EFFECT OF VARIABLE FEEDBACK DELAY ON VISUAL TARGET-ACQUISITION PERFORMANCE

A thesis submitted in partial fulfillment

of the requirements for the degree of

Master of Science

By

JULIO CHRISTIAN MATEO
B.A., Universidad Pontificia de Salamanca, 2003

2007
Wright State University

WRIGHT STATE UNIVERSITY

SCHOOL OF GRADUATE STUDIES

April 6, 2007

I HEREBY RECOMMEND THAT THE THESIS PREPARED UNDER MY SUPERVISION BY Julio Christian Mateo ENTITLED Effect of Variable Feedback Delay on Visual Target-Acquisition Performance BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF Master of Science.

Robert H. Gilkey, Ph.D.
Thesis Director

John M. Flach, Ph.D.
Department Chair

Committee on Final Examination

Robert H. Gilkey, Ph.D.

John M. Flach, Ph.D.

David LaHuis, Ph.D.

Joseph F. Thomas, Jr., Ph.D.
Dean, School of Graduate Studies

ABSTRACT

Mateo, Julio Christian. M.S., Department of Psychology, Wright State University, 2007. Effect of Variable Feedback Delay on Visual Target-Acquisition Performance.

Traditionally, private communication channels with stable characteristics have been used in teleoperation situations. However, recently there have been a few attempts at using public communication channels such as the Internet. In spite of their convenience, very little is known about the effect of the variable delays inherent in this type of channel on motor performance. In this thesis, we provide empirical data on the impact of variable feedback delays on a 3D visual target-acquisition task performed in a virtual environment. Target size, distance between targets, mean feedback delay, and feedback-delay variability were manipulated and the number of errors and movement time (MT) were measured. Results showed that feedback-delay variability affected the closed-loop part of visual target-acquisition movements, even though its effect was weaker than the effect of mean feedback delay. Our results advise against using techniques that reduce feedback-delay variability at the expense of increasing mean feedback delay. In addition, we found that target size was critical for visual target-acquisition performance in the presence of feedback delays and this should be considered when designing teleoperation situations. Issues associated with studying feedback-delay variability are identified and lines of future research are suggested.

TABLE OF CONTENTS

TABLE OF CONTENTS (CONTINUED)

TABLE OF CONTENTS (CONTINUED)

LIST OF FIGURES

LIST OF FIGURES (CONTINUED)

ACKNOWLEDGEMENTS

In spite of the fact that no other authors are listed on the front page of this thesis, the work presented here is not the result of the work of an individual but, rather, a collective effort. In this section, I attempt to recognize some of the other people without whom this work would have never seen the light.

First and foremost, I want to thank my advisor: Dr. Robert H. Gilkey. In the last 3 years, Gilkey has become not only my academic mentor, but also an extraordinary friend. I do not have here the space or the time to express my gratitude toward him, but I hope he realizes how important he has become in my (both academic and personal) life. His insightful comments, his creativity, and his problem-solving skills have been critical for the completion of this project.

This work would have never been possible with the help of Jeff Cowgill, Jr., our systems engineer. Jeff not only implemented the virtual environment used for this thesis, but his suggestions and technological support were crucial throughout the process.

I am also grateful to all the people at the Battlespace Acoustics Branch of the Air Force Research Laboratory that helped me during the data collection. Especially, I would like to thank Brian Simpson for arranging the participants for my thesis and for his continuous support throughout my graduate-school work (including, but not limited to, this thesis), Mike Ward and Mary Ann Barbato for their assistance during the long data-collection process, and of course, the participants of my study for their hard work.

I would also like to thank Joe Manning, our undergraduate research assistant, for his assistance during the late stages of the thesis design and all the data collection; Dr. Thomas Moore, for his always insightful comments; my committee members, Dr. John

Flach and Dr. David LaHuis, for their patience and assistance during the development of this work; the staff in the Psychology Department, especially Linda Blatz, for their support and assistance during all these years; and to my friends and fellow graduate students, who helped me get through the non-academic part of these last few years.

Last, but not least, I want to thank my family for always being there for me, in spite of how rarely I let them know how important they are to me. I especially want to thank my parents, Jose Mari Mateo and Cuca Prado, for their continuous encouragement and unconditional support. I also want to let my parents, my siblings (Jose and Vanessa), and all my extended family (both from Spain and Puerto Rico) know that they are always in my mind, even when I fail to communicate it to them.

I. INTRODUCTION

Most of the activities people perform in their daily lives, such as writing, picking up a glass, or turning off a light switch, involve coordination between motor actions and the sensory feedback resulting from those actions. Among the many types of sensory feedback (e.g., proprioceptive and haptic feedback) that are important for the successful completion of motor tasks, the focus in this thesis is on *visual feedback* and visuo-motor coordination.

In everyday situations, most people perform simple motor activities easily and without realizing the important role that visual feedback plays in their successful completion. Only when visual feedback is deteriorated (e.g., low or no light) are people likely to realize that they were using it to perform their actions. During most everyday activities, people can directly see the consequences of their actions, since visual feedback is available to the person from the moment the action occurs without delay. For example, most people can directly see their own hand moving (e.g., when picking up a coffee mug), as well as the objects that the hand might contact during the movement (e.g., a glass of water next to the mug). This visual feedback allows the person to avoid spilling the glass of water when reaching for the coffee mug and, if the glass was touched during the movement, the person could use the instantaneous feedback (visual and haptic) to correct his/her action and avoid the spillage.

