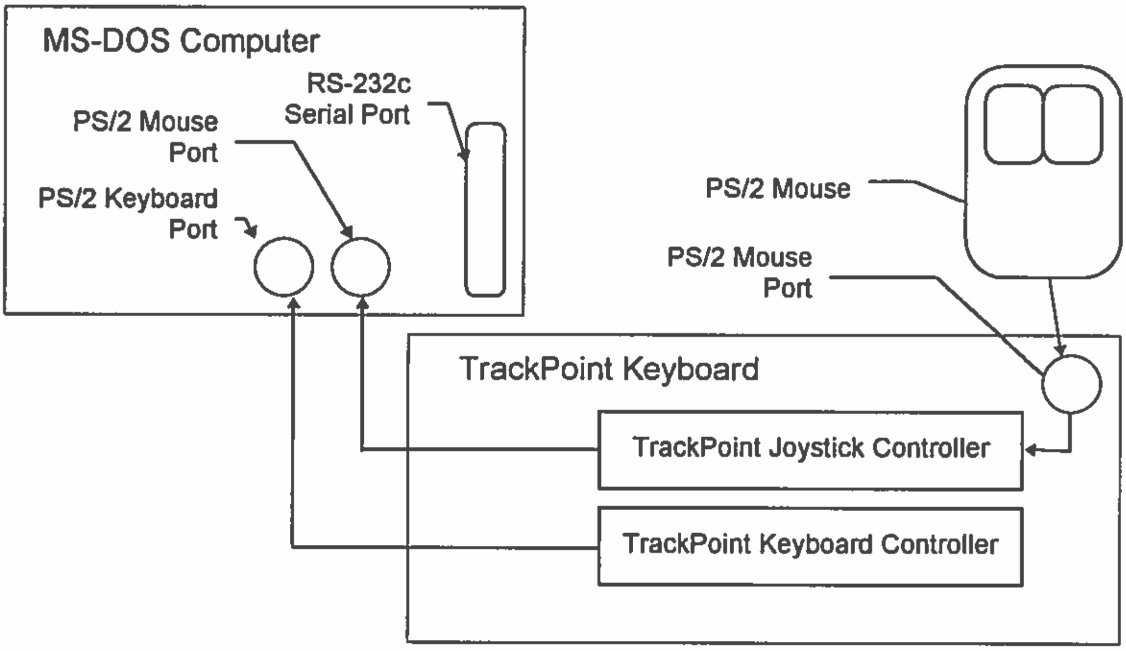

Figure 17. Connections for the TrackPoint keyboard.

The connections from this mouse goes through the TrackPoint joystick controller and feeds signals to the PS/2 mouse port on the CPU through the same set of wires that the TrackPoint's own signals are fed to the computer. When a user uses the TrackPoint joystick, there is a pause of a few seconds before the mouse becomes active, i.e., there is a gap of a few seconds before mouse signals are fed to the computer. The same PS/2 software driver running inside the computer responds to these signals, and these are picked up by the same program presenting the experimental task.

A PS/2 mouse can, of course, be connected directly to the CPU in conjunction with at PS/2 keyboard, and the PS/2 mouse can be substituted with a serial mouse. For the main study, the configuration shown in Figure 17 was used. Tests were conducted with pointing devices in all configurations, serial mouse, PS/2 mouse with standard keyboard, PS/2 mouse with TrackPoint keyboard, and TrackPoint keyboard with and without PS/2 mouse.

The monitor, keyboard and mouse were placed on an ergonomic workstation table with the monitor's center at eye level, and the keyboard and mouse at a position where they could be accessed by subjects while keeping their arms straight. The monitor was 33" (82 cm) away from the subject's eyes. Subjects were given a height adjustable chair and a footrest and were asked to adjust the experimental equipment to maximize their comfort.

Experimental Procedure

During the first session, subjects were asked to read, sign, and keep a copy of the informed consent form which described their rights as subjects in the experiment. They were then described the goal of the experiment, i.e. to gather data about cursor movement as they completed the experimental task. The operation of the pointing devices was also explained to them. In order to level the playing field, this explanation was given to them even if they were already familiar with the operation of the devices. Finally they were shown the controls on the chair and were asked to make any adjustments they felt necessary to the height and angle of the chair and to the lighting of the room.

Subjects were then asked to perform 10 warm-up trials with the experiment. These warm-up trials were repeated at the start of every session. They were discarded from the data analysis. During the warm up trials subjects were also shown the computer's behavior when an error occurred, and what they had to do to complete the trial.

The program provided a score as motivation. Subjects were shown the score bar at the bottom of the screen, and were encouraged to get as high a score as possible. This could be done by hitting the target quickly and keeping errors low. In order to increase motivation the scores were recorded and they were told how their scores varied across trials.

They were told that they could stop for a break at anytime, and that if they stopped between trials, they would not affect the data collection process. At the end of every block of 120 trials, a dialog box appeared telling the subjects that a block was completed. They were told that this was also an appropriate place to pause. Some subjects rested when they felt the need, while others preferred to get through the session as quickly as possible.

Subjects did not seem to have any trouble with these instructions. After the 10 warm up trials, the system was reset, and the experiment started. Sessions averaged 25 minutes.

Experimental Task

The task used for the experiment was based on the experimental configuration used in an earlier study of the applicability of the SOS model to mouse movement (Walker et al., 1993).

The experimental task presented to the users appeared as depicted in Figure 18. It presents a starting location, the home square and an ending location, the target ribbon. These elements were presented against a white background. 'A' represents the distance from the center of the home square to the center of the target ribbon. 'W' represents the width of the target ribbon. The values taken on by these are explained below.

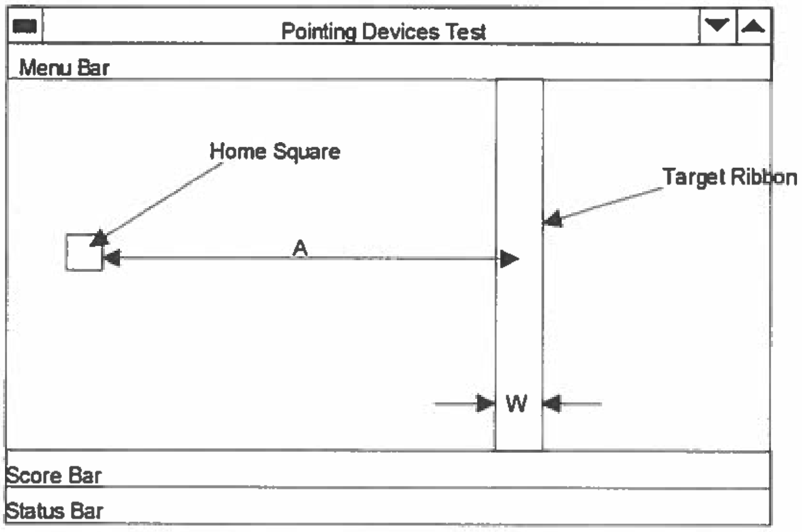

Figure 18. The experimental task

Subjects were asked to first click in the home square, and then in the target ribbon. A trial began when the subject clicked in the home square, and ended when the subject clicked in the target ribbon. At the start of the trial, the home square was colored red and the target ribbon white. When they clicked in the home square to start a trial, the home square was turned white and the target ribbon was turned red indicating both where they should aim, and that the trial was in progress. At the start of every trial, the software repositioned the cursor in the center of the home square. This was done to prevent the subjects from performing this movement on their own, and thereby performing a movement task that was not part of the experiment.

The time at the start of the trial and the time at the end of the trial were recorded. In addition, the software sampled the position of the mouse cursor between the clicks, keeping track of position and time, generating a sample on the average of every 16 ms.

After the completion of a trial, the screen was blanked, and the data from the trial was written to disk and then the next trial was presented. The cursor was re-centered on the home square prior to each trial, and the home square always appeared in the same location. The width and location of the target ribbon were varied as indicated in Table 3. These values generated the following six amplitude to width ratios (A/W): 48, 24, 12, 6, 3, 1.5.

Table 3. Values for amplitude (distance to target) and width (of target) used in the experiment.

Amplitude (A)		*Width (W)*	
Pixels	mm	Pixels	mm
768	232.5	16	4.84
384	116.3	32	9.69
192	58.1	64	19.38
96	29.1		

The software grouped 10 sets of the 12 combinations of width and distance together, randomized the order, and presented the 120 trials in this random order as a single block. 10 sets of trials were collected and randomized as opposed to randomizing and presenting a single set of 12 trials out of concern that if only one set was randomized, the subjects would learn the order of presentation. The order of presentation was exactly repeated from block to block, except for error trials. Subjects did 5 sessions per block, and three sessions per device.

When a subject made an error, i.e., clicked outside the target ribbon, the computer beeped, indicating the error. The subject then had to complete the trial by clicking in the target ribbon. The computer recorded the incorrect clicks and marked the trial as an error trial. The next trial in the sequence was then presented. The error trial was repeated at the end of the block, so that subjects had to complete all trials correctly.

Two status bars presented a score and a count of the total number of correct trials, and the number of blocks completed for the session.

Design

The experiment was a single factor (device) within-subjects design. Subjects were randomly assigned to one of the subject slots shown in Table 4. Each subject performed three sessions with one pointing device, followed by three sessions with the other device. The order was balanced among the subjects to remove order effects.

Table 4. Order in which subjects did the experiment.

Subject	Device used during session					
1	Mouse	Mouse	Mouse	Joystick	Joystick	Joystick
2	Joystick	Joystick	Joystick	Mouse	Mouse	Mouse
3	Mouse	Mouse	Mouse	Joystick	Joystick	Joystick
4	Joystick	Joystick	Joystick	Mouse	Mouse	Mouse
5	Mouse	Mouse	Mouse	Joystick	Joystick	Joystick
6	Joystick	Joystick	Joystick	Mouse	Mouse	Mouse

A session consisted of 5 blocks of 120 trials, or 600 trials per session. Subjects did three sessions per device, or 1800 trials per device. Earlier work (Card et al., 1978; Douglas & Mithal, 1993), has shown that in experiments with far more variability in the task presented (e.g. 192 different trial combinations in the Douglas study as opposed to 12 in this study), subjects showed minimal learning effects after about 1200 trials. It is therefore reasonable to assume that the data from the last block of the last session represents expert level performance even though most of the subjects were expert mouse users, and had little experience with the joystick prior to the study.

Sessions lasted approximately 25 minutes, and were scheduled so that a single subject did not do more than one session a day. On the average, subjects did sessions every alternate day, as made possible by their schedules. The experiment was conducted during the first three weeks of February, 1995.

Summary

Earlier research on pointing devices has not studied the characteristics of movement as a subject points towards a target on screen. This information is necessary in order to model pointing devices, and in order to explain performance differences among them. This experiment was designed to gather data on these characteristics.

It is important to note that this study measures practiced performance, and looks at the microstructure of movement after a large number of trials. It is not designed to make performance comparisons between the mouse and the isometric joystick. Such comparisons have been extensively made in the past, showing that the mouse is substantially faster than the isometric joystick.

CHAPTER IV

RESULTS

The experiment described in Chapter III was conducted in two stages. A pilot study was conducted between October and December 1994 to check the program and the equipment. The pilot study produced valuable results which had a significant effect on the subsequent research. Following this, the main experiment was conducted during the first three weeks of February, 1995. The six subjects all completed three sessions with each device, generating a total of 36 data files, which occupied 20.1 Mb of disk space.

This chapter discusses the results obtained from both the two studies. Broadly speaking, the results can be divided into two areas, depending on the kind of analysis conducted. The first kind of analysis looked at gross parameters of performance such as the total trial time, the total movement distance, and the target width. The second kind of analysis looked at the movement microstructure of individual trials.

The gross level analyses looked at trends in the overall trial completion time over the duration of the experiment. It grouped together the data from all the trials in a block. A Fitts' law analysis was conducted to determine the Indices of Performance (IPs) of the two devices, and to study the effect of movement distance and target width on trail completion time.

In order to understand why there are performance differences between the two devices, a microstructure level analysis was performed. This revealed differences between the two devices, as well as differences in performance with practice, and individual differences between the subjects which were linked to their performance with the devices.

We begin with the results from the gross level analysis. This is followed by the results from the microstructure level analysis, which revealed tremor in the isometric joystick and individual differences between the subject, which are discussed next. The last two sections discuss the applicability of the SOS model to the joystick's microstructure of movement and an analysis of tremor in the isometric joystick.

Gross Level Analysis

While the primary goal of this study was to perform a microstructure analysis of movement, an analysis of the gross movement characteristics was performed to ensure that there was nothing unusual about the data.

Table 5 presents the mean and standard deviations for each device, for each of the blocks of 120 trials, averaged over all subjects. As expected, the trial times drop with practice. The mouse also has quicker trial times than the joystick. Figure 19 plots this data.

Table 5. Mean and standard deviations of trial blocks, n=6.

Joystick Trials

Block	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Mean	1324	1216	1178	1116	1077	1072	1028	1009	984	978	1047	992	970	966	946
Std Dev	180	191	236	200	188	174	152	153	158	157	212	165	189	156	143

Mouse Trials

Block	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Mean	990	927	898	877	856	866	912	933	912	904	937	888	917	874	860
Std Dev	152	147	141	108	69	99	224	248	249	231	227	215	243	220	187

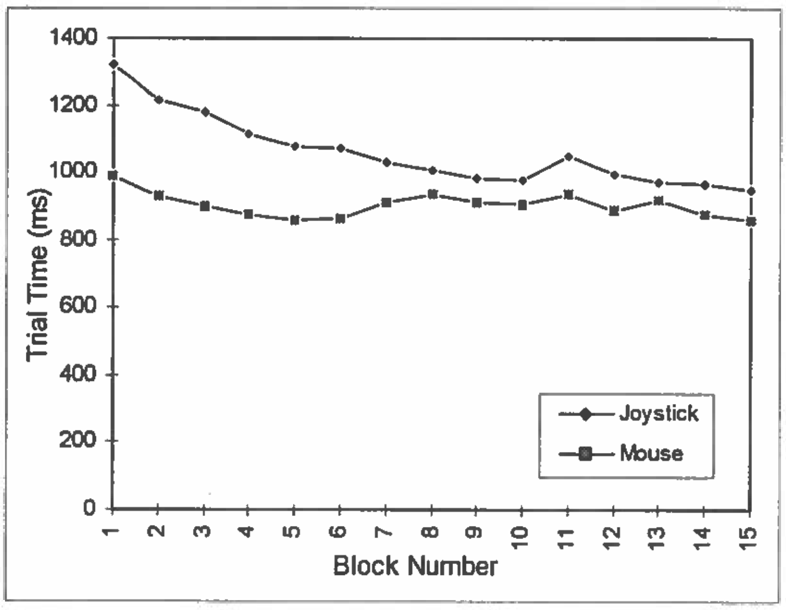

Figure 19. Mean trial times by block across subjects. n = 6.

Figure 20 and Figure 21 present the data from Table 5 in more detail. Figure 20 plots the average block time by block for each subject for the joystick. Subjects DD5 and

DD6 had prior joystick experience, though DD6 did not have one of the two fastest joystick times. DD6 was in the upper median, and shows a practice effect whereas DD5 does not. Of the two DD6 had much more experience (two years) than DD5 (six months), so these results are somewhat contrary to expectations.

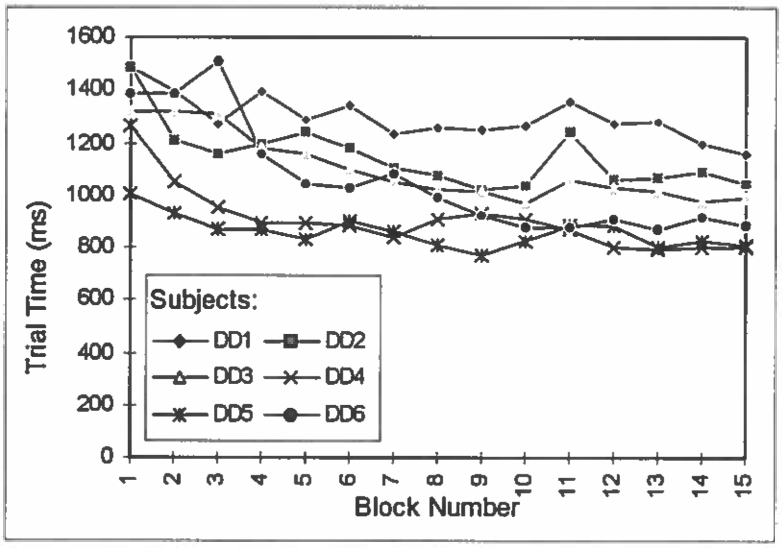

Figure 20. Mean block times for joystick by subject.

Figure 21 shows the average block times for each subject with the mouse. Note here that subjects DD1 through DD5 were all experts with the mouse prior to the start of the experiment and show little practice effect. Subject DD6 had very little mouse experience, and throughout the study expressed a dislike of the mouse. The sharp increase in average trial time for this subject from block 6 to block 7 is difficult to understand. That session begins with block 6, and at block 6 the subject's performance was not out of the ordinary. It is possible that the subject changed his goals (perhaps to minimize errors rather than to optimize between minimizing errors and maximizing speed), which accounts for the increase in trial time. An attempt was made to bring the subject in for one additional mouse session, but as the subject's thesis defense date was nearing, this was not possible.

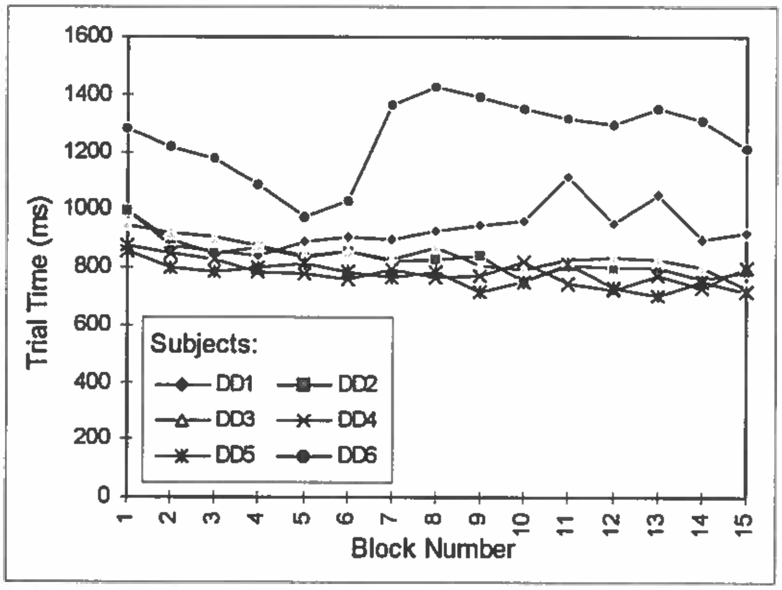

Figure 21. Mean block times for mouse by subject.

Due to the strange nature of this subject's mouse trials, this subject's data was dropped from the analysis that follows. The trial means and standard deviations for each block without this subject's data are shown in Table 6 and this data is plotted in Figure 22. The effect of dropping this data can be seen by comparing Figure 22 to Figure 19. The result is that there is a wider separation between the mean trial times for the two devices.

Table 6. Mean and standard deviations of trial blocks without Subject DD6, n=5.

Joystick Trials															
Block	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Mean	1312	1181	1112	1106	1084	1081	1017	1013	996	998	1080	1009	991	976	959
Std Dev	199	191	193	223	209	193	167	170	173	167	218	179	204	172	156
Mouse Trials															
Block	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Mean	933	869	843	835	832	833	822	835	815	815	861	807	830	788	790
Std Dev	63.5	45	44	41	41	59	48	66	85	85	145	94	132	67	81

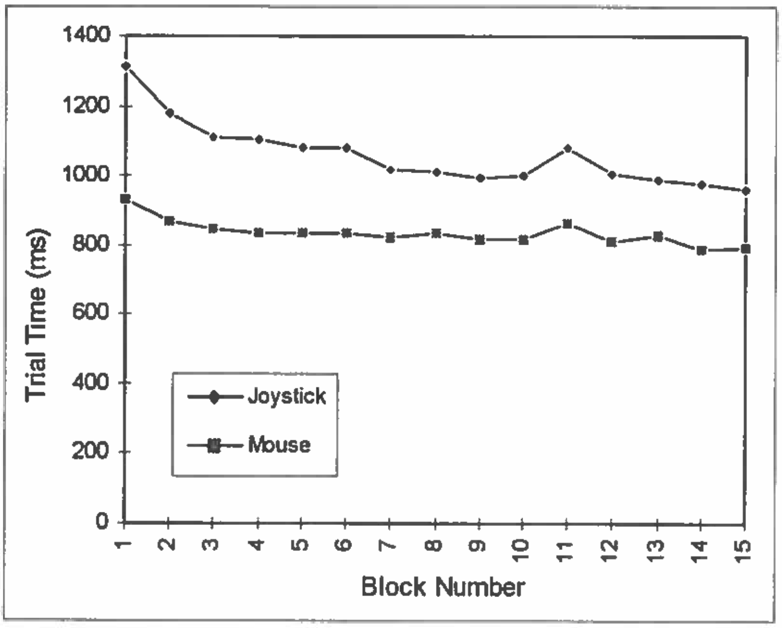

Figure 22. Mean trial times by block without Subject DD6. n = 5.

One of the issues that played a role in the design of the experiment was that volunteers were likely to have a lot of experience with the mouse and very little experience with the joystick. Therefore it was important to give the subjects enough practice so that they reached a point in the learning curve where there would not be significant improvement from block to block with each device, particularly with the joystick. Earlier studies (Card et al., 1978; Douglas & Mithal, 1993) have shown that after approximately 1200 trials, the learning from block to block becomes statistically insignificant ($p \geq .05$), so the subjects were given 1800 trials to ensure that they reached this stage with each device.

In order to determine if there was statistically significant learning from block to block, a pair-wise *T* test was performed between successive blocks. These results are listed in Table 7.

Table 7. Pair-wise *T* test between successive blocks, n=5.

Block	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12	12-13	13-14	14-15
Joystick	0.03	0.01	0.45	0.21	0.45	0.00	0.42	0.13	0.45	0.06	0.06	0.17	0.26	0.12
Mouse	0.01	0.00	0.28	0.42	0.46	0.22	0.18	0.17	0.50	0.14	0.08	0.19	0.13	0.47

At $p \leq .05$ we accept the hypothesis that the two distributions from Block(n) and Block(n-1) are from different populations. When the p-value is higher than .05, we reject the hypothesis and can say that there was no significant learning between the two blocks.

For the mouse, the only two blocks which showed significant differences were blocks 1 through 3, after which there was no significant learning. For the joystick, there was no significant learning after block 7. Thus we can assume that by the end of the experiment subjects were sufficiently practiced on both devices.

Fitts' Law Analysis

The Fitts' law equation indicates that the distance to the target (A) and the target width (W) should have strong effects on trial time. The last block was examined for this effect and this data is plotted in Figure 23 and Figure 24. Figure 23 shows that the trial completion time increases with the log of the distance to the target, while Figure 24 shows that the trial completion time decreases with the log of the target width.

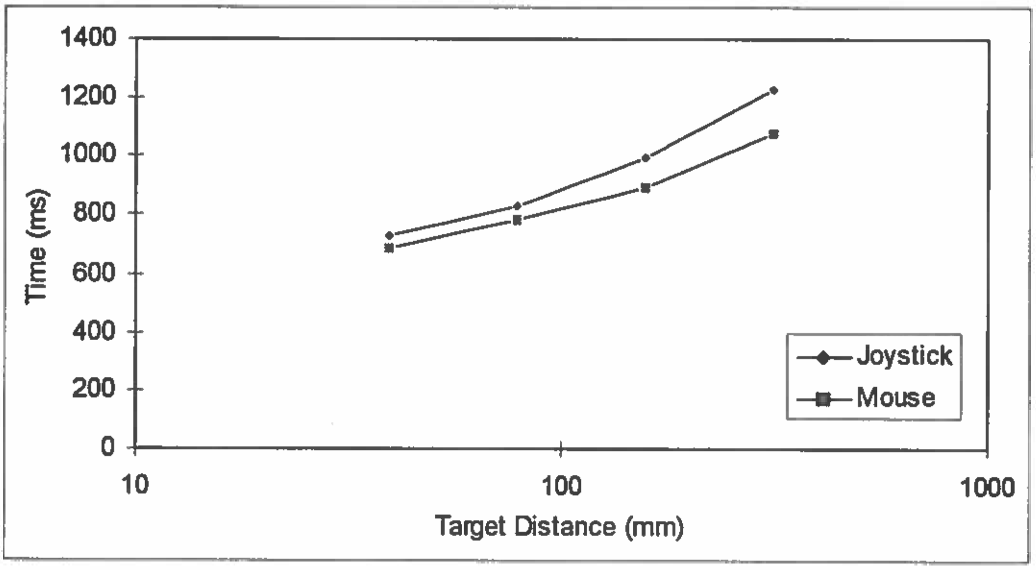

Figure 23. Effect of distance on trial time, n = 5.

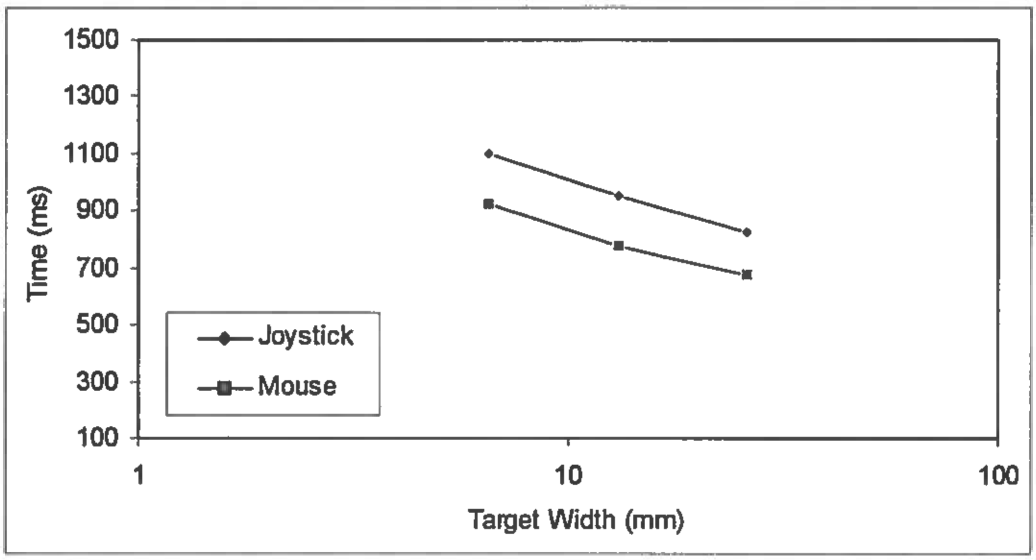

Figure 24. Effect of target width on trial time, n = 5.

Fitts' law combines the effect of distance and width by using the ID. Figure 25 plots trial completion time against the Welford formulation of ID (equation 10). We see a good match to a straight line. Performing a regression analysis of movement time against ID gives the values in Table 8. These results show that the joystick has a Fitts' IP lower than that of the mouse, which was expected.

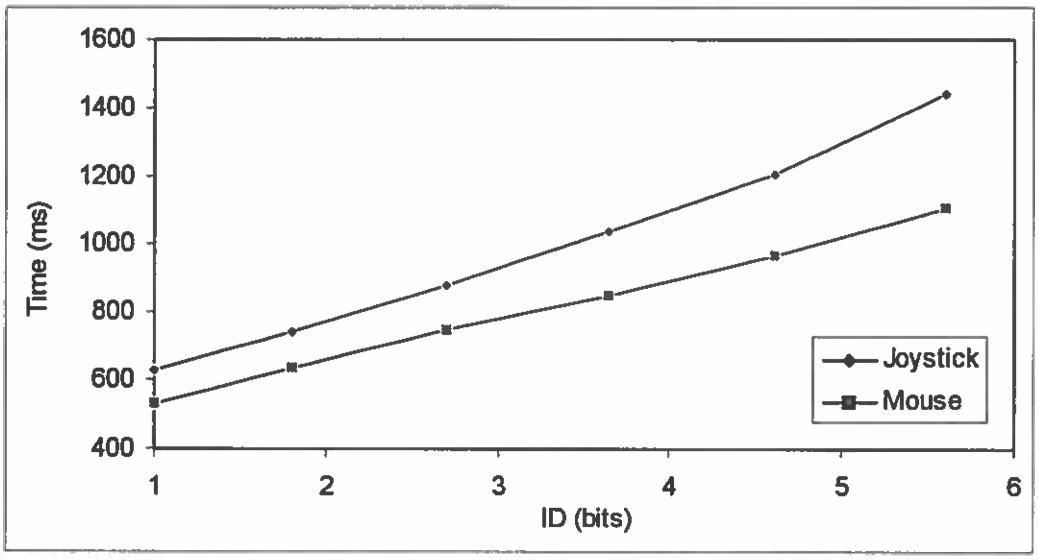

Figure 25. Plot of trial completion time with the Welford ID, n = 5.

Table 8. Results of the Fitts' law regression analysis

	r^2	IP
Mouse	0.998	8.13
Joystick	0.991	5.75
Ratio Mouse:Joystick		1.41

Generally speaking, the analyses of gross movement did not reveal any surprises. The joystick had a lower IP than the mouse, and there were effects due to target width and target distance. There was no statistically significant learning after block 4 for the mouse, and block 8 for the joystick, and thus we have confidence that the last (15th) block represents practiced performance.

Microstructure Level Analyses

The analyses of gross movement from the previous sections were conducted to develop confidence in the data so that in turn we would have confidence in the microstructure level analysis presented in this section. The microstructure level analysis was done in two stages, a pilot study and the main study.

The pilot tests were run on a 33-Mhz 486 computer built by Gateway 2000, with a Paradise 80c31 video chipset with 1 MB of video ram, running Windows for Workgroups version 3.11. When the program that presented the experiment was run on this machine it produced cursor position updates every 26 ms. As opposed to the software and equipment for the main study, which produced updates every 16 ms. When other hardware was used for the pilot study, it is described in the relevant context.

Raw Data - Velocity vs. Distance curves

This section describes the raw data obtained from the program that presented the experimental task and collected the movement data. A later section describes the data obtained by smoothing the raw data.

Mouse Movement Characteristics

We begin by looking at velocity vs. distance data from both the pilot study and the main study for the mouse.

Figure 26 is an example of data from the preliminary tests. It represents a single trial, with the time in milliseconds plotted on the x-axis, and distance from the home square plotted on the first y-axis. This curve is difficult to interpret, and is generally not used for analysis. Velocity is plotted on the second y-axis and provides more information. We can see how the velocity varied over the duration of the trial. There was almost no movement for 200 milliseconds, which corresponds well to the reaction time found in other studies (Card et al., 1978). There was then a large, rapid submovement, followed by a two smaller submovements. The first of these overshot the target, and the second was a corrective submovement that moved back to the target.

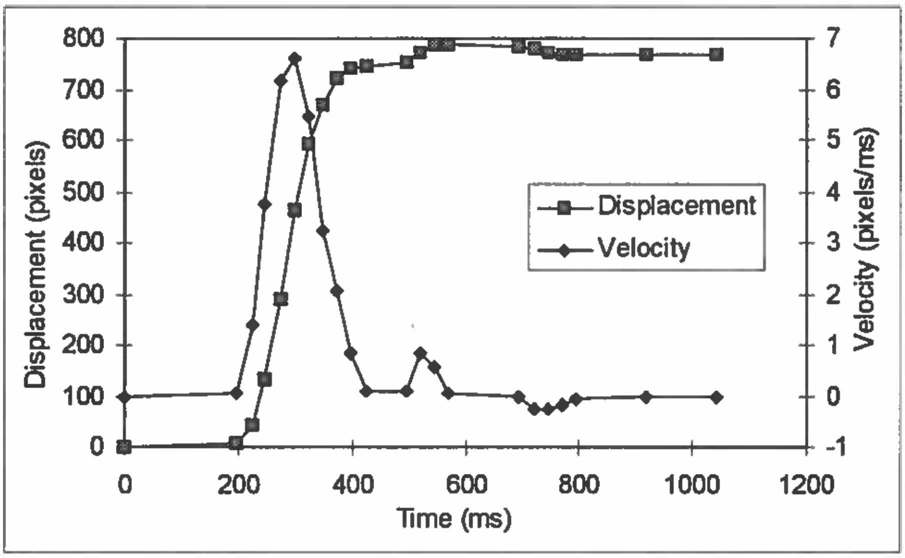

Figure 26. PS/2 mouse — trial from pilot study, A = 768 pixels.

A much better way of looking at this data is shown in Figure 27, with displacement on the x-axis, and velocity on the y-axis. From this graph, we can clearly see what the speed of the cursor was at any point along the path to the target. This graph also shows clearly that the movement is a sequence of submovements, and that in this case there are three submovements. The first submovement covers most of the distance to the target and reaches a high peak velocity. The next two submovements are smaller, covering much smaller proportions of the distance to the target. The second submovement actually overshot the target and the third submovement is a correction, bringing the cursor back to the target.

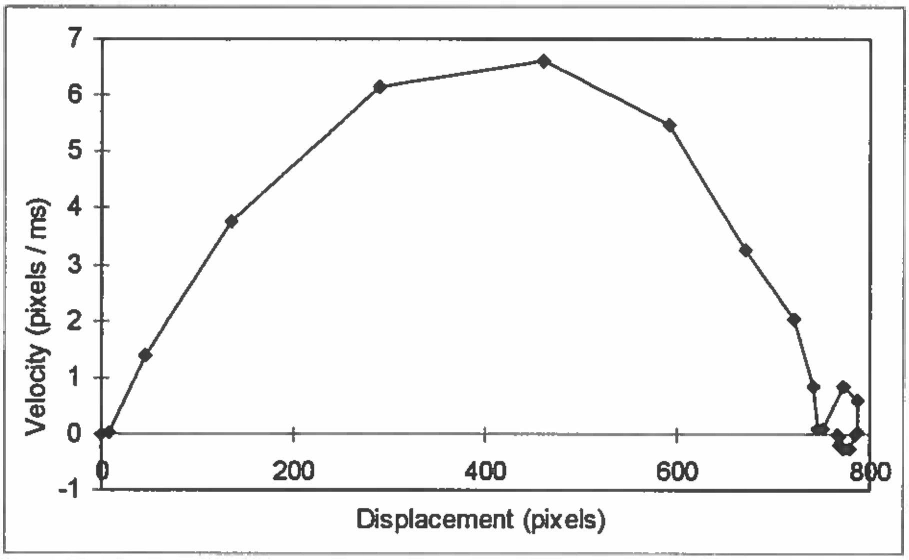

Figure 27. PS/2 mouse trial with velocity as a function of distance. A = 768 pixels.

This trial is a good match to the SOS model's predictions, and is fairly representative of mouse trials. It should be noted here that even with practiced performance, the same subject, pointing at the exact same combination of distance and width will perform each trial differently. For instance, in some trials, the subject can hit the target in one movement. In other trials, the subject will take more than one movement to reach the target. All this falls within the bounds of the SOS model, which is a stochastic model, and accommodates variations in movement.

The plot in Figure 27 is representative of the data obtained for a number of different types of mice in the pilot study. Mouse data from the main study was similar.

Data from the main study was examined in two stages. To begin with three trials for each device were examined to see what kinds of patterns emerged. These were the first trial (the first trial that the subject performed with a device), the middle trial (the middle trial from the 2nd session, i.e., the middle trial from the 8th block), and the last trial of the last session. This made a set of 18 trials for each device. Following this, the rest of the data was examined.

Figure 28, Figure 29 and Figure 30 represent data for mouse movement. The trials shown here have been taken from three different subjects. These trials have been selected for presentation because they are representative of mouse movement trials from the main study.

Figure 28 is a trial that can also be considered a prototypical representative of the SOS model. The subject made an initial rapid movement towards the target that covered most of the distance to the target. A series of movements of low velocity and duration followed that finally reached the target.

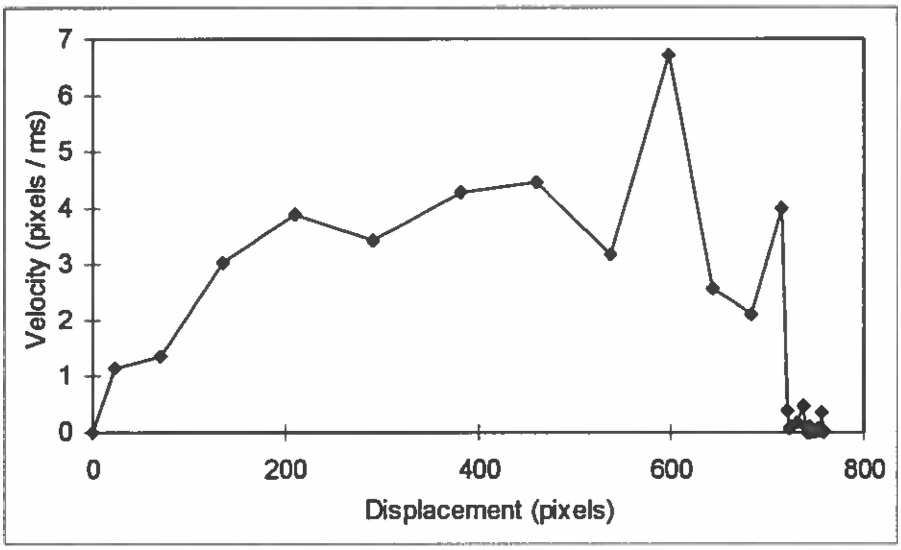

Figure 28. Last mouse trial, subject DD3 - raw data. A = 768 pixels.