However, in some less common situations (e.g., on-orbit satellite repair, deep underwater exploration, or radioactive waste removal), it is impossible or too dangerous to have the person (i.e., operator) performing the action in the same location as the object on which they are performing the actions. In these cases, as illustrated in Figure 1, it is often convenient to place an effector (e.g., a robotic arm) in the remote site where the action needs to take place while the human operator controls (operates) it from a local, safer site. In these teleoperation situations, it is usually impossible to provide the human operators instantaneous with visual feedback about their actions and the consequences of these actions. Actions performed by the human operator in the local site are not executed in the remote site until some time has elapsed (i.e., *control delay*) and actions executed by the effector (e.g., robotic arm) in the remote site are not displayed to the operator in the local site until some time has elapsed (i.e., *display delay*). In the telerobotics literature, the terms *transmission delay* and *round-trip delay* have often been used to refer to the total delay between the moment the human operators act in the local site to the moment when they receive feedback from their actions (i.e., control delay plus display delay). In the case of visual feedback, round-trip delay corresponds to the time elapsed between the moment the operator moves in the local site to the moment his/her movement is depicted in the display in that same location.

In this thesis, we are not interested in the source of the delay (i.e., control vs. display delay) but rather we are interested in how delayed visual feedback experienced by operators affect their motor performance. For this reason, the term *feedback delay* is used almost exclusively to refer to the time elapsed between the moment the operator performs an action and when this action and its consequences are visually displayed to the operator. The feedback delays used in our experiment were technically display delays (i.e., the operator's actions were implemented immediately but the display of these actions was delayed).

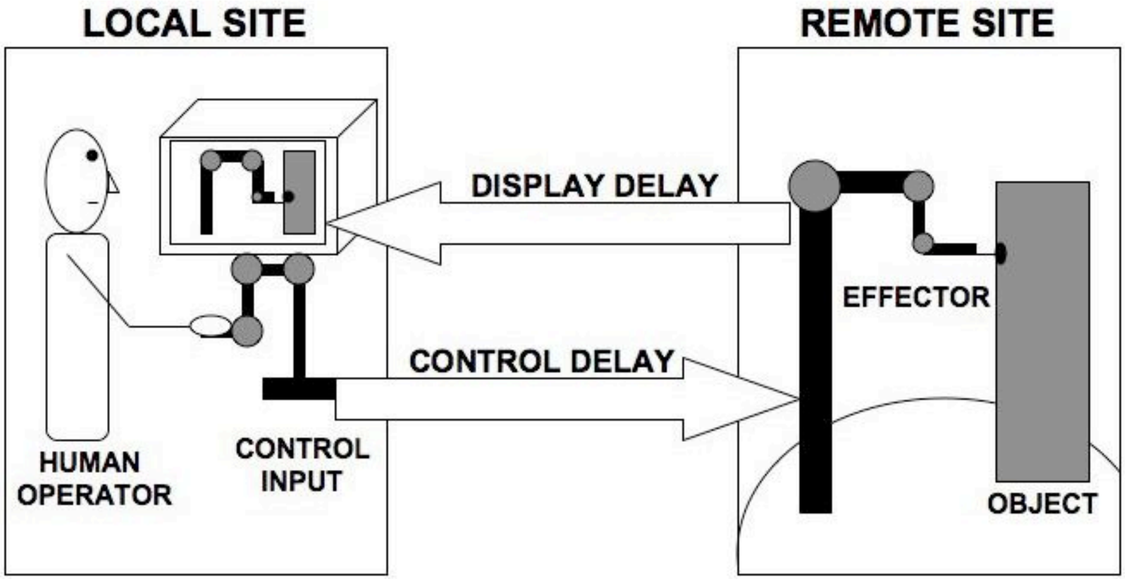

Figure 1. Illustration of a typical teleoperation situation with both control and display delays. The round-trip delay is the total delay between the moment an operator acts and when the resulting actions are displayed (i.e., control delay plus display delay). We refer to this round-trip delay in this thesis as feedback delay.

The type of tasks performed in teleoperation situations often resemble (or can be subdivided into) simple movements between two points followed by a particular action (e.g., unscrew a hatch). For this reason, researchers (e.g., Ferrell, 1965) have often used *target-acquisition* tasks to study human motor performance in teleoperation situations. In target-acquisition experiments, participants are instructed to move to a specific location (i.e., target) determined by the experimenter. Most often, participants are required to minimize the time to acquire the target (i.e., to move to the target as quickly as possible). The time taken to complete a trial (i.e., movement time, MT) and the number of errors (e.g., number of times the participant stops outside the target area) are typical dependent measures in target-acquisition studies. Given our interest in teleoperation situations and the usefulness of this type of tasks to study teleoperation performance, we used a visual target-acquisition task to investigate how feedback-delay manipulations affect performance.

Most studies exploring the effect of feedback delays in visual target-acquisition tasks have explored the impact of *constant feedback delays*. Given that most of the current teleoperation applications use private, exclusive communication channels resulting in constant delays (e.g., Hu, Yu, Tsui, & Zhou, 2001), studying constant feedback delays has been appropriate. As might be expected, researchers (e.g., MacKenzie & Ware, 1993) have consistently found a negative impact of increasing feedback delays on human motor performance in these studies (i.e., longer feedback delays resulted in longer MTs and more errors).

In the last decade, many researchers (see Goldberg & Seigwart, 2002, for a review) have suggested using public, shared channels (e.g., Internet) for some telerobotic applications. The Internet has many convenient features (e.g., its ubiquity, low cost, and easy access) when compared to private, exclusive channels. However, public channels also have an important drawback: transmission delays (and therefore, feedback delays) do not remain constant over time but vary continuously depending on many factors (e.g., number of users or routing paths). This property of the Internet results in *variable feedback delays* that, if large enough, lead to discontinuities in the movement of the effector (e.g., it may freeze and then jump to catch up with its latest location). Obviously, the representation of the effector provided to the user through the display will also show discontinuities in these cases. When the *delay variability* (or jitter) is large enough, the operator can easily notice it and his/her performance is expected to be negatively affected by this variability.