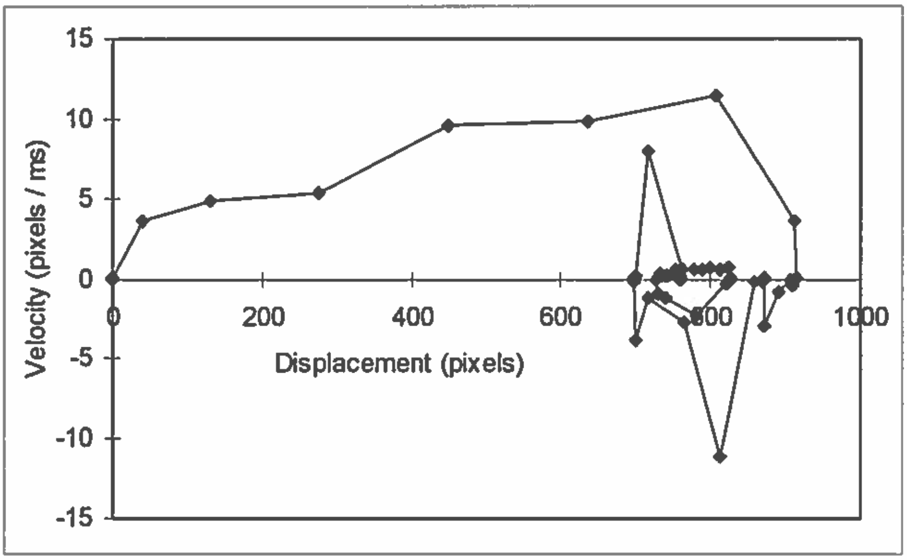

Figure 29. First mouse trial subject DD5 - raw data. A = 768 pixels

Figure 29 is interesting from the point of view of overshoots. It shows that the subject made a rapid primary submovement towards the target and overshot. The subject doubled back with the second submovement and overshot again. The third submovement also overshot, and the subject oscillated around the target till it was reached. This trial is also a prototypical representative of the SOS model in that it has a large primary submovement, followed by smaller secondary submovements. Figure 30 deviates from the SOS model prototype in that there appear to be two large movements followed by smaller movements. It serves to reinforce the point that the SOS model is stochastic in nature.

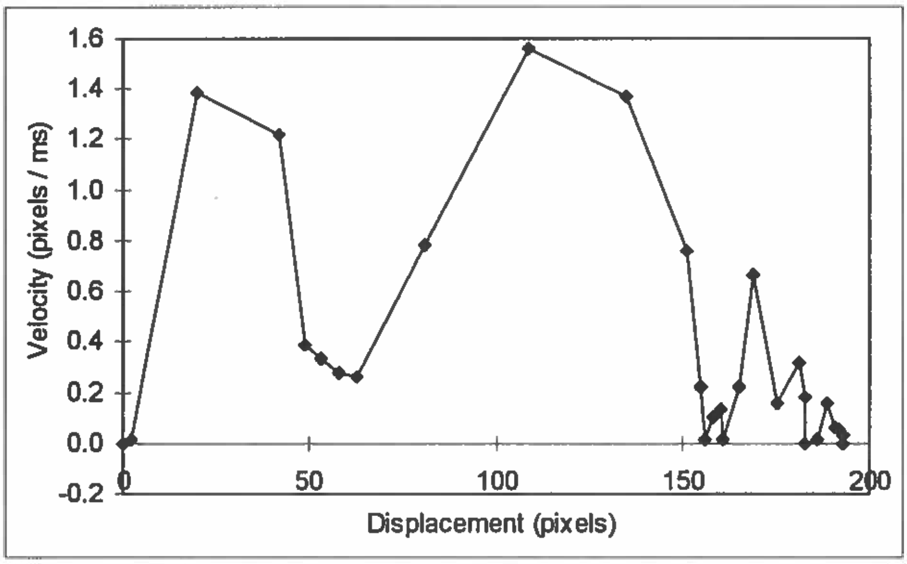

Figure 30. Middle mouse trial subject DD1 - raw data. A = 192 pixels

These figures show that in general, mouse movements have a large primary submovement, followed by smaller submovements. Note that each graph has been picked out of 3600 such graphs of mouse movement for each session. (Each of 6 subjects did 600 trials per session, and only one graph from each session has been shown.) While at every single trial has not been examined, a large percentage of the trials have, and generally speaking, for mouse movement we see a large primary submovement followed by smaller secondary submovements.

Joystick Movement Characteristics

Data for the isometric joystick was also obtained both in the pilot study as well as in the main study. The pilot study revealed an unusual characteristic in the joystick trials, and the main study also showed this phenomenon.

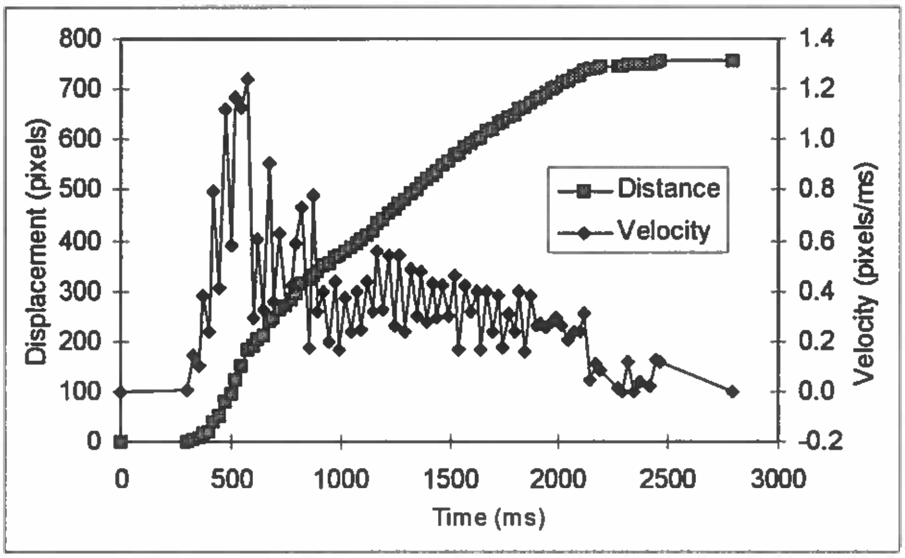

Figure 31. Isometric joystick - trial from preliminary tests.
A = 768 pixels.

Figure 31 and Figure 32 show a joystick trial from the pilot study. Note that this trial took a very long time, 3000 ms (3 seconds), and there was a significant delay (300 ms) before movement occurred. The distance vs. time sub-plot in Figure 31 does not reveal anything out of the ordinary, but the velocity vs. time sub-plot shows a strange pattern.

This pattern is more clearly evident when the velocity vs. distance plot (Figure 32) is examined. This jitter appeared in all the trials with the TrackPoint. This particular trial was one of the first attempted by the test subject (myself), but persisted after 50 trials, and also occurred in trials by others. Note that again, as in the case of the mouse, this trial is representative of multiple trials with the TrackPoint, but that individual trials are different.

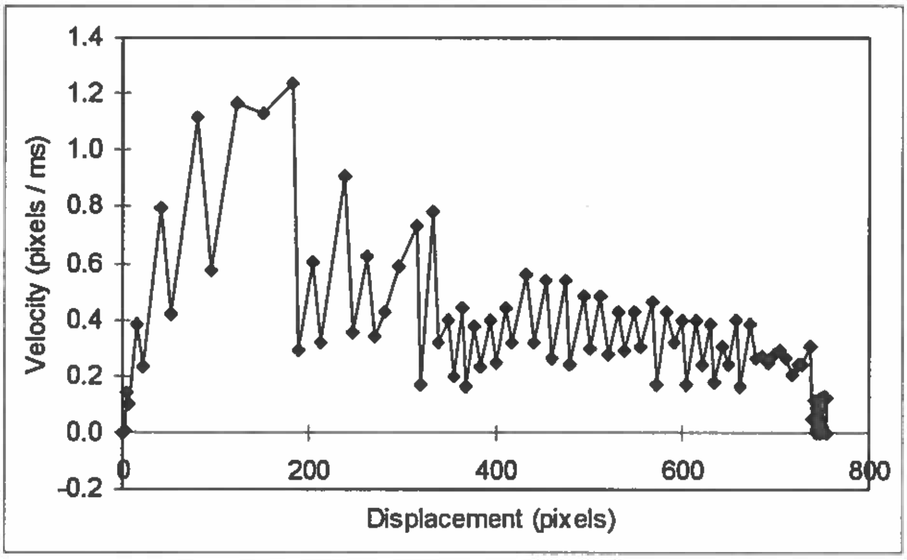

Figure 32. Isometric joystick preliminary tests. (Same trial as Figure 31.) A = 768 pixels.

There could be a number of things that could account for this jitter. Among other things, it could be caused by an interaction between the joystick and the computer, it could be an artifact of the program gathering the data, it could be the subject (myself), and it could be caused by the computer.

To eliminate the possibility that the error was caused by an interaction between the computer and the joystick, the TrackPoint keyboard was connected to another computer (a Dell P-90 Pentium based machine), with the same result, i.e., the tremor was still present in joystick trials.

Chapter III described the different ways in which pointing devices can be connected to MS-DOS computers. To eliminate the possibility that the tremor was due to the PS/2 connection used by the TrackPoint, tests were performed with four different pointing device configurations with the TrackPoint keyboard or a separate keyboard as appropriate. These configurations were the TrackPoint keyboard by itself, the PS/2 mouse by itself, the PS/2 mouse connected through the TrackPoint's PS/2 port, and an identical RS-232c serial mouse. In all cases, graphs for the mouse did not show jitter, with graphs similar to Figure 27, and jitter was present in the TrackPoint joystick.

Tests were also conducted on computers running Windows NT, which showed that the jerkiness appeared even with the faster sampling period (16 ms under NT., vs. 26 ms under Windows 3.11.). The same jerkiness was not apparent in mouse trials run on the Windows NT machines.

To eliminate the possibility that there was a flaw in this particular joystick, the software was recompiled to run on a smaller screen (640x480, the standard size for notebook computers), and the test carried out on two IBM ThinkPads (which have built-in

TrackPoints) and on an NEC Versa notebook (that has a similar isometric joystick, but is not called a TrackPoint), and the tremor was still present.

The last isometric device tested was a PortaPoint. It is a separate isometric joystick that uses the same electronics that are used in the Home Row 'J' mouse keyboard (Douglas & Mithal, 1993; Douglas & Mithal, 1994). The PortaPoint is mounted in a housing similar in shape to a mouse, and attaches to the computer through the PS/2 mouse port. Figure 33 is a trial from this device, showing jitter. The PortaPoint's isometric joystick has a 'squishy' rubber cap which deforms under the application of force, and, as we shall see in the discussion in Chapter V, this deformation might help damp the jitter.

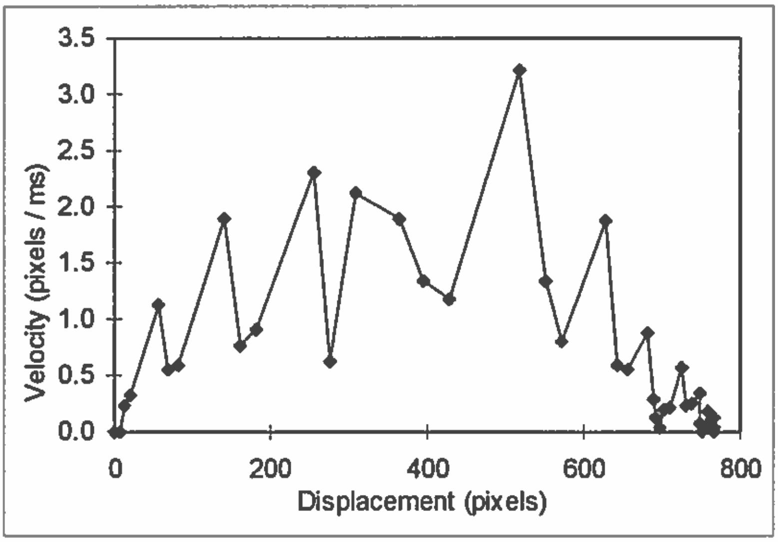

Figure 33. PortaPoint trial showing jitter. A = 768 pixels.

In all the tests conducted, none of the mouse configurations showed jitter, and all of the isometric joystick configurations showed jitter. If we go back to the base condition, the PS/2 mouse, by being connected to the TrackPoint's PS/2 connector, feeds data to the CPU through the same cable to the same software.

After carefully considering all possible sources that could cause this jerkiness, we came to the conclusion that the isometric joystick, which senses force, was picking up physiological tremor in the finger. Briefly, it is a series of random variations in force that occur in the finger, and the isometric joystick translates these changes into variations in cursor velocity. Tremor is discussed more fully in Chapter V.

The effect of tremor on the joystick's movement characteristics was a totally unexpected phenomenon, and considerable time, between October and December 1994, was spent in analyzing the source of the tremor, and conducting the tests listed above.

When the main study was conducted, tremor continued to be evident in joystick trials throughout the experiment.

The graphs shown in Figure 34, Figure 35 and Figure 36 are representative samples of joystick trials from the main study. Figure 34 shows that the movement towards the target was made up of a sequence of movements that had rapidly changing velocity. Figure 35 shows a period in the beginning of the movement which might represent a primary submovement with rapid movement towards the target, relatively unaffected by tremor. The rest of the movement is quite badly obscured by tremor.

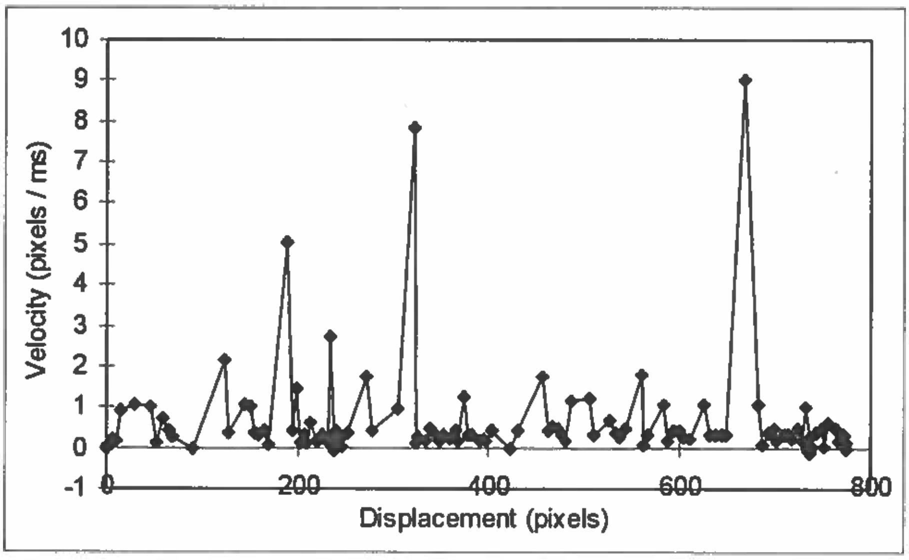

Figure 34. First joystick trial subject DD4 - raw data.

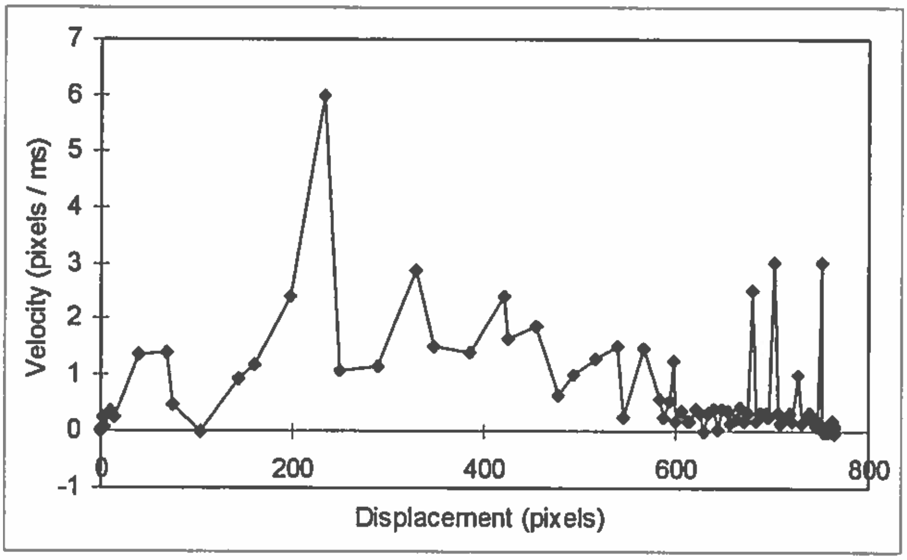

Figure 35. Middle joystick trial subject DD2 - raw data. A = 768 pixels.

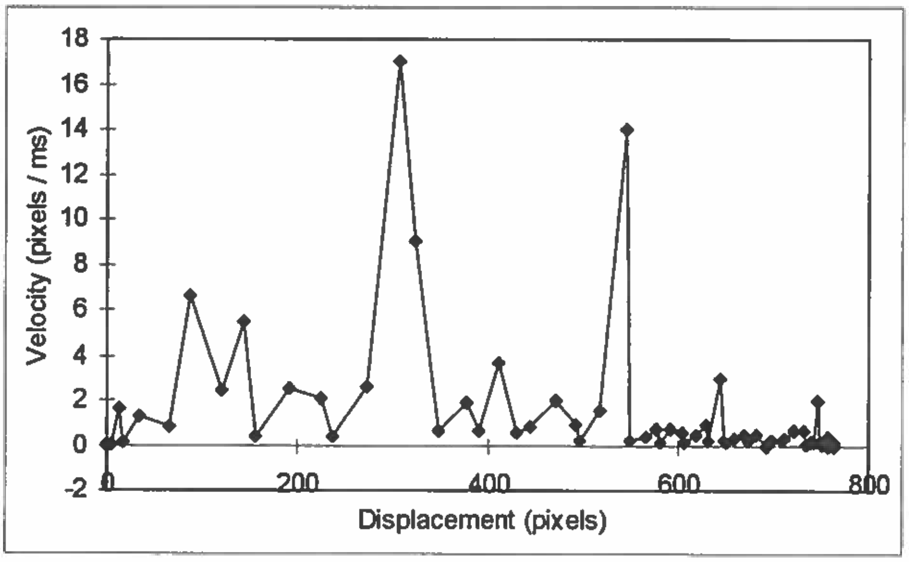

Figure 36. Last joystick trial subject DD1 - raw data.

Joystick trials were thus characterized by jerkiness that is visible only when the velocity of the trial is studied either over time or over movement distance. The jerkiness was evident in early trials in the pilot study and persisted throughout the duration of the main study.

Effect of Filtering Data

The process of taking a digital sample of a continuous waveform, in the manner in which the position of the cursor was sampled by the program, introduces sampling error. In addition, taking the numerical differential of time series data introduces high frequency error. Knowledge about the signal being measured can be exploited in order to reduce the amount of error. Other research has shown that human movement tends to lie in the frequency range below 10 Hz, and the upper limit of the range is 20 Hz (Stein & Lee, 1989). The sampling error can therefore be removed by running the data a low pass digital filter to remove all the frequency components higher than 20 Hz (Kaiser & Reed, 1977).

Prior to starting with the details of this analysis, we discuss a problem in the data gathering routine. The software that was written to gather the data sampled the position of the cursor on the average every 16 ms. It was not, however, written to make uniform samples at every 16 ms, only to average 16 ms. The reason for this oversight was inexperience with numerical analyses and Fourier methods. Most Fourier techniques assume that the sample rate is uniform. Fourier techniques exist for non-uniform sampling, but in general, they are more difficult to implement and are slower. In addition, standard math packages include only fast Fourier transforms that assume uniform sampling.

The data generated by the program was non-uniform for two reasons. First, it was written to respond to MOUSE-MOVE messages generated by the operating system (OS). These messages are, obviously, not generated when the mouse is not moving, which can happen between submovements. This results in large 'gaps' in the data (when we know that the cursor maintains the same position), but the time between samples is much greater than 16 ms. Second, because the messages are buffered by the OS and presented to the application program asynchronously, and because adjacent messages combined into a single message, MOUSE-MOVE messages are not sent to the application program uniformly.

The non-uniformity of the data required the addition of an expansion step to fill in gaps in the data when the cursor was not moving, and to then extrapolate the data to generate uniform samples. Generally speaking, it is best to massage the data as little as possible, and it would have been better not to have had to perform these steps. Having said that, plots of the data with and without the expansion and extrapolation steps were very similar. In addition, both mouse and joystick received exactly the same treatments, and it is reasonable to assume that any differences that we find at the end of the numerical analysis are due to differences in the devices.

The digital filtering of the data was able to remove a large amount of the effect of sampling error. The resulting plots for both devices were smoother than the plots of their raw data, but the plots for the joystick remained more jerky than the plots for the mouse. In other words. tremor persisted in the plots despite the low pass filtering. This is consistent with the literature which says that the frequency range of tremor lies below 20 Hz., and the filtering removed frequencies higher than 20 Hz.

Figure 37 and Figure 38 are samples of mouse data after filtering. They represent the same trials as shown in Figure 28 and Figure 29 respectively. They show a filtered approximation of the subject's trajectories towards the target. Figure 37 shows that the subject made 8 submovements towards the target, with 4 overshoots, and hit the target on the 5th attempt. Figure 38 shows a fairly straightforward movement consisting of three rapid adjustments while at speed towards the target, with additional small adjustments towards the end.

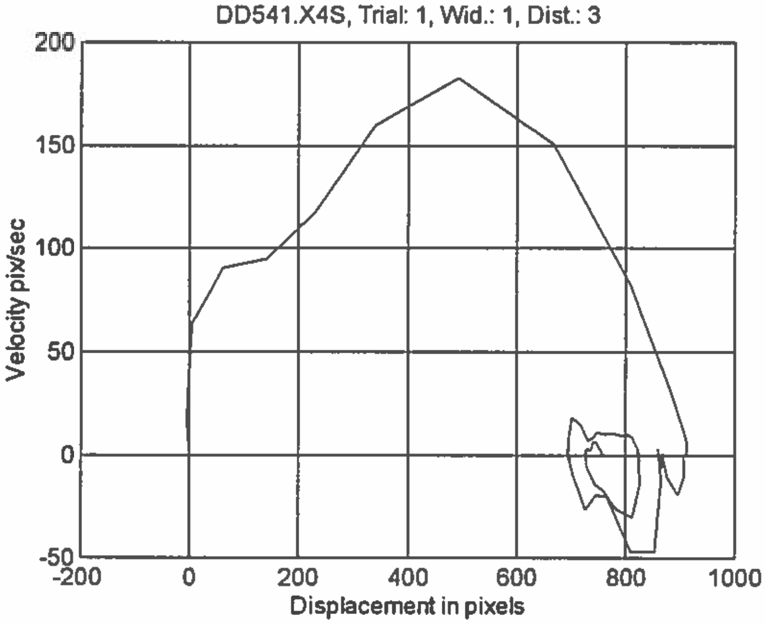

Figure 37. First mouse trial subject DD5 - filtered data. Same trial as Figure 29. A - 768 pixels.

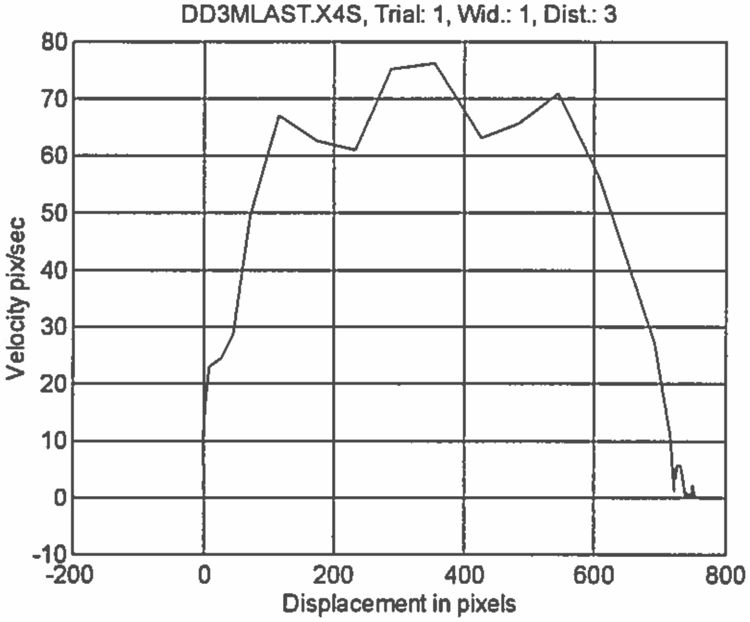

Figure 38. Last mouse trial subject DD3 - filtered data. Same trial as Figure 28.

Figure 39, Figure 41 and Figure 40 are smoothed joystick trials. From Figure 39 it is obvious that tremor persists despite the smoothing effects of the filtering algorithm. Tremor is visible in Figure 40. Figure 41 depicts a trial where the effects of tremor are less visible. This trial also shows practice effects, which are discussed later in this chapter.

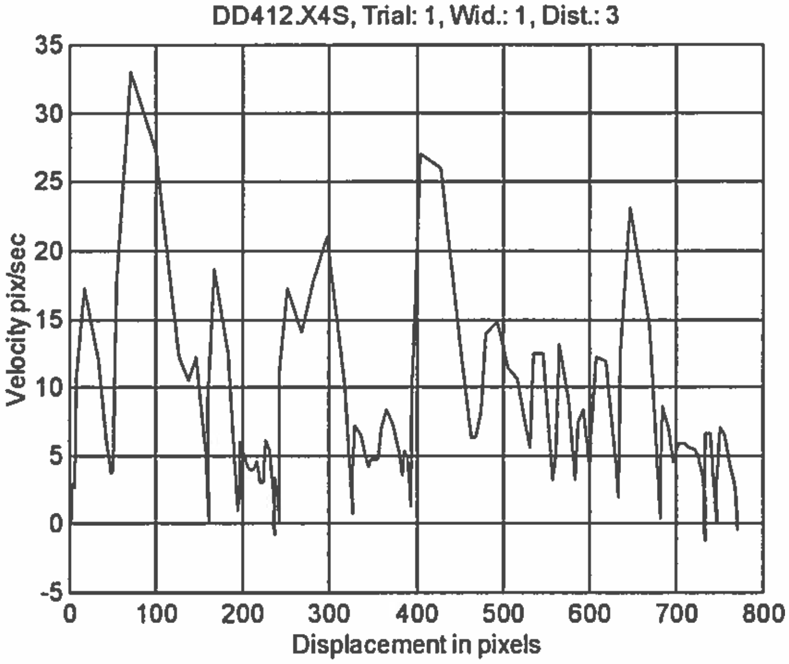

Figure 39. First joystick trial for subject DD4 - filtered data. Same trial as Figure 34. A = 768 pixels.

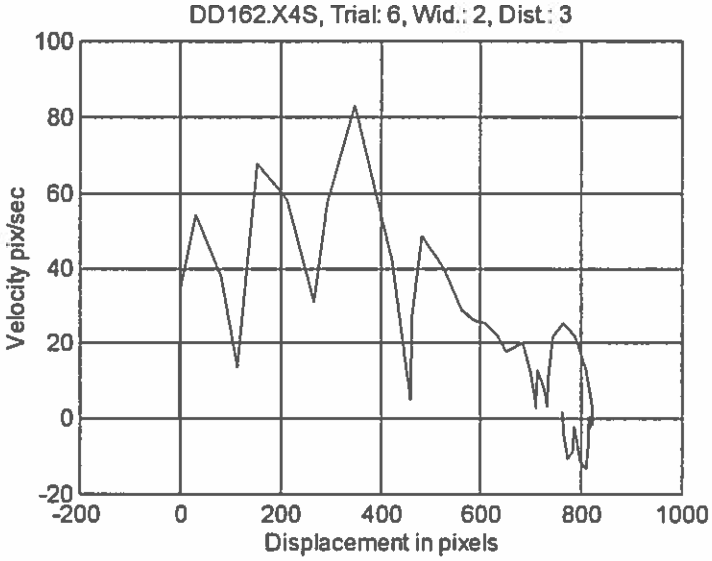

Figure 40. Trial from last session by subject DD1-Smoothed. A = 768 pixels.

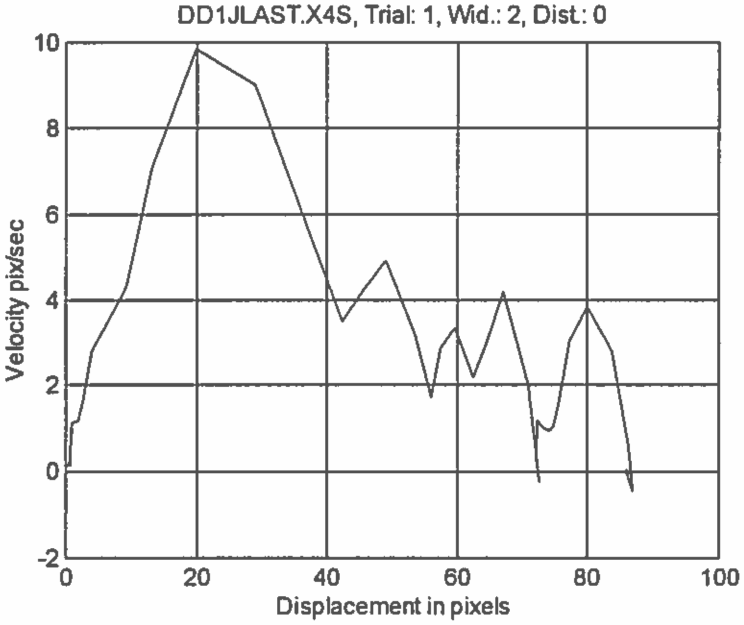

Figure 41. Last joystick trial by subject DD1 - Smoothed.
A = 96 pixels.

Individual Differences

Two subjects, DD4 and DD5 had high performances with the isometric joystick (see Figure 20). We were interested to see whether there were any differences between their microstructure of movement and those of the other subjects. Subject DD4 reported having had a lot of experience playing video games that used an isotonic joystick, but had never used an isometric joystick before. Subject DD5 owned an IBM ThinkPad and had used a TrackPoint for six months prior to the experiment. Interestingly, subject DD6 had the most experience with the TrackPoint. He had owned an IBM ThinkPad for approximately 2 years before the experiment and was much more comfortable using it than using a mouse. However, he did not have the highest performance with it relative to the other subjects in the experiment, though his final trial times lay in the top half.

Figure 42 and Figure 43 show two representative graphs of trials by subjects DD4 and DD5. Their trials appear to have a mouse-like quality, i.e., episodes of movement that reach large velocities and are relatively free of jerkiness. This suggests that individual differences may account for differences in performance. These implications are discussed in more detail in the next chapter.

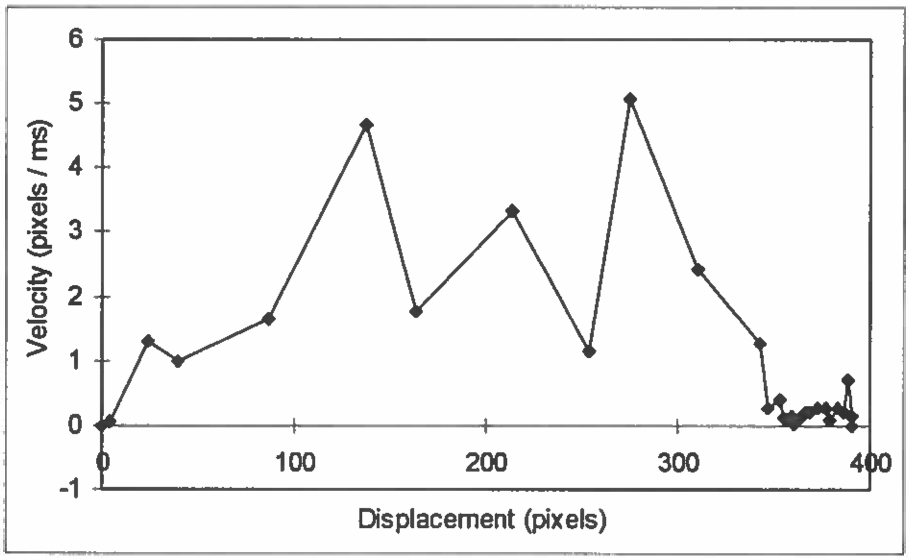

Figure 42. Middle joystick trial for subject DD4

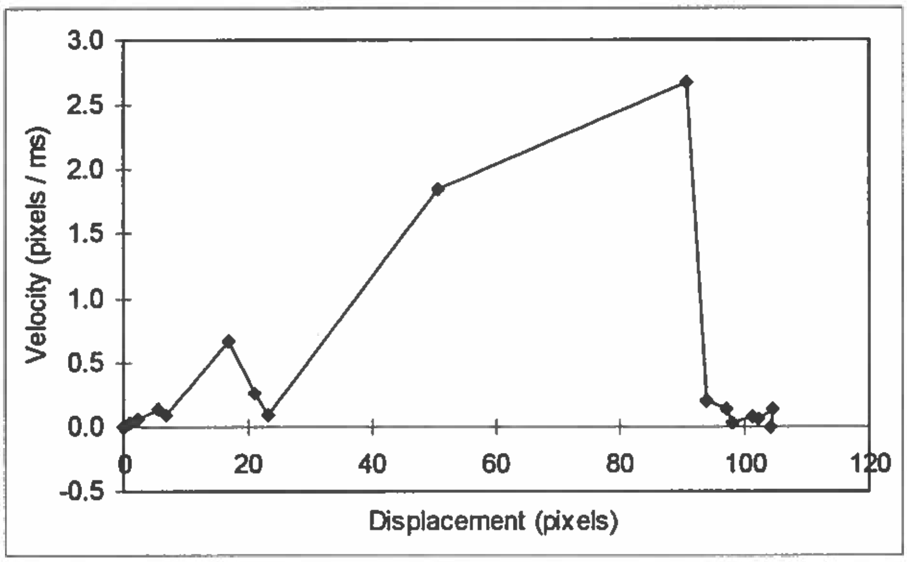

Figure 43. Subject DD5's last joystick trial.

Applicability of the SOS Model to Joystick Movement

One of the goals of this dissertation was to extend the SOS model to the joystick. It had already been shown to be fairly good at modeling movement with the mouse (Walker et al., 1993). If the model could be extended to the joystick, then, because of the differences between the joystick and the mouse discussed in Chapter II, we would be more inclined to believe that the SOS model extends to all pointing devices.

Unfortunately, the jerkiness in movement due to tremor that was uncovered by this study, while exciting from the research point of view, also caused analytical problems. The problem is that we have been unable to separate the voluntary (programmed) movements from the involuntary movements (tremor/noise).

The parser used by earlier researchers depends on transitions in velocity and acceleration to determine the start and end of a submovement. Unfortunately, the tremor appears as exactly such transitions in velocity and acceleration, and the parsing algorithm cannot distinguish between what is a submovement, and what is tremor.

Because of this analytical problem, the SOS model analysis could not be performed for the isometric joystick. It is possible that some analytical technique might be found in the future which will allow researchers to distinguish between tremor and voluntary movement, but such a method was not at our disposal.

One feature of this experiment and its analysis has been that for every treatment and analytical technique that was applied to one device, the exact same treatment and analysis has been applied to the other device. Partly for this reason, and partly because it would not substantially increase our knowledge, it was also decided not to do the SOS model analysis for the mouse.

Having said that, many weeks of looking at velocity vs. distance plots for the mouse has shown that generally speaking there is one larger initial submovement, some smaller subsequent movements, and thus there is every indication that the SOS model holds for the mouse.

The same cannot be said for the impressions obtained by looking at plots of the joystick data, because in most trials the tremor completely masked from the eye the underlying pattern such as it was, and it cannot be said with any certainty whether or not the SOS model holds for this device.

Characterizing Tremor

The analysis up to this point has been able to show that a jerkiness in the velocity of the isometric joystick exists, but so far these have only been qualitative measures. This section describes a series of analyses that were aimed at getting a better picture of the effect that tremor has on movement with the joystick.

A comparison of a mouse trial (e.g. Figure 27) with a joystick trial (e.g. Figure 32) suggests that the jerkiness of the joystick trial is caused by the superposition of a higher-frequency signal on a low-frequency signal. Tremor is a form of random noise, and the presence of frequencies in the data is stochastic in nature. If a single trial with the joystick is made up of a series of 4 or 5 submovements towards the target (a supposition that is not unreasonable based on the data obtained), and the average time that a joystick trial takes during the last block is 950 ms, then the approximate frequency of the submovements is about 5 Hz. The joystick data shows the tremor occurring at somewhat higher frequencies.

This sort of data can be analyzed in the frequency domain using Fourier techniques. Based on the analysis above, we would expect that a greater percentage of the power in a signal would exist at higher frequencies for the joystick than the mouse, a situation depicted in Figure 44.

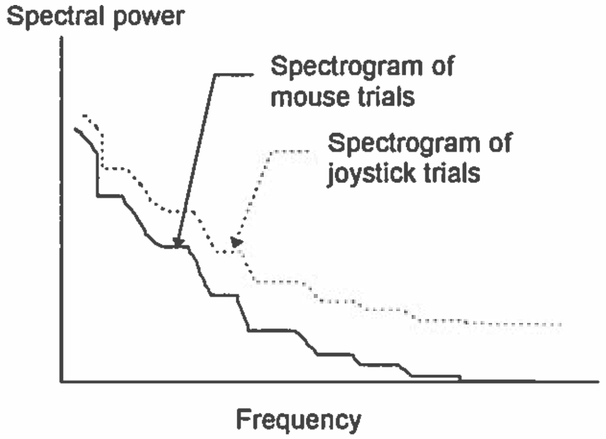

Figure 44. Comparison of expected normalized spectrograms of mouse and joystick.

The following analysis was performed using a mathematics package called Matlab, published by Mathworks, and available in a Student Edition from Prentice Hall (Student Edition of MatLab for Windows, version 4, ISBN 0-13-184995-6). The frequency domain analysis of the trials from the last block for all subjects was conducted using the numerical and signal processing functions in Matlab.

First, each trial was stretched to cover gaps in the data when the cursor was stationary, then it was interpolated to generate uniform samples. The resulting data was passed through a low-pass Butterworth filter and then analyzed in the frequency domain.

As expected, most of the power of each trial is in very low frequencies. The power also varies widely from trial to trial, which is not informative. The spectrograms of individual trials differed greatly from subject to subject and between devices, and no pattern was apparent either between subjects or between devices, or even within devices, apart from the very general pattern depicted in Figure 44. This pattern tended to have greater power in the low frequencies, and lower power in the high frequencies.

In order to get a better understanding, the spectrograms for each trial of the last block for each device were aggregated by subject. This produced an averaged spectrogram for the last block for each subject with each device. The power was normalized prior to addition by dividing the power at each frequency by the power of the frequency with the highest power. This generated a power spectrum where the highest power was 1. When aggregated over all trials for a subject, the maximum power rose to 120. The results are presented in the next few pages.

Figure 45 represents the spectrogram for the last block of mouse trials for each subject. Notice a slight bulge at approximately 4 Hz for DD6. This suggests that DD6 might have been using a different strategy for controlling the mouse from the rest of the subjects. Similarly, Figure 46 is the spectrogram of the last block of joystick trials for each subject. Here subject DD1 has a slight bulge at approx. 4 Hz. DD6 had the poorest mouse scores and DD1 had the poorest joystick scores.