A detrimental effect of feedback-delay variability on motor performance (above and beyond the effect of mean feedback delay) seems reasonable given that the displayed visual representation of the operator's original movement is no longer a consistent representation of that original (input) movement. That is, a smooth input movement by the operator may be displayed as a sequence of cursor jumps, making it difficult for the operator to know where his cursor really is and to adapt to it. In spite of the logical argument to expect this effect, empirical research exploring the effect of delay variability (and movement discontinuity) on motor performance is sparse and very little is known about this relationship. Given the potential of Internet-based teleoperation, it is important that more research is done in order to improve our understanding of the relationship between feedback-delay variability and motor performance. Understanding this relationship is a necessary step for the development of operator models for Internet teleoperation situations and appropriate strategies to deal with the variable delays resulting from the use of the Internet as a communication channel. The main objective of this thesis is to improve our understanding of the relationship between feedback-delay variability and motor performance. For this purpose, we conducted an empirical study in which feedback-delay variability was systematically manipulated and its effects on motor performance were measured.

In the first and second sections of the literature review, the main frameworks traditionally used to model human motor performance in the presence of feedback delays are briefly described and some of the empirical studies that have explored the impact of constant delays on human performance are presented. Due to the large number of studies fitting this criterion, those studies that explored visual target-acquisition tasks and attempted to model human performance are emphasized. In the third section, the few studies that have explored the impact of variable feedback delays on human performance in individual and collaborative tasks are described. The fourth section includes a review of the strategies that human operators use to deal with constant and variable feedback delays and, in the fifth section, some of the techniques proposed by designers to ameliorate the negative effects of constant and variable feedback delays on performance are introduced. Finally, the goals and predictions of this thesis are presented.

Motor-Performance Models

The models that have been proposed to explain human performance in movement tasks can generally be categorized under one of two frameworks: *Information theory* and *Control theory*. In broad terms, models based on Information theory compare the human operator to a communication channel of limited capacity, whereas models based on Control theory compare the human operator to a feedback control system. The models account for the effect of feedback delays in slightly different ways.

Models Based on Information Theory

Since Shannon and Weaver (1949/1963) formulated the mathematical theory of communication, many psychologists have used information measures derived from Shannon and Weaver's work to characterize different aspects of behavior (see Welford, 1960, for a review). In the context of movement control, models based on Information theory have assumed that the human motor system acts as an information channel of limited capacity (e.g., Fitts, 1954). That is, a certain amount of information is transmitted every time a person performs an aimed movement, and the amount of information that the human motor system can transmit per unit of time (i.e., the channel capacity) is limited. In this context, information refers to a reduction of uncertainty, and the information transmitted by a specific movement depends not only on the movement characteristics (e.g., how far the person moves) but also on the number of possible alternative movements with the body part used to carry out the movement. That is, moving a body part that allows many possible alternative movements (e.g., an arm) transmits more information than moving a body part with less possible alternative movements (e.g., a finger).

Fitts (1954) was the first researcher to apply these Information-theory ideas to the study of the human motor system. One of the tasks used by Fitts in this initial study, and the most widely known, is the *reciprocal-tapping* task in which participants moved a stylus horizontally back and forth between two plates separated by a certain distance, stopping within the area of the target plate each time. Fitts varied the distance between plates and the size of the plates across conditions. (Fitts referred to these two variables respectively as movement amplitude, A, and target width, W, in his one-dimensional task; however, in this paper the terms *distance between targets*, D, and *target size*, S, are used instead because they are more general and apply better to three-dimensional situations.) Using an information measure, Fitts defined the *index of difficulty (ID)* of a specific movement as: $\mathrm{ID} = \log_2\left(\frac{2 * \mathrm{D}}{\mathrm{S}}\right)$. From an Information-theory standpoint, S can be interpreted as the uncertainty after the execution of the movement (i.e., how much error is tolerated in an 'accurate' movement), and $2 * \mathrm{D}$ was arbitrarily selected by Fitts to represent the uncertainty before starting the movement.

Another measure used by Fitts (1954) was the *index of performance* (*IP*), which is analogous to the capacity of a communication channel in Information theory (i.e., the higher the IP, the more information can be transmitted per unit of time). Index of performance was defined as: $\mathrm{IP} = \frac{\mathrm{ID}}{\overline{\mathrm{MT}}}$ where $\overline{\mathrm{MT}}$ is the average movement time for the corresponding index of difficulty. Fitts found that the human index of performance for the tasks and body parts considered in his study (i.e., arm and hand) ranged from 10 to 12 bits/s. For the same task and body parts, changes in ID did not affect the index of performance (within limits). The following equation was obtained by regressing movement time (MT) on index of difficulty: $\mathrm{MT} = a + b * \mathrm{ID}$, where a and b are empirically determined constants (i.e., intercept and slope, respectively), and $\mathrm{IP} = \frac{1}{b}$ (Fitts & Peterson, 1964). MacKenzie (1992) interpreted the intercept (a) in Fitts' law as the extra time needed to perform "the *select* operation, which typically follows pointing" (p. 98). For example, in a traditional desktop computer interface, the intercept would correspond to the time required to select an icon after placing the cursor over the desired one. If a singe click (instead of a double click) were needed to select an icon, the intercept in Fitts' equation would be expected to decrease, but no change would be expected in the slope (b). Using this information-based model of movement control, Fitts (1954; Fitts & Peterson, 1964) and many others (e.g., Langolf, Chaffin, & Foulke, 1976) have accounted for more than 95% of the variance in movement time for visual target-acquisition tasks. Due to its robustness, this prediction equation is known as *Fitts' law*.

Fitts and Peterson (1964) used a discrete task in which participants, instead of serially moving between two targets as in Fitts' initial study (1954), started at a central resting position and moved to one of two lateral targets (i.e., right or left of the resting position) in response to a light onset (i.e., slightly right or left in front of the participants). Fitts and

Peterson measured both the time elapsed from the appearance of the light to the moment the participants started moving (i.e., reaction time, RT) and the time from the beginning of the movement to the moment the person touched the surface of the target (i.e., MT). They found, as expected, that RT was unaffected by index of difficulty and that MT followed Fitts' law. That is, when using visual target-acquisition tasks in which the participant does not know the target location before the beginning of the trial, Information theory predicts a linear effect of index of difficulty on MT (i.e., Fitts' law), but no effect of index of difficulty on RT.