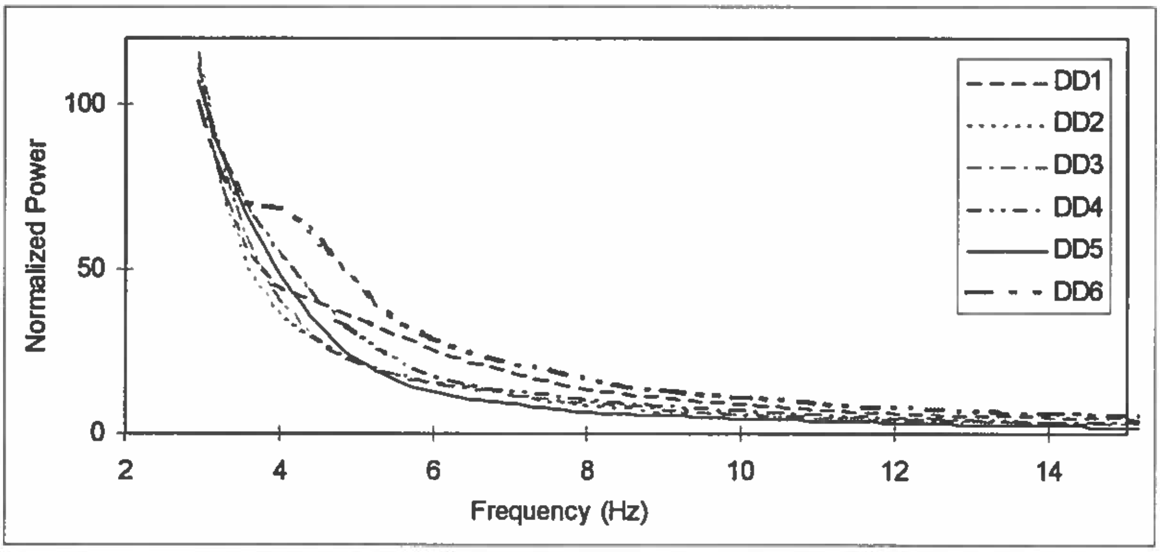

Figure 45. Spectrograms for last block of mouse trials.

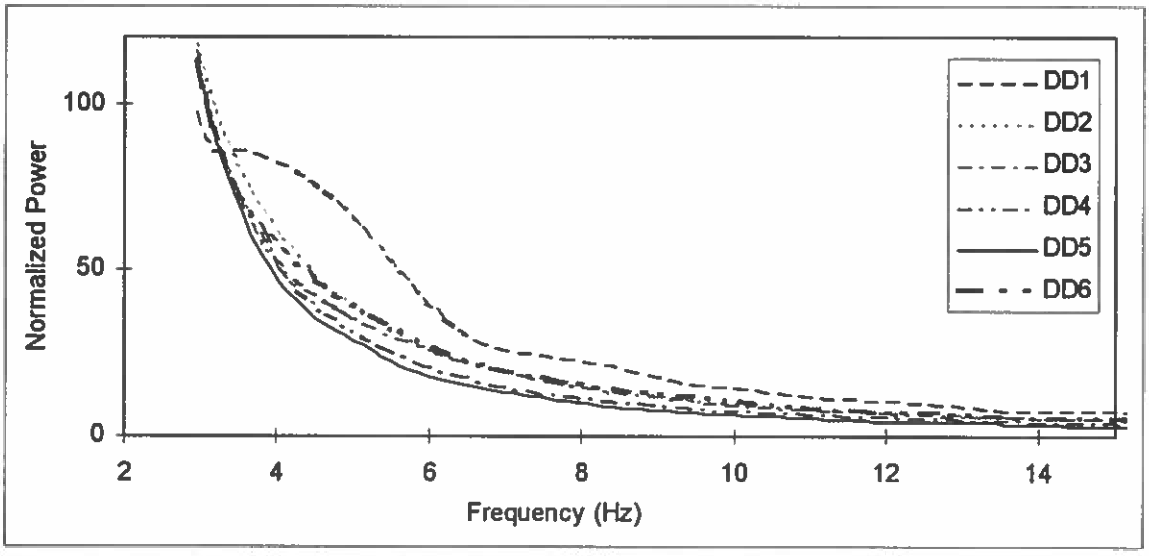

Figure 46. Spectrograms for last block of joystick trials.

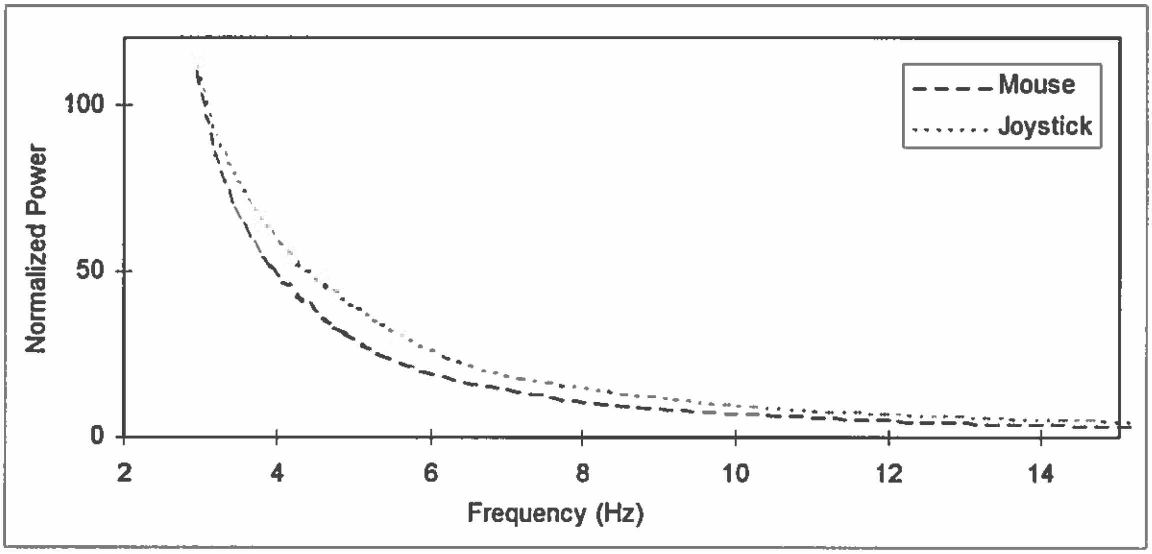

Figure 47. Comparison of spectrograms for both devices.

Figure 47 presents the frequency data for the mouse for all subjects compared to the frequency data for the joystick all subjects. It clearly shows a that there is a greater percentage of the power for higher frequencies (between 3 and 10 Hz) for the joystick than for the mouse, which is what the earlier analysis predicted.

The shape of the curves in Figure 45, Figure 46 and Figure 47 suggest an exponential relationship, and plotting the data on a log-log chart produces Figure 48 which shows that the normalized power and the frequency for both mouse and joystick appear to have a polynomial relationship. A regression analysis produces the following results:

$$\log(m) = 2.64 - 1.73 \log(f); \quad r^2 = .98 \tag{22}$$

$$\log(j) = 2.71 - 1.68 \log(f); \quad r^2 = .99 \tag{23}$$

To summarize, the spectral analysis showed that there are more high-frequency components to the joystick movement, which is consistent with the presence of tremor.

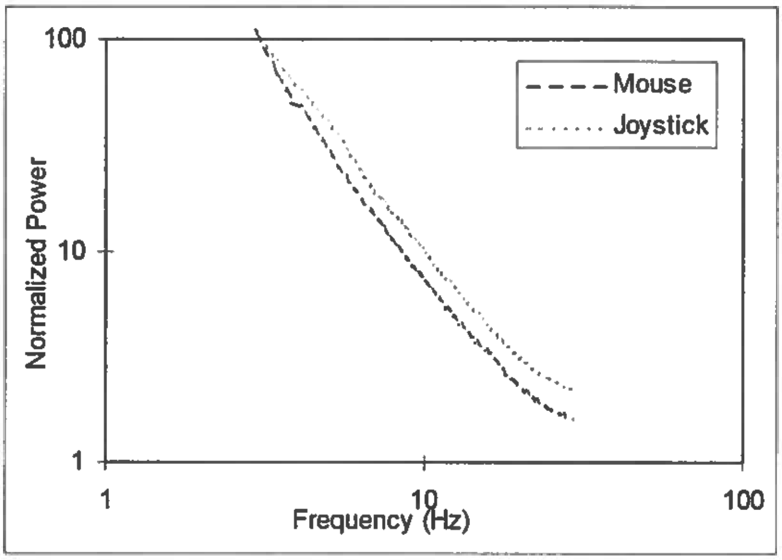

Figure 48. Log-log plot of frequency components for both devices.

Individual Differences

An earlier section in this chapter discussed the manner in which individual differences appeared to play a role in the performance of the joystick. It noted that that subjects who had low trial times with the joystick had mouse-like movement curves. To explore this further, mouse and joystick spectra for each subject were compared.

Figure 49 through

Figure 54 plot the spectral analyses of mouse and joystick trials for the last block of each subject. The plots for subjects DD1, DD2 and DD3 show a difference between the mouse data and the joystick data. There is more power in the joystick data than the mouse data between 3 to 15 Hz. This indicates that there are more high frequency components in the joystick data.

In this light, the plots for subjects DD4 and DD5, who had high performances with the joystick are very interesting. (Figure 52 and Figure 53). For them, these figures show a great deal of similarity between the spectra for mouse and joystick movement. Their plots do not show the separation that can be seen in the graphs of the other subjects. This suggests a relative lack of the presence of tremor in their trials. This issue is discussed further in the next chapter.

Figure 54 is also of interest. It shows the spectra for subject DD6, who did not like the mouse, and in fact had a slower performance with the mouse than the joystick. This subject had more power at higher frequencies for the mouse, and the bulge at around 4 Hz suggests some sort of stop-and-go pointing strategy.

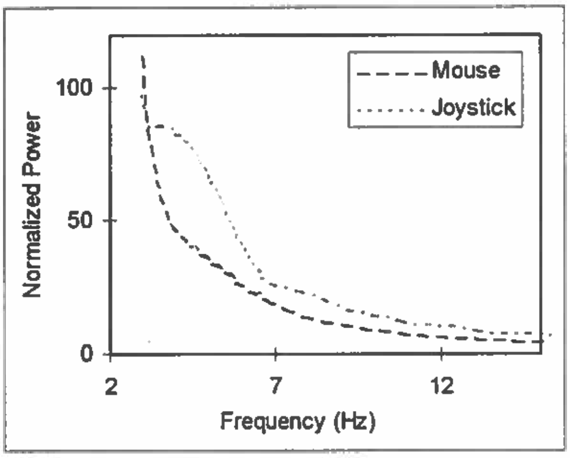

Figure 49. DD1's last block.

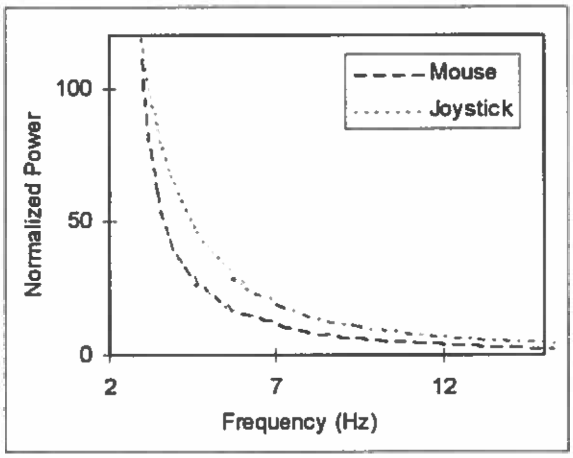

Figure 50. DD2's last block.

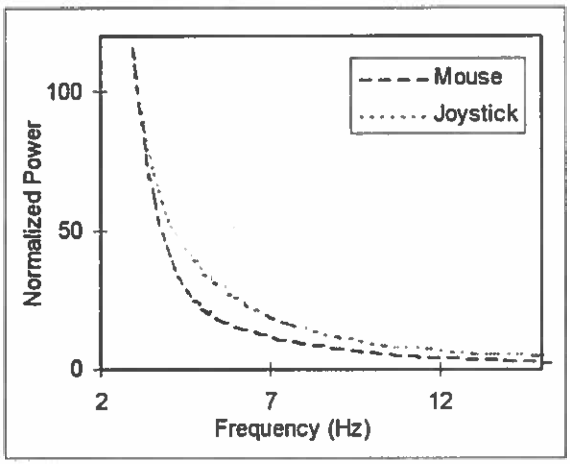

Figure 51. DD3's last block.

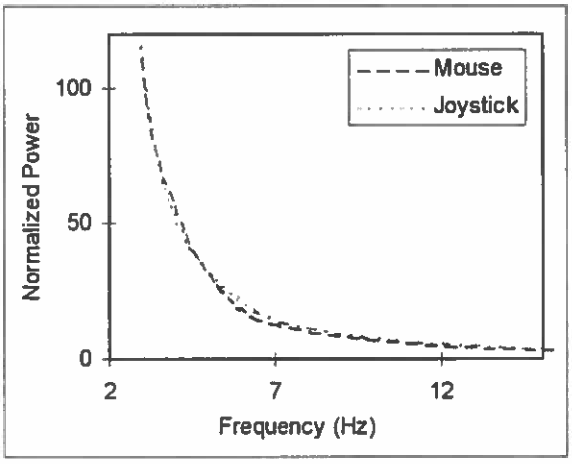

Figure 52. DD4's last block.

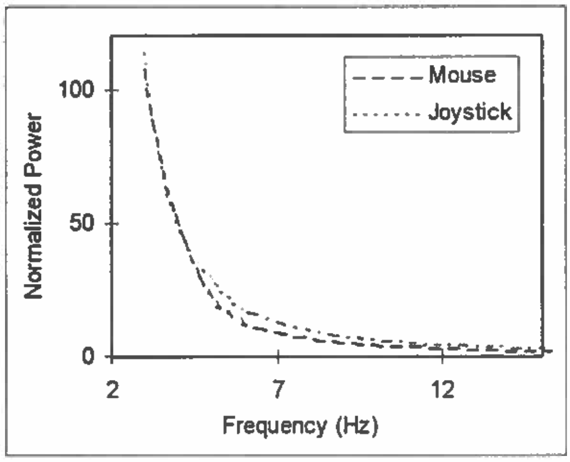

Figure 53. DD5's last block.

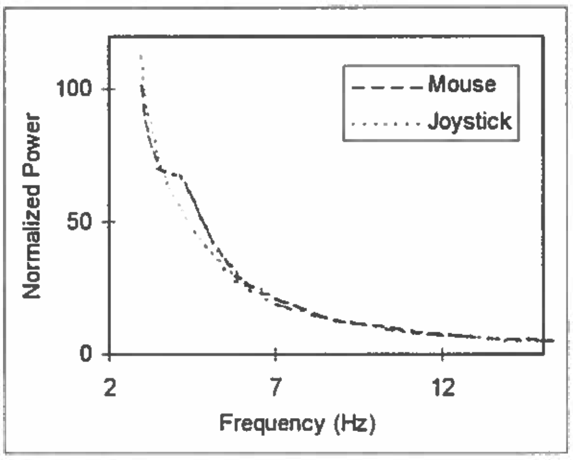

Figure 54. DD6's last block.

Summary

In summary, two kinds of analyses were carried out on the data from the experiment, gross level analyses and microstructure level analyses. The gross level analysis, which included a Fitts' law analysis, showed that there was nothing unusual about the data. An analysis of learning over time showed that we can be confident that the last blocks for both devices represent practiced performance.

The microstructure level analysis showed that the tremor that was found in the preliminary tests also appeared in the main experiment. This tremor persisted after the data was filtered to remove sampling error. Unfortunately, the tremor also made the SOS model analysis impossible. Comparing trials of subjects with good performance on the joystick against those with poorer performance suggested that there were qualitative differences between the manner in which these subjects used the joystick.

Spectral analyses of the trials from the last block with each device showed that there are more high frequency components for the joystick than for the mouse, providing analytic evidence over and above visual evidence that the joystick is affected by tremor. The subjects who had good performance with the joystick showed smaller amounts of this difference.

The implications of these results are discussed in the next chapter.

CHAPTER V

DISCUSSION

There are a wide range of pointing devices currently in use, and they have widely varying performances. This study was aimed at understanding why these differences in performance exist, and why some pointing devices are more difficult to use than others. An analysis of the microstructure of movement was viewed as the means for getting this answer. A set of research questions were then developed to guide the study.

We begin this chapter with a discussion of the research questions and the answers obtained by the study. The discovery of tremor lead to an additional line of study which was not part of the original research agenda, and the second section summarizes what we know about tremor. The following section discusses the steps taken to ensure that the jerkiness in the velocity of joystick trials was not an artifact of the experiment. After establishing the validity of the results, we discuss how tremor develops in the isometric joystick. The last section discusses the results of the Fourier analysis that was used to quantitatively analyze the tremor.

Review of Research Questions

There were three groups of research questions, questions about Fitts' law, questions about the microstructure of movement, and questions about the applicability of psychomotor models to pointing devices. This section groups the questions under these headings followed by a discussion of each question.

Fitts' Law and its Applicability to the Devices

1. Do the devices (mouse, and isometric joystick) follow Fitts' law? What are the constants in their Fitts' law equations? How do the indices of performance for these devices compare with one another?

The Fitts' law question had two goals. First, both pointing devices, the mouse and the isometric joystick are known to follow Fitts' law, and it is important to know whether the data from this experiment shows the same results.

The analysis of gross movement showed that indeed both the mouse and isometric joystick follow Fitts' law. The mouse had an IP of 8.13 bits/sec, while the isometric joystick had an IP of 5.75 bits/sec. This gave a ratio of IP_{mouse} : $IP_{joystick}$ of 1.41 (see Table 8, page 94) which is somewhat low. Earlier research indicated the ratio of the IPs for a mouse to a finger-controlled isometric joystick was close to 2.0 across a number of studies (Douglas & Mithal, 1994).

Another variation from earlier research was that the mean trial completion time for the last block was 860 ms for the mouse, and 946 ms for the joystick. The completion times are lower than those obtained in other studies. The Keystroke Level Model suggests a pointing time with the mouse of approx. 1100 ms (Card et al., 1980; Card et al., 1983), while our earlier research suggests a time of 1700 ms for joystick pointing tasks (Douglas & Mithal, 1993).

A number of reasons could account for this difference. The most important of these is the simplicity of the task in this study, i.e., movement in only one direction, and fewer combinations, target positions and decisions that the subjects had to make. In addition the size of the targets used was larger than those used in the earlier study (Douglas & Mithal, 1993). The present study used targets of 16, 32 and 64 pixels, while the earlier study used targets of 8, 16 and 32 pixels. Finally, the large number of repetitions and the use of computer experts could have contributed to the low trial completion times.

Some studies have shown that movements with the right hand moving directly to the right have higher performance than in other directions (Boritz et al., 1991). The present study exclusively examined movement in that one direction, while other studies combined data of movements from two or more directions.

From a broader point of view, the gross movement analysis showed that both devices followed Fitts' law ($r^2 = 0.998$ for the mouse and 0.991 for the joystick), and the mouse had a higher IP and lower mean trial completion time.

Analysis of the Microstructure of Movement

The second group of questions focused on the microstructure of movement for the two devices.

2. What is the microstructure of movement for the isometric joystick in terms of the duration and accuracy of the first and subsequent submovements? What is the microstructure of movement for the mouse?

The microstructure of movement for the mouse followed the primary assumptions of the SOS model. On the average the first submovement was larger, faster and more accurate than the rest of the submovements. This qualitative conclusion was reached by a visual examination of the velocity vs. distance graphs for the mouse.