Although the original formulation of Fitts' law (1954) presented above has consistently yielded good fits to the data, MacKenzie (1992) proposed a modification of Fitts' index of difficulty to make it more faithful to the original Information theory of communication (Shannon & Weaver, 1949/1963): MT = $a + b * \log_2\left(\frac{\mathrm{D}+\mathrm{S}}{\mathrm{S}}\right) = a + b * \log_2\left(\frac{\mathrm{D}}{\mathrm{S}}+1\right)$. This form clarifies the role of the distance (D) as the signal (i.e., desired movement) and of the target size (S) as noise of the channel (i.e., allowable variability in the endpoint of the movement) and it consistently yields slightly better fits to the data than Fitts' original index of difficulty. For this reason, MacKenzie's modification is often used instead of Fitts' law.

Fitts (1954) hypothesized that the index of performance (i.e., channel capacity) would change from one task to another depending on the central demands imposed by the specific task. For example, the presence of longer feedback delays is expected to increase the central demands on the operator. Due to increasing central demands, the human index of performance for a specific task is expected to decrease (i.e., steeper regression lines) with increasing delay (e.g., Hoffmann, 1992) or, in other words, the same increase in index of difficulty is expected to result in a greater increase of MT with longer feedback delays than with shorter ones.

Models Based on Control Theory

Models based on Control theory characterize the performance of the human operator in movement tasks with analytical techniques used to describe feedback-control systems (Craik, 1947). An example of such a system would be an antenna with a tracking system that allows it to rotate to face in the desired direction (Jagacinski, 1977). The components of such a system are: a desired direction as the input signal (analogous to the target in a movement task), an antenna (analogous to the hand), a motor that allows the antenna to move (analogous to the muscles in the arm and hand), and information about the current direction of the antenna that can be compared to the desired one (analogous to sensory feedback).

Control theory uses quantitative measures such as *gain*, *lag*, *time delay*, and *lead* to model human performance. Gain in Control theory describes how fast the system moves to correct for error (i.e., difference between the current and desired position). Movements with higher gains will be faster, whereas movements with lower gains will be slower. Jagacinski and Flach (2003) proposed that gain could be interpreted as an "index of the sensitivity of the system to error" (p. 30). In Control theory, lag and time delay are two distinct concepts: a system with a lag begins its response to an input instantaneously and gradually approaches a steady-state output, whereas a system with a time delay does not begin its response to an input until after a period of time (i.e., the delay) has elapsed, and then, it reaches a steady-state output instantaneously. Although it is not always the case in the literature (e.g., MacKenzie & Ware, 1993), in this paper, the terms lag and time delay are used following the Control-theory convention and, therefore, are not interchangeable (actually, we will talk almost exclusively about time delays). Finally, a lead refers to the prediction of future inputs based on past observations. For example, in order to successfully catch a moving ball, it is necessary to anticipate the position of the ball into the future (i.e., when it will reach the receiver's position). In some cases, the lag (or time delay) of a system may be partially cancelled by a lead.

Researchers in Control theory are generally interested in the *stability* of systems, and the time delay between input and output signal is an important parameter in this context. For the orienting-antenna feedback-control system, a specific combination of time delay and gain will allow the antenna to rotate to the desired direction in an optimal amount of time. If gain is too low (for a specific delay), it will take longer than optimal to reach the desired position, whereas if gain is too high, the antenna will overshoot first and then oscillate around the desired position (i.e., an example of an unstable system). A modification of the delay or gain of the system in the appropriate direction would be necessary to stabilize the system. Delayed systems are expected to require lower gains than systems with the same characteristics but no time delay; therefore, according to Control theory, longer movement times are also expected with increasing feedback delays.

Most of the studies exploring the effect of feedback delays on visual target-acquisition tasks emphasize Information-theory over Control-theory models and the literature review in this thesis reflects this bias. Control-theory models have mostly been used to model tracking performance (i.e., where the human operator is trying to minimize the error between their output and a continually changing target position). Although tracking tasks differ from target-acquisition tasks in some respects, it is important to acknowledge that either approach (i.e., Information or Control theory) could potentially be used to model target-acquisition performance. Control theory's predictions of MT in target-acquisition tasks are equivalent to Fitts' law predictions (e.g., Jagacinski & Flach, 2003) and, in addition, Control theory also predicts the time histories of movements for a given gain, lag, time delay, and lead. Although Information-theory terminology is used when comparing results with previous research, the data from this thesis could be used to develop human-performance models from either (or both) framework(s).

Studies Exploring the Impact of Constant Feedback Delays on Motor Performance

Numerous studies have looked at the effect of constant feedback delays on motor performance during the last 50 years. While early studies were motivated by an interest in the impact of feedback delay on handwriting and drawing, interest in telerobotics and virtual environments has motivated more recent studies.

Early Studies on Feedback Delay and Behavior

The discovery of the interference of delayed auditory feedback on speech (Lee, 1950) motivated an initial interest in the effect of feedback delay on handwriting. van Bergeijk and David (1959) explored the effect of 6 levels of feedback delay between 0 and 520 ms on a handwriting task, finding a gradual increase in completion time with longer feedback delays. They also observed that some participants changed their performance strategy when faced with feedback delays. For example, they described how some participants, instead of writing continuously, often wrote one letter at a time without paying attention to feedback, then stopped to check feedback, and repeated this sequence for the following individual letters until they completed the task. Participants using this strategy took longer to complete the task, but they also produced fewer errors and obtained better "neatness" ratings than participants that attempted to perform the task continuously.