For the joystick, tremor prevented a quantitative analysis from being performed because it was not possible to remove the involuntary tremor from the voluntary submovements. Because the microstructure level analysis could not be performed, it was not possible to compare the properties of the first and subsequent submovements for the joystick. The tremor was so pronounced in the joystick that it prevented even qualitative assessments of the nature of the primary and subsequent submovements by a visual examination of the velocity vs. distance plots. As the quantitative analysis could not be performed on the joystick, it was also dropped for the mouse.

3. Are there differences in the microstructure of movement of the mouse and isometric joystick, and do these differences explain the difference in performance?

In part we have been able to answer this question, and in part we have not. If the jerkiness caused by the tremor is considered part of the microstructure, then the answer is yes, differences in movement can be linked to differences in performance by saying that the jerkiness in velocity makes it more difficult to control the joystick, which degrades its performance.

On the other hand two issues were kept in mind while framing the original question. One was the lack of kinesthetic feedback that users get with an isometric device. The other was that the isometric joystick is velocity controlled, while the mouse is a displacement control device. From this point of view, we wanted to know if the user might have different strategies for controlling the two devices, and so the part of the microstructure of movement that we were interested in comprises voluntary movements by the user. From this point of view the question cannot be answered it could not be determined what the voluntary movement for the isometric joystick was.

Individual Differences

4. Are there changes over time in the microstructure of movement for individual subjects' performance?

There was not much change in the movement microstructure for subjects using the mouse over the duration of the experiment. This is not surprising because they were all experts in mouse use prior to the study (except DD6). The mouse trials showed prototypical SOS model movement throughout the experiment.

At the start of the experiment, all subjects using the joystick had trials where the velocity of the cursor was jerky . Over time, the shape of the velocity vs. distance plots of the trials began to appear more mouse-like. This was particularly marked for subjects DD4 and DD5. This aspect is discussed below.

5. Are there differences in the microstructure of movement between individual subjects?

Examining the plots of velocity vs. distance between individuals can suggest reasons why there might be differences in individuals' movement microstructure that can account for differences in their performance with a device.

This idea was examined for the isometric joystick. Subjects DD4 and DD5 had mean trial times of 797 and 805 milliseconds respectively, as compared to the group mean of 946 ms. On examining the joystick trials for these subjects (e.g. Figure 42 and Figure 43), it appears that many of their trials have episodes of high cursor velocity that are relatively unaffected by tremor. This suggests that these subjects might have been able to overcome the effects of tremor in some manner. In addition, some trials by other subjects had similar characteristics. These trials would have smaller trial completion times (a mathematical property of the velocity-distance graph - if the velocity is high over a large distance, the time taken is low). So it can be suggested that when trials are more mouse-like, or more in accordance with the assumptions of the SOS model, they are faster.

If the two faster subjects are able somehow to overcome the tremor, then we expect to see smaller amounts of the evidence of tremor in the spectrograms of their trials.

The spectrograms of the last blocks with both devices for all subjects are shown in Figure 49 through

Figure 54.

Figure 53 and

Figure 54 are the spectrograms for subjects DD4 and DD5 respectively. For subject DD4, the two curves are remarkably similar, and for subject DD5, there is a distinct overlap between the curves. The curves for the rest of the subjects show a separation between the joystick and mouse curves as the frequency increases. This indicates that there are not as many high frequency components in subject DD4 and DD5's joystick trials as there are for the rest of the subjects and hence less tremor. Which in turn supports the argument that subjects with higher performance with the joystick might be able to suppress or overcome tremor in some fashion.

These results about individual differences have been shown only for a small sample, and need to be replicated with a larger sample before they can be generalized.

Modeling Isometric Movement

The third group of questions looks at modeling movement from the data of the microstructure level analysis.

6. Do models of movement that describe isotonic devices describe isometric devices? In particular, does the SOS model describe the movement of both the mouse and the isometric joystick?

7. If the SOS model does not describe the movement of the isometric joystick, can we modify the model so that it does?
8. Can we use our model for the isometric joystick and the knowledge about its movement microstructure to improve its design?

The SOS model analysis is based on the analysis of the microstructure of movement. As a result, the problems that made the microstructure level analysis difficult also plagued questions about the SOS model and the application of any other model to isometric movement. Because we cannot determine the underlying movement, we cannot perform an analysis that would refute or support the applicability of the SOS model to the joystick. This also means that at the present time we have not been able develop any alternative model to explain isometric movement.

We were, however successful in the underlying goal of learning more about how the isometric joystick is used, and as the following discussion will show, there are opportunities to use this knowledge to improve its design.

The Effect of Tremor on the SOS Model Analysis

The original idea of trying to model pointing behavior is a very strong idea, and is valid despite the fact that the method of analysis failed due to the tremor. A major contribution of the SOS model is that it hypothesizes an intentional plan on the part of the user. This hypothesized plan is the initial programmed (ballistic) submovement and the and the subsequent corrective submovements. The plan might be unconsciously developed and executed, but it is intentional in its nature, and the SOS model's great contribution is in trying to make it explicit.

With such model of the user's intentions we are better equipped to design pointing devices to meet that intentionality, and we need to develop such a plan for isometric devices.

What broke down in the modeling process was not the model, or the idea that there should be a model, but the analysis technique. The analysis technique, in looking at the microstructure of movement, measured the subject's intentionality indirectly and looked at the displacement, velocity and acceleration. What we need to do to model the joystick is to use a more direct measure of intentionality. Possible techniques that come to mind include studies of neural firing and muscle activations. These might help us get closer to understanding what the user intended to do rather than what actually happened.

Tremor and the Voluntary Application of Force

Tremor was the most important finding of this study, and this section is a brief discussion of what is known about tremor.

All forms of voluntary movement have some time period associated with them. For instance, the early work by Woodworth noted time periods of 300 ms. for the initial movement (Meyer et al., 1990). It has also been noted that in trying to maintain a steady posture or perform a voluntary movement, involuntary variations in force occur, these are described as tremor (Stein & Lee, 1989).

Tremor is visible in most individuals when holding a limb outstretched. If you hold your arm out with your fingers extended, you are likely to see a ‘twitch’ in the fingers, and possibly even see that you cannot hold your arm completely steady. The twitch is caused by tremor, which is an inability of the human neuro-muscular system to maintain a constant force. Stein and Lee (1989) report that the range of frequencies generated by an outstretched limb span the range of frequencies present in voluntary movement.

The largest components of tremor occur in the low frequency range, usually below 5 Hz, and these movements are quite random. Tremor occurs during the application of steady force in the absence of movement, such as when holding down a button, or pushing an isometric joystick. In such situations, some amount of damping for the tremor is caused by the friction between muscle fibers. Tremor also occurs when the applied force causes movement, such as when the arm moves across the body in an arc, or during mouse movement. Tremor even occurs at rest, because of the variations caused by the blood pulsing through the veins and arteries (Stein & Lee, 1989). Finally, tremor tends to increase with increases in applied force. A number of pathological conditions such as hyperthyroidism and Parkinson’s disease tend to increase the amount of tremor exhibited by an individual.

While a large portion of tremor is quite random and generally occurs at low frequencies, some components are more oscillatory and occur at slightly higher frequencies, generally between 8 - 12 Hz. This oscillatory tremor appears to be caused by feedback mechanisms between the skeleton, muscles and the nervous system.

In general, tremor occurs with all application of force, and is a random noise that occurs in the same frequency range as voluntary movement.

Isolating the Source of the Tremor

Discovering tremor and its affect on the joystick was a surprise and a number of issues were taken into account to ensure that the jitter was not an artifact, and that it was indeed caused by the joystick picking up tremor in the finger.

Sampling Rate

A major concern at the outset of the development process for this experiment was that the rate at which mouse position updates are received by the program were at a rate high enough so that the underlying continuous signal could be accurately reproduced. This became even more of an issue when the tremor was discovered because it became important to be sure that the jerkiness observed was not a result of the data gathering process.

The Nyquist theorem tells us that if a signal is comprised of frequencies up to N Hz, it must be sampled at a frequency of 2N Hz, in order for the signal to be accurately reproduced. Most human movement occurs at frequencies of lower than 10 Hz, corresponding to a time period of 100 ms. The smallest period associated with human movement is approximately 60 ms, which corresponds to a frequency of 16.7 Hz. The sampling rate needed to measure this signal is 33 Hz, or approximately once every 30 ms. This was used as a benchmark.

The original study by Jagacinski used external analog recording devices that measured joystick displacement every 1 ms (Jagacinski et al., 1980a). Unfortunately, we did not have access to this kind of equipment. The recent study that examined the applicability of the SOS model to the mouse tracked the mouse with updates occurring every 17 ms using software alone (Walker et al., 1993). This particular software could not be used for this study because it was not compatible with the TrackPoint joystick.

The IBM PC has a timer chip that is updated 17 times a second, or every 58 ms. The first version of the data-gathering software was written in the standard Windows 3.1 SDK (software development kit) which relies on the IBM PC timer, and therefore produces time samples that are accurate only to 58 ms, with a sampling frequency of 17 Hz and a Nyquist sampling frequency of 9 Hz. This was obviously unsatisfactory, leading to the development of faster routines.

The multimedia extensions for Windows require more accurate timing mechanisms than those available in the standard Windows 3.1 SDK, and rewriting the software to use the Windows multimedia API produced mouse position updates at approximately every 26 ms, which corresponds to a frequency of 38 Hz, and a Nyquist sampling frequency of 19

Hz. This started to approach an acceptable value, but did not leave much margin for error.

Microsoft Windows NT is an advanced operating system that provides preemptive multi-tasking, threads, and thread priority. It is a 32-bit operating system as opposed to Windows 3.1, which is a 16 bit OS. When a 16-bit application is run under NT, it runs in a high-priority, virtual Windows 3.1 environment. Running the code under Windows NT increased the sampling rate to 16 ms.

This prompted a trip to Microsoft to use their development tools and the hardware they had at their disposal to run some tests. One test was to recompile the code to 32-bit Windows code. Tests running with this software did not, surprisingly, produce faster code, but slower code. Apparently the manner in which NT virtualizes 16-bit code leads to a higher priority.

A second approach was to modify the driver code to produce a pointing device driver that generates faster interrupts. The PS/2 pointing device port uses the i8042 driver. The driver was modified to generate interrupts every 10 ms. Running the 32-bit software with this modified driver generated far more mouse position interrupts, but also resulted in a great deal of relative error.

Closer analysis showed that this was in fact not surprising. Time stamping a cursor position requires two actions, one to get the current position of the cursor, and another to get the current timer value. In a computer with a single CPU, these processes have to take place in different CPU cycles. This is illustrated in Figure 55, where the sampling interval is i. At some finite time c after i, (c can be zero), the computer gets the current cursor position. At some other finite time t after i, the computer gets the value of the clock. t cannot be equal to c, and there is relative error $r = |c-t|/i$ in the measurement. As the sampling rate goes up, i becomes smaller, but c and t stay the same. As a result, r becomes bigger. The relative error is not much of an issue when plotting cursor position as a function of time, but becomes an issue when we are plotting velocity, which was a central part of this study.

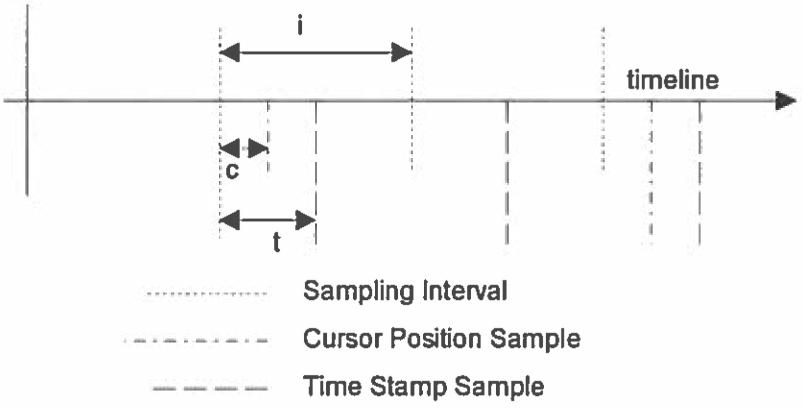

Figure 55. Sampling error

The ideal situation would have been to have some external device measuring cursor displacement as a function of time. This would have made the measurement more accurate, though at the expense of having to synchronize the stream of displacement data and time data.

Given the range of options tried, the experiment was conducted using the second version of the software that used multimedia APIs and timers, and run under Windows NT. This generated a mouse position update approximately every 16 ms. The sampling was therefore at a rate of 62.5 Hz, with a Nyquist sampling frequency of 31 Hz. This was almost twice the 16.7 Hz of the highest frequency observed in human movement.

The relative error in time, r, has a smaller time period than the sampling frequency, and thus occurs at a higher frequency, i.e., greater than 62.5 Hz. A low pass digital filter was used to reduce the effect of the sampling error.

We were thus confident that the data was being sampled at a high enough frequency, and that the jerkiness in the velocity that appeared was as a result of tremor, and not as a result of the manner in which the data was gathered.

Other Possible Sources of Tremor

When the jerkiness in the velocity was discovered during the pilot study, an effort was made to determine the exact situation under which tremor plays a role. We found that tremor was observed only when an isometric joystick was used and in all situations when an isometric joystick was used. The other possible sources of tremor that were eliminated were:

1. Subject - multiple subjects showed tremor.
2. The test computer - multiple other computers showed tremor when using the TrackPoint as the pointing device.

3. The test TrackPoint - TrackPoints in other computers and a similar joystick in an NEC Versa notebook showed tremor. In addition the PortaPoint, a separate isometric joystick showed tremor.
4. PS/2 connection - the serial mouse did not show tremor, but when connected to the serial port, the PortaPoint did.
5. Sampling rate - tremor appeared both at a sampling rate of 26 ms (Windows 3.11) as well as at a sampling rate of 16 ms (Windows NT 3.5).

The electrical connections of the mouse and TrackPoint provide the most compelling evidence that the isometric joystick is the source of the tremor. The TrackPoint has a PS/2 mouse port. The main study was conducted with the mouse connected to this port. In this configuration, both mouse and isometric joystick feed their signals to the computer over the same set of wires. These signals are then interpreted by the same software driver, and then passed on to the data gathering program via the OS. Subsequently, the data gathered for both mouse and joystick receives the same treatment. Any differences are therefore differences between the devices.

Tremor and Microstructure of Movement with the Isometric Joystick

The jerkiness in the velocity domain caused by tremor did not appear either in the analyses of mouse movement, nor in plots of cursor position vs. time for the joystick. But when the velocity of the joystick was plotted, either against time or against distance, the tremor leapt out. On reflection, this is not surprising.

Consider Figure 56 which shows a block diagram of an isometric joystick. The joystick is a force sensitive device that translates input force into changes in cursor velocity. Note that because the joystick is using the same hardware connections and software drivers as the mouse, it must send data to the OS in the same form as the mouse. The mouse sends the OS data about displacement, and so the isometric joystick performs an integration of velocity over time in order to send similar displacement information.

Figure 56. Block diagram of isometric joystick. (Modified reproduction of Figure 9).

When there is tremor present in the finger, the tremor is perceived as changes in applied force. These in turn are mapped onto changes in velocity which of course get integrated into displacements and sent to the OS. As long as we look at plots of displacement, nothing is apparent. The moment we look at plots in the velocity domain, the tremor appears.

The isometric joystick happens to be unique in the manner in which it picks up tremor and the way in which it is affected by it. Consider for a moment the finger that controls the joystick. It is applying a lateral force on the joystick, and it is unsupported. It also does not move. Therefore there is little to damp the tremor. In other types of movement, tremor is damped by two things, inertial mass and friction. Because the finger does not move, inertial mass does not play a role and so does not damp tremor. Again, because nothing moves, friction does not play a role. In fact, the only damping factor that I have been able to uncover is the friction that exists between adjacent muscle fibers in the finger (Stein & Lee, 1989).

Tremor also occurs during mouse movement, but it does not have the same effect for two reasons. First, the mouse operates in the displacement domain, and is not sensitive to changes in applied force in the manner that the joystick is. It responds only to changes in displacement. The mouse is therefore not directly affected by tremor. Second, when a mouse is moving across a mousepad, any tremor is damped by the friction of the mouse and the hand on the mousepad, as well as by the mass of the hand, wrist and mouse (and perhaps the forearm). Due to these reasons, tremor does not appear in plots of mouse movement. So the physical construction of the joystick makes it susceptible to tremor in a manner that other devices such as the mouse are not.

The effect of tremor on other pointing devices is still an open question. Two devices that use relatively unsupported limbs are the trackball and the touch tablet. The trackball is typically controlled by a single thumb or single finger, while touch tablets are controlled either by a finger, or by a stylus. While inertial mass and friction still play a role in these devices, the lack of support for the controlling limb suggests that a similar study needs to be performed on these devices to determine how they are affected by tremor.

The joystick also accentuates tremor. This tremor is picked up by the joystick and passed on to the joystick's hardware controller. The joystick's controller employs an 'accelerated' design, i.e. it employs a non-linear translation between input force and output displacement. This transfer function accentuates the effect that changes in input have on the output (Rutledge & Selker, 1990). So the changes caused by tremor are exaggerated in the displacement data that is sent to the software driver. The software driver is in turn accelerated, and accentuates changes in displacement into larger changes in the position of the cursor. The net result is that the velocity of the cursor varies widely as it moves across the screen. Thus the design of the joystick makes it unique in the way that it amplifies tremor.

Slight changes in design might reduce this susceptibility to tremor. Figure 57 reproduces the velocity-distance plot of the PortaPoint shown in Figure 33. While jerkiness appears in this plot, it is much less pronounced than in similar plots for the TrackPoint (e.g. Figure 32). A possible that the reason for this is that the cap on the isometric joystick of the PortaPoint is not completely rigid, and deforms with the application of force. This small amount of deformation might bring the mass of the finger into play to damp out some of the tremor. Note however, that this device was not extensively studied, and this theory needs to be substantiated with additional research.

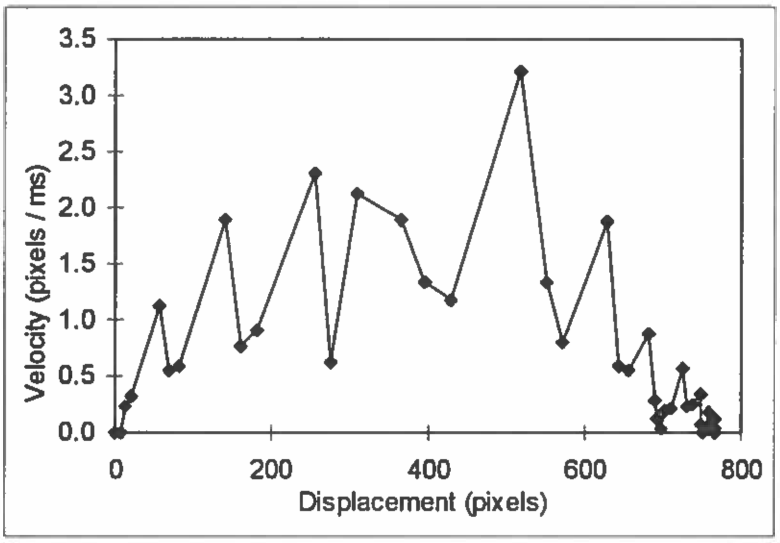

Figure 57. PortaPoint trial showing jitter. A = 768 pixels.

To summarize, studying the movement microstructure for the joystick revealed that tremor affects the velocity at which the cursor moves across the screen. Taking a closer look at the characteristics of the joystick and its connections to the computer showed that the isometric joystick is unique in the way that it is affected by tremor and the way in which it accentuates the effects of tremor.

Fitts' Law Analysis of Tremor

The Fitts' law equation indicates that the distance to the target (A) and the target width (W) should have strong effects on trial time. The last block was examined for this effect and the plots for this data are reproduced in Figure 58 and Figure 59.

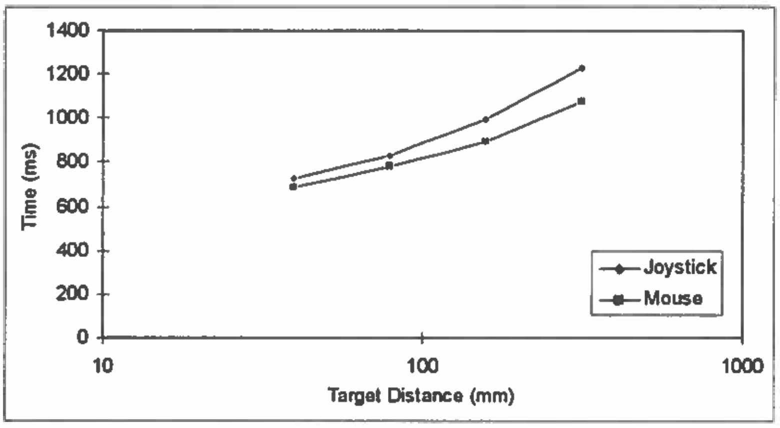

Figure 58. Effect of distance on trial time, n = 5.

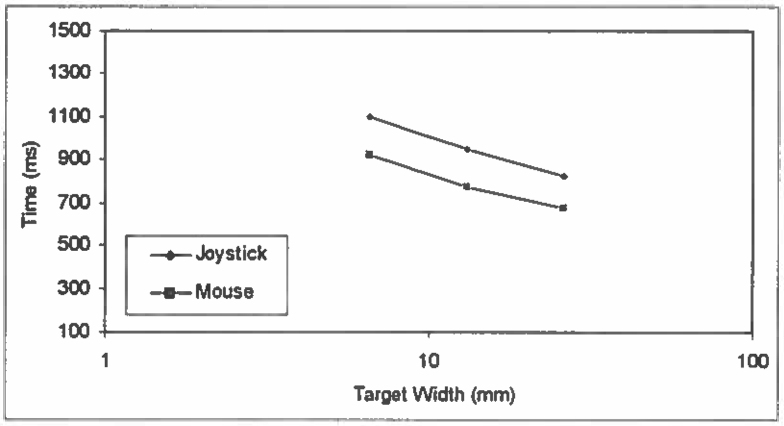

Figure 59. Effect of target width on trial time, n = 5.

Figure 58 shows that the trial completion time increases with the log of the distance to the target, while Figure 59 shows that the trial completion time decreases with the log of the target width. There is a form of the Fitts' law equation developed by Kerr (1978) that looks like:

$$MT = a + b_d \log_2 D - b_w \log_2 W \qquad (24)$$

If we assume that the tremor makes it hard to stop in a target, then there should be a larger effect due to the target width. That is to say that the smaller targets should take a disproportionately long time to reach. Figure 59 however shows that this is not the case, and the two lines have very similar slopes.

On the other hand, Figure 58 compares the effect of distance. The plots there show a curvature, with an increasing slope for targets that are far away. In addition, the slope for the joystick is steeper than the slope for the mouse.

There are a number of things that can account for these phenomena. First, the experiment was set up to measure only three widths. Based on the earlier experiment we conducted with an isometric joystick (Douglas & Mithal, 1993), we decided not to use the smallest target size (8 pixels), and the smallest width used was 16 pixels. If we had used more target widths, or if we had used a target width of 8 pixels, then it is possible that the plot of movement time versus distance would have shown a greater effect for smaller targets.

Second, the curvature evident in the velocity curve might be a consequence of the form of the Fitts' law equation being used. Figure 58 assumes that there is a logarithmic relationship between movement time and target distance. On the other hand, Kvalseth suggested a relationship of the form:

$$MT = aA^bW^c \tag{25}$$

(Kvalseth, 1981) which takes the form

$$MT = a(A/W)^b \tag{26}$$

when b = -c. This is the form of the equation that is suggested by the SOS model. If such an approach is taken, it might flatten the curve in Figure 58. In fact, MacKenzie's work (MacKenzie, 1992) has focused on forms of the Fitts' law equation that would lead to straighter lines when plotting movement time versus distance.

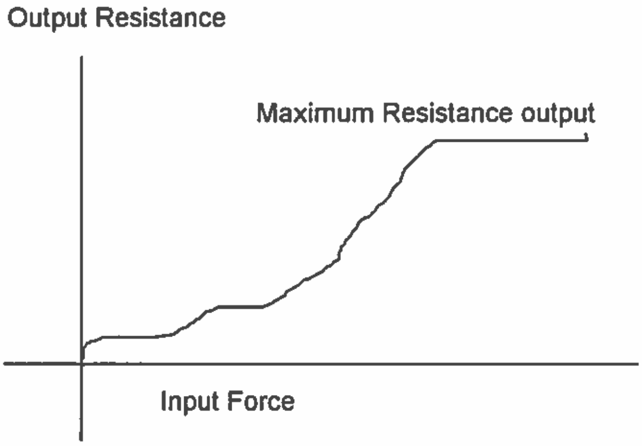

Figure 60. Hypothetical mapping of input force into output resistance on the transducers in the isometric joystick.

Finally, the manner in which the isometric joystick maps input force onto movement velocity could play a role in making the targets that are further away more difficult to reach. There are two aspects to this, both of which involve limiting the maximum velocity that can be achieved. First, the transfer function that is implemented in the joystick's hardware controller might have a peak velocity. In this approach, the cursor is prevented from moving at a velocity greater than this peak velocity, a value that could either be hard coded, or that can be varied by the user.

The second way in which the peak velocity can be limited is in the manner in which the transducers map input force into resistance. For example Figure 60 shows a hypothetical transducer that maps input force into output resistance. It has an initial range which has a somewhat linear mapping of force into resistance, and then reaches a peak after which regardless of the amount of force applied, the resistance does not change. If this is true for the transducers in the joystick, then the velocity output from the joystick will reach a peak beyond which it will not increase.

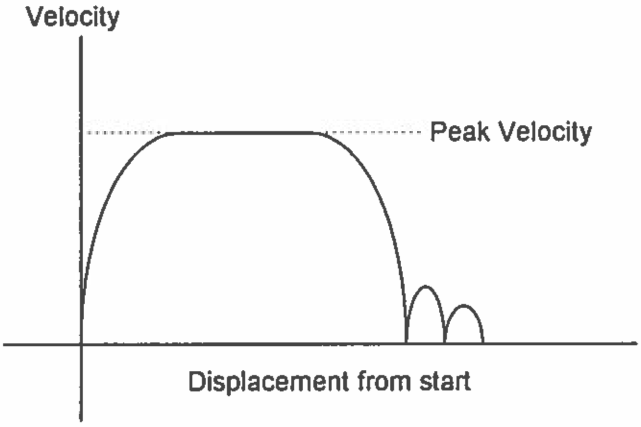

Figure 61. Depiction of velocity vs. displacement curves under the effect of a peak velocity

For large movements, assuming that the SOS model holds, this will result in a velocity-displacement curve that looks like Figure 61. Figure 61 shows an initial submovement that attempts to reach a velocity that is larger than the peak velocity allowed by the joystick subsystem, and therefore the velocity gets clipped at the peak velocity. This could also account for the observation the MT vs. log distance plots are curved, and that more distant targets take longer to reach than more narrow targets.

Frequency Domain Analysis of Movement

The velocity vs. distance plots provided a qualitative understanding of tremor and its affect on the movement microstructure. This could have been analyzed quantitatively by a quantitative analysis of the movement microstructure, i.e. parsing the data into submovements and studying their properties. The tremor made this infeasible. So, in order to develop a more quantitative analysis, the data was analyzed in the frequency domain.

As discussed earlier, comparing the plots of velocity vs. distance for the last block of mouse trials to the last block of joystick trials suggested that there was a greater portion of high-frequency components in the joystick trials.

The spectrogram of the last blocks of mouse and joystick trials in Figure 47, shows that there are in fact more high frequency components in the joystick trails than there are in the mouse trials which is consistent with the presence of tremor in the joystick trials. This suggests the possibility that a low pass filter with gentle roll-off could reduce the effect of tremor, and this idea is discussed further in Chapter VII.

Figure 48 shows that the relationship between the power and the frequency is somewhat linear in the log-log domain for both devices, suggesting a polynomial roll-off of power as frequencies increase. This finding is not accounted for by the SOS model.

Thus the frequency analysis confirmed the presence of tremor in the joystick. The characteristics of the spectrograms suggests a regularity that should be accounted for in psychomotor models of movement.

Summary

The lessons learnt from this experiment fall into two categories, those addressed by the research questions, and those that fell outside the scope of the research questions. Not much progress was made in the attempt to extend psychomotor models of movement to isometric devices because the analytical technique that was employed was adversely affected by tremor, which was a factor that had not been previously considered. Outside the scope of the research questions, we have learnt a great deal about what happens when people use isometric devices.

CHAPTER VI

CONCLUSIONS, AND FURTHER RESEARCH

In order to understand the underlying characteristics that make human performance using pointing devices different, this dissertation studied the microstructure of movement of two very different pointing devices. Differences in the microstructure were used to explain differences in the performance of the two devices. The microstructure level analysis was also used to see how well a psychomotor model of movement applied to both devices. This final chapter discusses the conclusions reached from this process, the importance of this work on the study of pointing devices and the implications for further research.

Conclusions

The study leads to a number of conclusions: about the effect of tremor on controlling pointing devices, about the interaction between tremor and the SOS model analysis, and about the frequency characteristics of pointing trials.

Tremor Makes Isometric Devices Difficult to Control

During an earlier study with a finger-controlled isometric joystick (Douglas & Mithal, 1993), subjects were interviewed to get qualitative feedback from them about their experience. A surprisingly large number of them said that they had a lot of trouble getting the joystick to stop in small targets (the smallest target was 8 pixels across).

The reason for this is now easy to understand. Tremor causes involuntary changes in the velocity at which the cursor moves. This makes it difficult for users to achieve fine control over the cursor making it difficult to stop the cursor at a desired point on the screen and explains why isometric joysticks are hard to control.

For a moment, let us assume that the SOS model is applicable to isometric movement, and that a large, fast ballistic submovement is followed by smaller corrective submovements. Tremor is probably less of an issue during the ballistic submovement where it is a small percentage of the total force being applied. In the second phase of movement, as the subject is making small corrective submovements, it is entirely possible that the voluntary corrections that the subject makes are the same or comparable order (in terms of force and duration) as the tremor. This would explain the difficulty that subjects felt while trying to stop in the smaller targets.

Two other hypotheses that explain the poor performance of the isometric joystick were the lack of kinesthetic feedback, and the lack of stimulus response-compatibility in using velocity control to control displacement. This study did not address these hypotheses, and these might be additional reasons why isometric devices have poorer performance than isotonic devices.

Verifying the SOS Model in the Presence of Tremor

It was not possible to verify the SOS model for the isometric joystick because tremor obscured the underlying movement. As a result, subject's trials could not be parsed into their component submovements.

The SOS model tries to establish what subjects are doing when they point, and this knowledge is very valuable. Establishing this by measuring submovements worked for the mouse, but other techniques such as monitoring EEG patterns or potentials of muscle activations might be better suited in the more general case where factors such as tremor play a role. If the user's intentions can be directly measured, then it might be possible to develop a pointing device that uses this measurement to control the device.

Alternative Model for Isometric Movement

A number of times during the course of the discussion in the last chapter and in this one, analyses have been made assuming that the SOS model holds for the isometric joystick. This is, however, only an assumption, and even the authors of the SOS model disagree on whether or not it might be applicable to isometric movement.

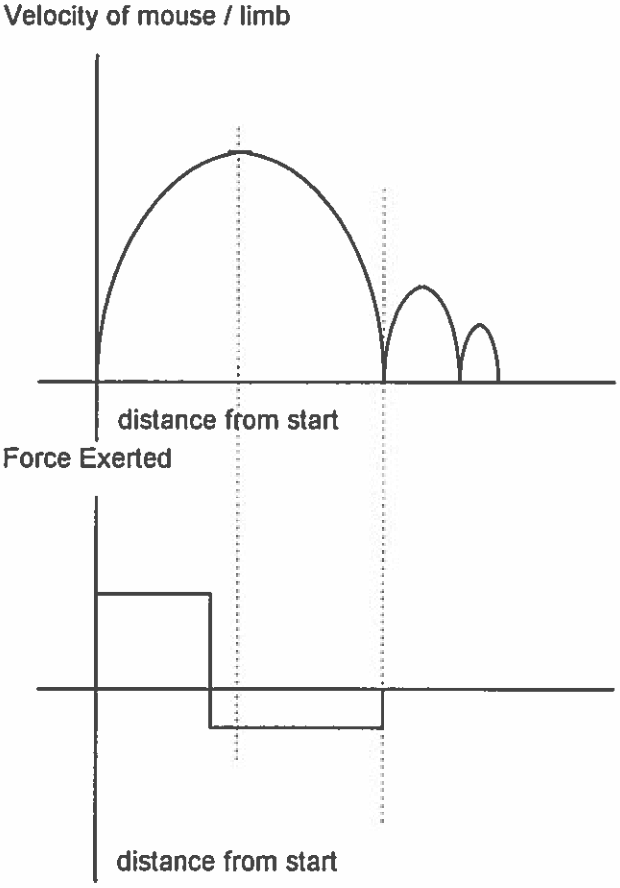

Figure 62. Force pulses involved in SOS model submovements.

In order explore this further, it is useful to understand what the user is doing from the point of view of the forces exerted to control the device. For mouse movement and for manual movement, a single sub-movement is made up of two force pulses, one that moves the mouse or limb in the direction of the target, and the second that causes a deceleration. This is depicted in Figure 62. This and the shape of the velocity vs. distance plots for the isometric joystick suggest that users of the isometric joystick might use a different movement strategy than that used to control a mouse. Note that the force applied to the isometric joystick is mapped directly onto the velocity of the cursor. In isotonic movement such as with the mouse, the cursor velocity is an integration of the force, and so force only indirectly affects velocity, while it directly affects acceleration.

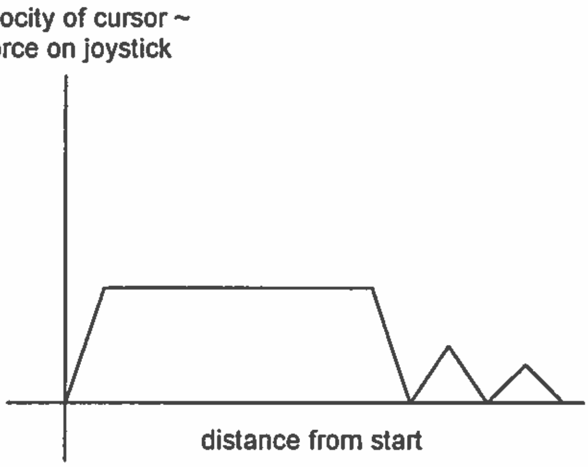

Figure 63. Model of isometric movement.

In the isometric joystick, subjects might push down on the joystick and try to reach a high velocity, and hold it at that peak till they reach close to the target. At this point, they would release the joystick and make smaller adjusting movements. This is depicted in Figure 63. This model is also supported by the finding that targets that were further away took disproportionately longer for the isometric joystick.

Frequency Characteristics and Implications for the SOS Model

The spectral analysis of the trials from the last block showed that for both joystick and mouse movement the relative power of the frequencies in the trials decreases polynomially as the frequencies increase. This method used to reach this conclusion was stochastic in nature, and averaged the frequency components of a large number of trials together. This approach is similar to that used in the SOS model, which is also stochastic in nature, and looks at the average characteristics of a large number of trials.

The SOS model does not make any predictions about movement characteristics in the frequency domain. This omission might be significant in the light that the exponential decay in frequency components holds for two widely differing devices, the isometric joystick and the mouse. Extending the SOS model analysis to explain this frequency characteristic is beyond my expertise, but it needs to be done, even if the explanation turns out to be trivial.

Relevance to Pointing Device Research and Design

This study was different from earlier work in pointing devices because its goal was not to compare the performance of pointing devices. Rather its goal was to understand why differences in performance exist. The tool used to gain this understanding was the microstructure level analysis, and this process has provided valuable lessons on the value of the microstructure level analysis, the need to examine the variable being controlled, and the need to consider individual differences.

The Microstructure Level Analysis as a Research Tool

Earlier research had shown the value of the microstructure level analysis in modeling isotonic devices (Jagacinski et al., 1980a; Walker et al., 1993). While it could not do the same for isometric devices, it was extremely effective in providing information about what happened during joystick trials. The next step development research in pointing devices should be to study movement microstructure. The manner in which this study uncovered tremor in the isometric devices strongly supports this suggestion.

In the absence of models to explain movement, an analysis of the microstructure of movement is best available tool. This sort of analysis does not shorten the design process, but it will provide the designer with a much better idea of the kinds of changes to make to improve a device.

We need to know how a device will perform in a novel environment, and this is becoming more important as we start to use virtual environments where people navigate in virtual spaces. For example, a recent presentation in the Doctoral Consortium at the CHI '95 conference described a virtual environment where subjects use a modal 3-D pointing device to navigate in a 3-dimensional space with multiple degrees of freedom (Boyd, 1995). One of the problems faced by the designer was to determine what kind of pointing device was appropriate for the situation, and how its input signals should be mapped onto the control parameters.

The availability of a corpus of knowledge about the movement characteristics of different pointing devices would have helped the researcher pick the most appropriate mapping, and tailor the interaction to make that device work better in that environment. We need inventories of knowledge of microstructure of movement with different devices just like we have a corpus of knowledge about gross movement and the applicability of Fitts' law.

The Need to Examine the Variable Being Controlled

A simple lesson from this study is that when a pointing device is being studied, we need to examine its effect on the variable it controls. The isometric joystick controls the velocity of the cursor (even though it integrates the velocity over time and sends the computer displacement information). The effect of tremor was visible only when velocity was plotted against time or displacement.

Taking this argument to the next stage, the implication is that we should look at graphs of the appropriate variable. For example, if we control acceleration, we should look at graphs of acceleration.

Individual Differences

Contemporary research on pointing devices has not taken individual differences into consideration. Standard Fitts' law analyses for example, compare group means and derive Fitts' law coefficients from entire samples. The results of this study showed that there might be a qualitative difference between subjects who perform well on the isometric joystick (smaller pointing times), and subjects who perform poorly. This difference was apparent in two ways. First, subjects with higher performance tended to have smaller amounts of tremor in their movements when the velocity vs. distance graphs were examined. Second, when the spectrograms of joystick movement were compared with the spectrograms of mouse movement, there was less difference between the two for subjects with higher performance than was apparent for the rest of the subjects.

This suggests that the isometric joystick might be more suitable for some users than for others. Motor control is considered to be made up of components, three important ones being timing, force and the sequence in which movements are made. These are obviously of importance to pointing device movement. A study of "clumsy" children has linked soft neurological signs with variability in the controlling of isometric force pulses (Lundy-Ekman, Ivry, Keele, & Woollacott, 1991). Soft neurological signs are very mild disorders of the nervous system, so minor that they are in fact difficult to categorize, and in no way make the individual disabled. The study describes children with such signs as "clumsy." In simpler terms, it indicates that clumsy children with soft neurological signs in the basal ganglia have difficulty in controlling the amplitude of isometric force pulses.