Smith, McCrary, and Smith (1960) studied nine movement tasks under two feedback-delay conditions (i.e., no delay and 520-ms delay). In describing the performance of their participants while trying to place dots in circles (i.e., one of the nine tasks), they wrote "what normally would be fast, smooth, placing motions becomes erratic and oscillatory movements which assume a characteristic jerkiness" in the presence of feedback delays (p. 1013). This task was similar to the visual target-acquisition tasks in which we are interested and Smith et al.'s qualitative observation closely resembles Control theory's description of an unstable system with oscillatory behavior.

Visual Target-Acquisition Performance in the Presence of Constant Feedback Delays

The development of the first telerobotic systems in the early 1960s motivated some research on the effect of feedback delays in visual target-acquisition and manipulative tasks. Ferrell (1965) used a task with 2 degrees of freedom (2DOF) plus grasp in which participants moved a manipulator, which had two parallel fingers, from the starting position to the target location and, then, closed the two fingers to grab the target. In his study, Ferrell introduced feedback delays of 0.0, 1.0, 2.1, and 3.2 s and observed that human operators tended to spontaneously adopt a *move-and-wait strategy* in the presence of feedback delays. This strategy consisted of making an *open-loop (OL) movement* (i.e., without using feedback), waiting for the feedback to show the result of their actions, and then repeating the sequence again until they completed the task. Ferrell hypothesized that movement times in the presence of feedback delay could be estimated knowing the performance of participants in the absence of delay and the number of OL movements participants needed for the task. He assumed that total MT in the condition with feedback delay was equal to the sum of reaction times (RT), open-loop movement times (MT_{OL}), and waiting times (corresponding to the duration of the feedback delay, *Del*) for the total number of OL movements (N).

In order to determine the number of OL movements (i.e., submovements) necessary for the specific task, Ferrell (1965) created an *OL condition* in which participants were instructed to use a move-and-wait strategy in the absence of feedback delay (i.e., they were asked to only open their eyes when they were not moving and only move with their eyes closed). Ferrell measured the movement times for this open-loop condition (MT_{OLC}). In addition, Ferrell compared the number of OL movements in the OL condition to the number of OL movements in the presence of feedback delays and concluded that the number of OL movements for a specific task remained relatively stable across conditions and levels of feedback delay.

Based on the movement time in the no-delay condition (MT_{NoDel}), the MT in the OL condition (MT_{OLC}), and the number of movements needed to complete the task in the OL condition (N), Ferrell proposed two estimates to predict movement time in the presence of feedback delays (MT_{Del}). The first estimate (MT_{Del1}) was based on the MT in the no-delay condition: $MT_{Del1} = MT_{NoDel} + N * (Del + RT)$. Because this estimate took the no-delay condition with continuous movements as its baseline, it was expected to underestimate MT_{Del}. The second estimate (MT_{Del2}) was based on the movement time in the OL condition, $MT_{Del2} = MT_{OLC} + N * Del$. Because it took the OL condition as its baseline and the MT_{OLC} included the time required to open the eyes, focus, and close the eyes, MT_{Del2} was expected to overestimate MT_{Del}. (Notice that MT_{NoDel} included the total MT plus the initial RT, whereas MT_{OLC} included all RTs and all MTs, making it unnecessary to include RT in the MT_{Del2} formula.) And so, Ferrell (1965) proposed using the average of these two estimates, $MT_{Del} = \frac{MT_{Del1} + MT_{Del2}}{2}$ to predict MT in the presence of feedback delays. Using this average estimate, he found good agreement between the predicted and observed MT_{Del}. Ferrell was the first researcher to systematically explore the impact of feedback delays in a visual target-acquisition task similar to teleoperation tasks and his was the first attempt to model human-operator performance in the presence of feedback delays. Most of the later studies exploring the effect of feedback delays used modified versions of Fitts' law to model the performance of human operators.

Both Hoffmann (1992) and MacKenzie and Ware (1993) proposed similar modifications of Fitts' law to account for the effect of constant feedback delays on MT in visual target-acquisition tasks. In Hoffmann's study, participants rotated a control knob whose movements were translated into one-dimensional pen displacement. The participants' task was to move the pen into the target area as quickly as possible. Hoffmann used feedback delays of 30, 200, 500, and 1000 ms, and his analysis of variance (ANOVA) showed significant effects of index of difficulty (ID), feedback delay (*Del*), and the interaction between index of difficulty and feedback delay (*ID∗Del*). In addition, Hoffmann performed a multiple regression of MT on ID, Del, and ID∗Del for each individual participant. He found that the predictor feedback delay (Del) was not significant in 7 out of the 8 participants, but the interaction ID∗Del was always significant. Therefore, he proposed a prediction model using index of difficulty and the interaction between index of difficulty and feedback delay as the predictors and MT as the outcome: MT = $a + b*(c + \text{Del})*\text{ID}$, where a, b, and c are empirically determined constants. This model accounted for 97% of the variance of MT in Hoffman's experiment and for 92% of the variance in Ferrell's (1965) data. However, in contrast to Ferrell, Hoffmann found that the number of OL movements did increase with delay.

MacKenzie and Ware (1993) used a one-dimensional task in which the participants controlled a cursor on a computer monitor using a mouse. The cursor always appeared on the left side of the screen and the target appeared to the right of the cursor. The participants' task was to move the cursor horizontally to the target area as quickly as possible as soon as the target appeared. MacKenzie and Ware used shorter feedback delays (8.33, 25, 75, and 225 ms) than Hoffman (1992), and their ANOVA showed a significant main effect of index of difficulty (ID), feedback delay (Del), and their interaction (ID*Del) on MT, error rate, and index of performance (IP). They tested four different models and finally proposed a model equivalent to Hoffmann's: MT = $a + (b + c*\text{Del})*\text{ID}$. In their study, MacKenzie and Ware's model accounted for 93.5% of the variance in MT. In contrast to Hoffmann (1992), MacKenzie and Ware did not test the regression model for individual participants, and therefore, did not justify the exclusion of Del from the prediction model. In both of these studies, the proposed model changed the index of performance depending on the feedback delay conditions: IP = $\dfrac{1}{b + c*\text{Del}}$.