In the terms of the present study, we could say that some subjects might have a lot of tremor in their movements and there might be strong physiological reasons why some people are better at using the isometric joystick than others. This is a line of research that is well worth investigating.

Such an investigation would start by gathering both movement data and neurological signs data from a large group of subjects, perhaps 30, and would then compare the characteristics of the trials and the neurological signs of the top 10 subjects against those of the bottom 10 to see what differences emerge.

If there is a link between individual differences and performance, it might be possible to devise a test. Users would walk up to a test machine, perform a series of pointing actions with one or more pointing devices. They would then be given a recommendation of the kind of pointing device best suited to that person. So, clumsy people might be told to avoid isometric devices.

The point is that pointing device research has ignored individual differences. Over the years the electronic newsgroup comp.human-factors has provided ample evidence that individuals have strong, almost religious likes and dislikes when it comes to pointing devices. The arguments that rage on that newsgroup about whether trackballs are better than mice are possibly second only to arguments about whether PCs are better than Macs.

By ignoring these device preferences, and comparing the performance of different pointing devices across the board, we do users a disservice. We need to be able to match pointing devices to users. For instance, in this study, subject DD6 had a poorer pointing time with the mouse than with the isometric joystick. It would be nice to be able to take such individual differences into account and help select the most appropriate pointing device for the individual.

Fate of the Isometric Joystick

One of the impressions arrived at in years of studying the isometric joystick is that while it has a number of benefits, such as small space requirements and low construction costs, there are problems associated with its use. We need to ask the question "Is it worth it?" about continuing research and development on improving the isometric joystick. Depending on the point of view there are two answers. From the point of view of research, continuing to study the isometric joystick is of great use because its properties provide an excellent counterpoint to isotonic devices in the search for understanding how humans point. From a more practical standpoint of manufacturing a pointing device with superior user comfort and performance, the answer is less clear.

There are other pointing devices available that have many of the properties of the isometric joystick and might not have its drawbacks. One such device is the relative-mode touch sensitive pad that has begun to appear in notebook computers by Apple, Sharp, Epson and others. Apple calls their product a TrackPad, and the offering by Sharp and Epson is called the Glidepoint. Other research has shown the relative mode touch tablet to be a very good pointing device though they did not test the TrackPad itself (MacKenzie et al., 1991). A recent study of the TrackPad from the point of view of postural comfort concluded that extensive use of the TrackPad did not cause discomfort, but did not measure its performance (Cakir, Cakir, Muller, & Unema, 1995).

Tests comparing the isometric joystick and trackpad need to be performed to get qualitative reactions of users as well as quantitative values of performance times and error rates. In such a test, both pointing devices would be controlled by the index finger, providing a more accurate way of comparing differences in isotonic and isometric movement.

Some groups have begun to consider the TrackPad good enough to replace the mouse. Windows Magazine recently wrote:

> "HOT TOUCH: Don't get too used to that mouse at your fingertips; within the next three years, touchpads will be all the rage. According to 'The Optimal Pointing Strategy for Mobile Computing," a new report from BIS Strategic Decisions, touchpads will command 70 percent of the market by 1998, leaving the rest to mice, trackballs and joysticks. Apple recently became the first large vendor to incorporate a trackpad in its products, namely the 520 and 540 PowerBooks." -- Windows Magazine, April 1995, vol. 6 #4, page 42.

The final determinant of market success is user acceptance, and if the trackpad feels better than the isometric joystick, it is likely to do better in the marketplace.

Opportunities for Further Research

This study raises a number of possibilities for further research. The ones that have been mentioned already are linking the SOS model to more direct measures like EEGs, explaining the polynomial relationship between power and frequency in the mouse and joystick spectra, establishing a relationship between clumsiness and joystick performance, and doing a microstructure level comparison of joystick and trackpad movement.

Other opportunities for research center around improved experimental methodology and improved designs for the isometric joystick.

Improving the Experiment

One important conclusion from this study is about the methodology. As mentioned in the literature review, a number of studies use the computer as both the measuring device as well as the experimental device. The computer is therefore playing two roles, presenting the experimental task, and measuring itself. This leads to errors in measurement. In cases such as standard Fitts' law studies of gross movement time, this is not a problem. The slight error in measurement is negligible compared to the magnitude of the times being measured and averaging together the values from a large number of trials reduces the effect of the error.

However, when the size of the measurement becomes small, and when the measurements are not averaged together error plays a bigger role. This is what happens when we analyze the microstructure of each trial. In such a situation we need to be able to make external measurements, perhaps using one device to determine cursor position, and another device to measure time. This is the approach used by Jagacinski et. al (1980) in their study, and is a superior approach to the one employed in the present study.

Improving the Finger-Controlled Isometric Joystick

At some level, the overall goal of research and model building is design. Finding ways to improve the design of the joystick was a goal of this study. It uncovered problems with finger controlled isometric devices, and a natural way of proceeding is for the design to directly address these problems. Each of these approaches needs to be evaluated to see how effective it is, and its effect on the movement microstructure.

Filtering out the tremor

The most obvious problem is the tremor, and the difference between the spectrograms for mouse trials and joystick trials raises the possibility that a filter with a gentle roll-off might eliminate or at least reduce some of the tremor. In any case, such a filter can be implemented either in the computer, or in the pointing device.

If implemented in the computer, it can use a software only solution, and use real-time digital signal processing (DSP) routines to filter out the tremor. It can also use on-board DSP chips such as those becoming increasingly common on sound boards which would speed up the software. A pointing-device only solution could take a similar approach and put a DSP chip into the device. A different approach would be to use analog filtering using capacitors and inductors which would have the advantage of being relatively easy to implement.

While the details of implementation lie outside my area of expertise, it seems clear that all four solutions can be presented to the user along with customization and training software. The software would present simple pointing tasks to the user, gather data from a series of trials and automatically set parameters on the filtering routines to match the user's individual movement characteristics. As the user improved with practice, she could manually run the customization routines. Alternately the driver could constantly run in the background analyzing the user's movements and resetting the filter's parameters.

One issue that will need to be considered in these approaches to reduce the effect of tremor is that of feel. Feel, or the responsiveness of the system to user input is subjectively important to the user (Barrett, Selker, Rutledge, & Olyha, 1995; Rutledge & Selker, 1990). Such approaches use "a transfer function which amplifies changes in control action" (Barrett et. al., 1995 page 316). Such an approach is likely to increase the effect of tremor, and a balance will have to be reached between the need to filter the tremor and the need to improve the feel of the device.

Minimizing the Tremor

The filtering approach tries to minimize the effect of tremor once it is introduced into the system. A different approach would be to minimize the tremor itself. One cause of tremor is the fact that the finger is unsupported. The finger could be supported say under the third metacarpal bone which is shown in Figure 64. This approach might reduce the tremor from the rest of the hand. It might also have the added benefit of reducing fatigue among the users. A number of subjects complained after the study was completed that they had found the joystick sessions tiring, including the two subjects who had their own IBM ThinkPads with built in TrackPoints.

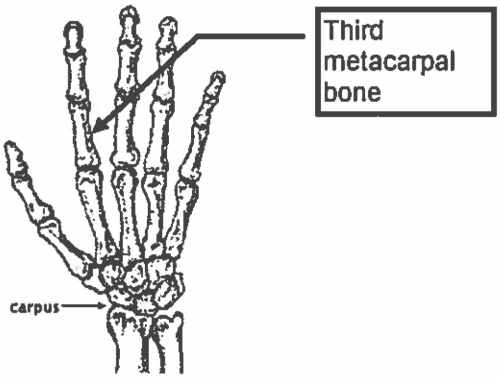

Figure 64. Possible location of a support for the finger controlling the isometric joystick

A second source of tremor is the fact that there is no friction associated with movement in isometric devices. An intriguing possibility is to re-engineer friction into the transfer function of the joystick controller.

A third source of tremor is the isometric nature of the device. Making the device slightly deformable might have the effect of associating an inertial mass which can damp the effects of tremor. A deformable joystick would have the added advantage that it would provide the user with some kinesthetic feedback, which gets lost in an isometric joystick.

Training Users to Use the Isometric Joystick

It was discussed earlier how some subjects in the experiment were able to make appropriate movements with the joystick a large percentage of the time, and that most users can make appropriate movements at least some of the time. Appropriate movements were movements which had episodes of high velocity, relatively unaffected by tremor. This raises the possibility that the appropriate movements might be something that can be taught to users. Users can be presented with a training program that gives them positive feedback (e.g. a happy sound) when they are able to make a movement with small amounts of tremor.

In summary, a number of research opportunities arise out of this study in the areas of model building, developing relationships between physiological characteristics and ability with the isometric joystick and improving the isometric joystick. Further work in any of these areas will significantly contribute to the knowledge about pointing and pointing devices.

Summary

Pointing devices are likely to increase in importance as new technologies such as virtual reality, interactive television and new forms of collaborative communication become more pervasive. These technologies will require, and are already requiring, new forms of interaction, new forms of communication with the computer. Pointing devices will provide a fertile research area for years to come.

Research on pointing devices will feed off research on the human psychomotor system, and is likely provide an impetus for new discoveries about how humans move. In this effort, Fitts' law must be replaced by microstructure level analysis as the primary research and design tool for pointing devices (MacKenzie, 1992). This dissertation is the first step in that direction.

BIBLIOGRAPHY

Abrams, R. A., Meyer, D. E., & Kornblum, S. (1990). Eye-Hand Coordination: Oculomotor Control in Rapid Aimed Limb Movements. Journal of Experimental Psychology: Human Perception and Performance, 16(2), 248-267.

Andres, R. O., & Hartung, K. J. (1989). Prediction of Head Movement Time Using Fitts' Law. Human Factors, 31(6), 703-713.

Barrett, R. C., Selker, E. J., Rutledge, J. D., & Olyha, R. S. (1995). Negative Inertia: A Dynamic Pointing Function. In Human Factors in Computing Systems CHI '95 Conference Companion (pp. 316-317). New York: ACM.

Bentley, T. L., Custer, G. M., & Meyer, G. M. (1975). An Alternative Editing Console for the Palo Alto Research Center, Xerox Corporation. Unpublished project report). The Design Division, Stanford University.

Boritz, J., Booth, K. S., & Cowan, W. B. (1991). Fitts' Law Studies of Directional Mouse Movement. In Graphics Interface '91 (pp. 216-233). Toronto, Ontario, Canada: Canadian Man-Computer Communications Society.

Boyd, C. (1995). Human and Machine Dimensions of 3D Interfaces for Virtual Environments. In Human Factors in Computing Systems CHI '95 Conference Companion (pp. 41-42). New York: ACM.

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

Buxton, W., & Myers, B. A. (1986). A Study in Two-Handed Input. In Human Factors in Computing Systems: CHI '86 (pp. 321-326). New York: ACM.

Cakir, A. E., Cakir, G., Muller, T., & Unema, P. (1995). The TrackPad - A Study on User Comfort and Performance. In Human Factors in Computing Systems CHI '95 Conference Companion (pp. 97-98). New York: ACM.

Card, S. K., English, W. K., & 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., Mackinlay, J. D., & Robertson, G. G. (1990). The Design Space of Input Devices. In Human Factors in Computing Systems CHI ‘90 (pp. 117-124).

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

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

Crossman, E. R. F. W., & Goodeve, P. J. (1983). Feedback control of hand-movement and Fitts' Law. Quarterly Journal of Experimental Psychology, 35A, 251-278.

Douglas, S. A., & Mithal, A. K. (1993). Human Performance Evaluation of a Finger-Controlled Pointing Device. CIS TR 92-11. Department of Computer Science, University of Oregon, Eugene, 97403-1202.

Douglas, S. A., & Mithal, A. K. (1994). The Effect of Reducing Homing Time on the Speed of a Finger-Controlled Isometric Pointing Device. In Human Factors in Computing Systems CHI '94 (pp. 411-416). New York: ACM.

Drury, C. G. (1975). Application of Fitts' law to foot-pedal design. Human Factors, 17, 368-373.

Epps, B. W. (1986). Comparison of Six Cursor Control Devices Based on Fitts' Law Models. Proceedings of the Annual Meeting of the 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(6), 381-391.

Fitts, P. M., & Peterson, J. R. (1964). Information capacity of discrete motor responses. Journal of Experimental Psychology, 67(2), 103-112.

Gillan, D. J., Holden, K., Adam, S., Rudisill, M., & Magee, L. (1990). How Does Fitts' Law Fit Pointing and Dragging. In Human Factors in Computing Systems CHI '90 (pp. 227-234). New York: ACM.

Gillan, D. J., Holden, K., Adam, S., Rudisill, M., & Magee, L. (1992). How should Fitts' Law be applied to human-computer interaction? Interacting with Computers, 4(3), 291-313.

Goodwin, N. C. (1975). Cursor Positioning on an Electronic Display using Lightpen, Lightgun or Keyboard for Three Basic Tasks. Human Factors, 17(3), 289-295.

Hoffmann, E. R. (1991). A comparison of hand and foot movement times. Ergonomics, 34(4), 398-406.

Jagacinski, R. J. (1989). Target Acquisition: Performance Measures, Process Models, and Design Implications. New York: Plenum.

Jagacinski, R. J., Hartzell, E. J., Ward, S., & Bishop, K. (1978). Fitts' Law as a function of system dynamics and target uncertainty. Journal of Motor Behavior, 10(2), 123-131.

Jagacinski, R. J., & Monk, D. L. (1985). Fitts' Law in Two Dimensions with Hand and Head Movements. Journal of Motor Behavior, 17(1), 77-95.

Jagacinski, R. J., Repperger, D. W., Moran, M. S., Ward, S. L., & Glass, B. (1980a). Fitts' Law and the Microstructure of Rapid Discrete Movements. Journal of Experimental Psychology: Human Perception and Performance, 6(2), 309-320.

Jagacinski, R. J., Repperger, D. W., Ward, S. L., & Moran, M. S. (1980b). A Test of Fitts' Law with Moving Targets. Human Factors, 22(2), 225-233.

Jellinek, H. D., & Card, S. K. (1990). Powermice and User Performance. In J. C. Chew & J. Whiteside (Ed.), Human Factors in Computing Systems CHI '90 (pp. 213-220). New York: ACM.

Kaiser, J. F., & Reed, W. A. (1977). Data smoothing using low-pass digital filters. Review of Scientific Instrumentation, 48(11), 1447-1456.

Kantowitz, B. H., & Elvers, G. G. (1988). Fitts' Law with an Isometric Controller: Effects of Order of Control and Control-Display Gain. Journal of Motor Behavior, 20(1), 53-56.

Kantowitz, B. H., & Knight, J. L. (1976). Testing Tapping Time-Sharing: Attention Demands of Movement Amplitude and Target Width. New York: Academic Press.

Keele, S. W. (1968). Movement control in skilled motor performance. Psychological Bulletin, 70(6), 387-403.

Keele, S. W. (1986). Motor Control. In J. K. Boff, L. Kaufman, & J. P. Thomas (Eds.), Handbook of Human Perception and Performance, Vol. II (pp. 30:1-60). New York: Wiley.

Kerr, R. (1978). Diving, Adaptation, and Fitts' Law. Journal of Motor Behavior, 10(3), 255-260.

Kvalseth, T. O. (1981). An experimental paradigm for analyzing human information processing during motor control tasks. Proceedings of the Human Factors Society, 25, 581-585.

Langolf, G. D., Chaffin, D. B., & Foulke, J. A. (1976). An investigation of Fitts' Law using a wide range of movement amplitudes. Journal of Motor Behavior, 8(2), 113-128.

Lin, M. L., Radwin, R. G., & Vanderheiden, G. C. (1992). Gain Effects on performance using a head-controlled computer input device. Ergonomics, 35(2), 159-175.

Lundy-Ekman, L., Ivry, R., Keele, S., & Woollacott, M. (1991). Timing and Force Control Deficits in Clumsy Children. Journal of Cognitive Neuroscience, 3(1), 367-376.

MacKenzie, I. S. (1991). Letter to the author.

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., & Buxton, W. (1992). Extending Fitts' Law to Two-Dimensional Tasks. In Human Factors in Computing Systems: CHI '92 (pp. 219-226). New York: ACM.

MacKenzie, I. S., Sellen, A., & Buxton, W. (1991). A Comparison of Input Devices in Elemental Pointing and Dragging Tasks. In Human Factors in Computing Systems: CHI '91 (pp. 161-166). New York: ACM.

Mackinlay, J., Card, S. K., & Robertson, G. C. (1990). A Semantic Analysis of the Design Space of Input Devices. Human-Computer Interaction, 5(2), 145-190.

Meyer, D. E. (1988). Optimality in Human Motor Performance: Ideal Control of Rapid Aimed Movements. Psychological Review, 95(3), 340-370.

Meyer, D. E., Smith, J. E. K., & Wright, C. E. (1982). Models for the speed and accuracy of aimed movements. Psychological Review, 89(5), 449-482.

Meyer, D. E., Smith, K. J. E., Kornblum, S., Abrams, R. A., & Wright, C. E. (1990). Speed-Accuracy Tradeoffs in Aimed Movements: Toward a Theory of Rapid Voluntary Action. In M. Jeanerod (Eds.), Attention and Performance XIII (pp. 173-226). Hillsdale, New Jersey: Erlbaum.

Pearson, G., & Weiser, M. (1986). Of Moles and Men: The Design of Foot Controls for Workstations. In Human Factors in Computing Systems CHI '86 (pp. 333-339). New York: ACM.

Pearson, G., & Weiser, M. (1988). Exploratory Evaluation of a Planar Foot-Operated Cursor-Positioning Device. In E. Soloway, D. Frye, & S. B. Sheppard (Ed.), Human Factors in Computing Systems CHI ' 88 (pp. 13-18). New York: ACM.

Rosenberg, D. J., & Martin, G. (1988). Human-Performance Evaluation of Digitizer Pucks for Computer Input of Spatial Information. Human Factors, 30(2), 231-235.

Rutledge, J. D., & Selker, T. (1990). Force-to-Motion Functions for Pointing. In Human-Computer Interaction - INTERACT '90 (pp. 701 706). Amsterdam: North-Holland.

Schmidt, R. A. (1979). Motor-output variability: A theory for the accuracy of rapid motor acts. Psychological Review, 86(5), 415-451.

Schmidt, R. A., Zelaznik, H. N., & Frank, J. S. (1978). Sources of Inaccuracy in Rapid Movement. In G. Stelmach (Eds.), Information Processing in Motor Control and Learning. New York: Academic Press.

Stein, R. B., & Lee, R. G. (1989). Tremor and Clonus. In Handbook of Physiology: The Nervous System II (pp. 325-342).

Walker, N., Meyer, D. E., & Smelcer, J. B. (1993). Spatial and Temporal Characteristics of Rapid Cursor Positioning Movements with Electromechanical Mice in Human Computer Interaction. Human Factors, 35(3), 431-458.

Ware, C., & Mikaelian, H., H. (1987). An Evaluation of an Eye Tracker as a Device for Computer Input. In J. M. Carroll & P. P. Tanner (Ed.), Human Factors in Computing Systems and Graphics Interface: CHI+GI 1987 (pp. 183-188). New York: ACM.

Welford, A. T. (1968). Fundamentals of Skill. London: Methuen.