In a later study, Mateo, Manning, Cowgill, Simpson, Moore, Weisenberger, and Gilkey (2005) explored the impact of constant feedback delays in a three-dimensional (3D) visual target-acquisition task (i.e., more similar to real telemanipulation tasks) performed in a virtual environment (VE). In their study, participants wore head-mounted displays (HMD) that displayed a cubical virtual workspace in 3D and controlled a cursor in all three dimensions of the workspace via a PHANTOM® haptic device. Their cursor was always visible in the virtual workspace and hollow cylindrical targets with one circular face facing the participant (see Figure 2) appeared one at a time in one of 33 possible locations within the 3D virtual workspace. The participants' task consisted of piercing the front circular side of the cylindrical target in the workspace as quickly as possible. Each time they acquired a target successfully, the target disappeared from its location and reappeared at a different location.

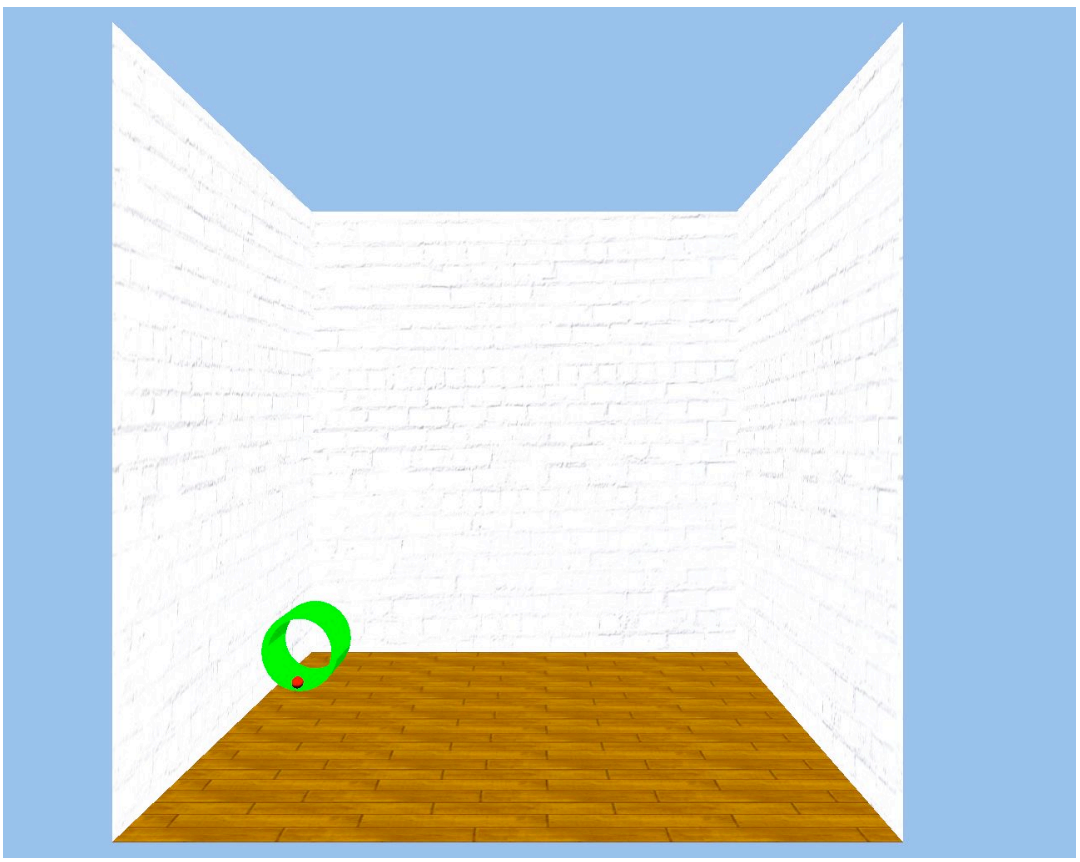

Figure 2. Screen shot of the virtual workspace used by Mateo et al. (2005), showing a hollow cylindrical target and the cursor on the bottom left corner of the workspace.

Using feedback delays of 32, 266, 532, and 799 ms, Mateo et al. found results (shown in Figure 3) in their 3D task that were consistent (i.e., same trends) with the models

Hoffmann (1992) and MacKenzie and Ware (1993) had proposed for one-dimensional tasks. That is, with increasing feedback delay, the index of performance of the task gradually decreased (i.e., the regression line became steeper). However, the amount of variance accounted for in Mateo et al.'s study was lower than in Hoffman's or MacKenzie and Ware's study. Nevertheless, Mateo et al.'s study suggested that the VE and experimental design they used could be used to study visual target-acquisition performance in the presence of feedback delays.

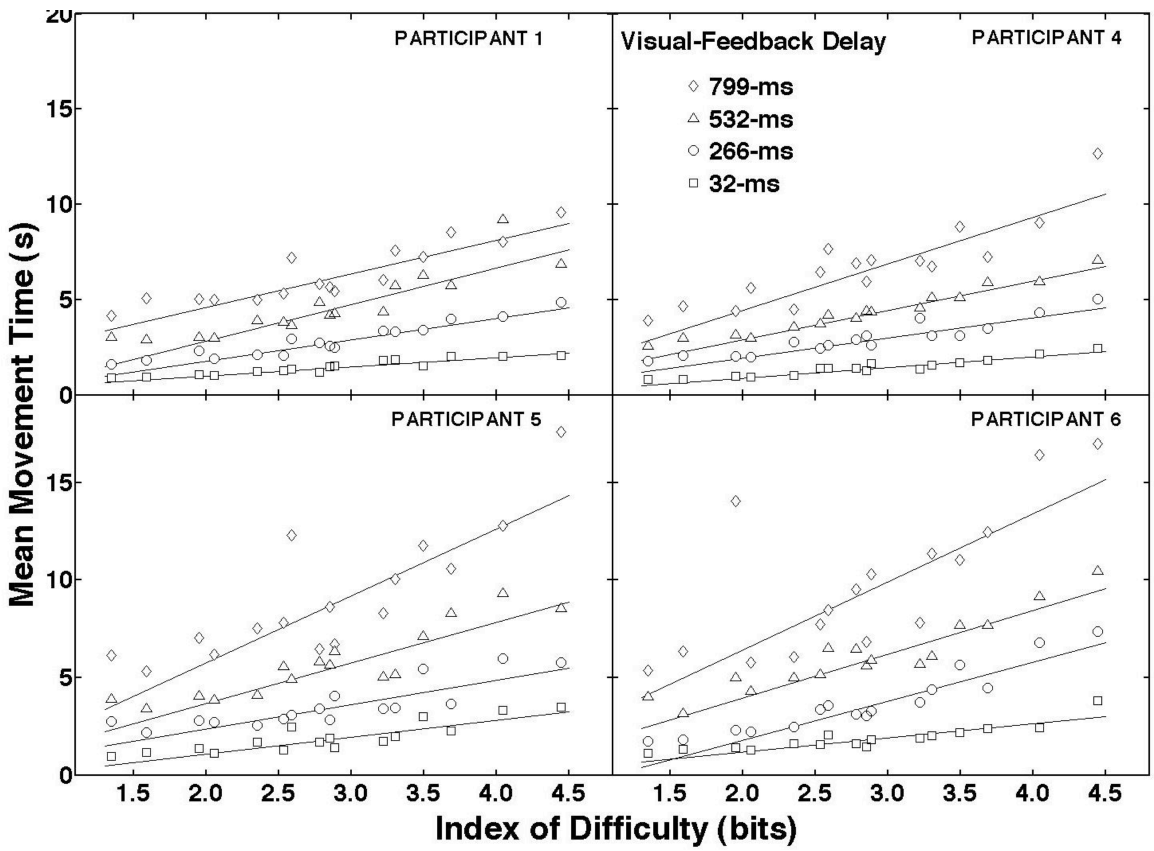

Figure 3. Mean MT (in seconds) as a function of index of difficulty: ID = $\log_2\left(\frac{D}{S}+1\right)$ (ID = 1.34, 1.59, 2.06, 2.35, 2.53, 2.86, 2.78, 2.85, 2.89, 3.22, 3.30, 3.49, 3.69, 4.05, & 4.45 bits) for each level of feedback delay (32-ms, 266-ms, 532-ms, and 799-ms) used by Mateo et al. (2005). Each panel shows the mean data of 1 of the 4 participants. As feedback delay increases, the slopes of the linear fit become steeper and the fits become poorer.

So and Chung (2002) also explored motor performance in the presence of feedback delays in an immersive VE. They tracked both the head and hand movements of the participants. Feedback delays of 0, 110, and 220 ms were introduced both between the moment the participants moved their head and the moment the scene view was updated (i.e., *head-related delays, HeDel*) and between the moment the participants moved their hands and the moment their virtual hand moved in the scene (i.e., *hand-related delays, HnDel*). These hand-related delays correspond to what we have previously referred to as feedback delays, whereas head-related delays were not addressed in any of the studies described above. (Mateo et al. used a VE and head-mounted displays, but they did not track head movements in their study.) So and Chung hypothesized that the relative impact of distance between targets and target size on MT would differ in the presence of HnDel or HeDel. Specifically, they hypothesized that, as HnDel increases, size would have a greater effect on MT than distance (first hypothesis) and that, as HeDel increases, distance would have a greater effect on MT than size (second hypothesis). For this reason, they did not combine distance and size into an index of difficulty (which assumes equal impact of distance and size), but instead used them as independent predictors.

So and Chung (2002) supported their first hypothesis, which is consistent with (and was, actually, motivated by) Woodworth's (1899) *two-component theory* dividing visual target-acquisition movements into two parts: an *initial-impulse* phase and a *current-control* phase. The initial-impulse phase involves an OL movement to get into the vicinity of the target (and therefore, is expected to be more affected by distance between targets), whereas the current-control phase involves closed-loop (CL) movements (i.e., using feedback) to acquire the target (and therefore, is expected to be more affected by target size). So and Chung's findings suggested that feedback delays affected the current-control phase more than the initial-impulse phase. Their second hypothesis was only marginally supported. So and Chung expected head movements to occur only when targets were very far apart, but apparently participants used head movements also when the targets were comfortably within their field of view in order to center the area of interest in their foveal region.

Based on their results, So and Chung (2002) concluded that combining target size and distance between targets into an index of difficulty may not be appropriate in the presence of feedback delays, as some of the previous models had done (e.g., Hoffmann, 1992). Using a stepwise regression of MT on the predictors HnDel, HeDel, $\log_2\left(\frac{2}{S}\right)$, $\log_2(D)$, and all their interactions, So and Chung obtained the following model, which accounted for 95% of the variance in MT:

$$MT = -0.82 + 3.47 * HnDel + 0.14 * \log_2\left(\frac{2}{S}\right) + 0.26 * \log_2(D) + 4.51 * HnDel * HeDel + 0.32 * HeDel * \log_2(D) + 0.50 * HnDel * \log_2\left(\frac{2}{S}\right)$$

Although this model has more parameters and does not account for more variance than the previous models, it addressed some issues that are relevant for the study of movement control in VE (where feedback delays affecting head movements, HeDel, may be an issue). In addition, this model emphasized the different relative impact of distance and size on MT when the task is completed in the presence of feedback delays affecting hand movement (i.e., HnDel).

Studies Exploring the Impact of Variable Feedback Delays on Motor Performance

Most traditional teleoperation systems have used private, exclusive communication channels with stable, predictable characteristics that resulted in transmission delays that remained constant. Not surprisingly, most studies exploring the effect of feedback delays on teleoperation performance have used exclusively constant delays. However, during the last decade, many researchers have proposed using the Internet as the communication channel for some telerobotic applications. Some of the properties of the Internet make it a desirable channel for some applications, such as telesurgery in the battlefield or in rural areas (e.g., Hanly & Broderick, 2005). However, the Internet is not a private, exclusive channel but, instead, it is public and shared. As a consequence, the transmission delay (and the resulting feedback delay) experienced by operators varies continuously depending on factors such as number of users, routing path, and so forth. Moreover, the factors affecting Internet delays are often unpredictable and difficult to model (Oboe & Fiorini, 1997).

Most of the research in the field of Internet-based teleoperation has been performed from an engineering perspective, addressing issues such as the technical feasibility of Internet-based teleoperation (e.g., Goldberg, Gentner, Sutter, Wiegley, & Farzin, 2002), the efficacy of different Internet protocols (e.g., Oboe & Fiorini, 1997), or interface design (e.g., Hu, Yu, Tsui, & Zhou, 2001). In general, it has been assumed that direct control of online robots is not possible and, therefore, most implementations have used supervisory control (e.g., Hu et al., 2001). Furthermore, research on the impact of variable transmission delay created by the Internet on motor performance has been concerned mostly with haptic-feedback delays (e.g., Chopra, Spong, Hirche, & Buss, 2003). Because haptic feedback is often provided through the control input (e.g., a joystick), directly affecting its position, haptic-feedback delay has been viewed as a more serious problem than visual-feedback delay, which does not directly affect the position of the operator's controller. Such haptic-feedback delays are very likely to destabilize the system, since the operator cannot ignore them or compensate for them (e.g., by using a move-and-wait strategy) to avoid instability. For that reason, some researchers (e.g., Anderson & Spong, 1988) have suggested that instability is only a problem when haptic feedback is delayed.

However, MacKenzie and Ware (1993), So and Chung (2002), and others clearly showed that *visual*-feedback delays do have a detrimental effect on performance and ignoring these findings seems unwise. Some researchers (e.g., Jay & Hubbold, 2005) have even found a greater impact of visual-feedback delays than haptic-feedback delays on performance. Exploring the impact of variable visual-feedback delays on human motor performance is especially interesting now that the Internet has become a candidate channel for teleoperation applications. Current models of human performance in the presence of feedback delays do not consider delay variability as a parameter (e.g., MacKenzie & Ware, 1993). For example, these models cannot address the differences among a constant 500-ms feedback delay, a variable feedback delay with a mean of 500 ms, and a variable feedback delay with a maximum magnitude of 500 ms. In the literature, however, it is generally assumed that hard-to-predict variable feedback delays result in worse performance than constant feedback delays, even when the latter are much longer (Lane, Carignan, & Akin, 2002). A model that can predict performance for both constant and variable feedback delays would be useful and the first step in the development of such a model is understanding how feedback-delay variability affects human motor performance.

Impact of Variable Feedback Delays on Individual Performance

An interest in virtual environments (VEs) motivated Watson and his colleagues (Watson, Spaulding, Walker, & Ribarsky, 1997; Watson, Walker, Ribarsky, & Spaulding, 1998) to explore the effects of feedback delays with different levels of variability on the performance in a movement task. These researchers were interested in variability occurring in VEs (not necessarily in teleoperation applications) and, as a consequence, their manipulations involved varying the frame rate during their trials. With these frame-rate manipulations, however, feedback-delay variability also varied across conditions and, therefore, their findings are relevant.

The task used by Watson and his colleagues consisted of the participants tracking an object moving from left to right in front of them, grasping it, and then moving the object to a placement box on the right side of the workspace. They divided this movement task into two components: the OL *grasping movement* and the CL *placement movement*. Not surprisingly, given that the manipulation affected feedback, they found that the CL movement was more sensitive to the feedback-delay-variability manipulation than the OL movement. The results of their studies suggested that frame-rate variations that resulted in standard deviations of less than 83 ms had no effect on performance, but larger ones could affect performance.

In addition, Watson et al. (1998) explored different ways of manipulating feedback-delay variability: either by adding a specific frame-rate variation (e.g., *SD* = 2.0 Hz) to all mean frame rates or by adding a frame-rate variation whose magnitude is proportional to the mean feedback delay (e.g., if *SD* = 20% of the mean, then *SD* = 1.8 Hz for the 9-Hz mean, *SD* = 2.6 Hz for the 13-Hz mean, and *SD* = 3.4 Hz for the 17-Hz mean). Watson et al. concluded that manipulating feedback-delay variability in magnitudes proportional to mean feedback delay (e.g., 20% of the mean) resulted in more consistent results than manipulating it in absolute magnitudes (e.g., 2.0 Hz for all mean levels). However, Watson et al. found that magnitude was not the only aspect of the feedback-delay-variability manipulation that made a difference: distribution shape (i.e., skewness) was also important. That is, feedback-delay variability seemed to affect performance more when the feedback-delay distribution was negatively skewed (i.e., biased toward longer delays) than when the feedback-delay distribution was symmetrical around the mean (or positively skewed).

Interested in telerobotic applications, Sheik-Nainar, Kaber, & Chow (2005) explored the effect of a number of variables (e.g., level of automation, delay) on the performance of a telerover-navigation task. Their network-delay manipulation was especially relevant to the present study. The three network-delay conditions were: no delay, constant 1000-ms delay, and random delay (varying between 750 and 1250 ms). Surprisingly, they found that participants completed the navigation task faster and more accurately under the random-delay conditions than under the constant-delay conditions. Sheik-Nainar et al. explained this counterintuitive result by pointing out that the mean feedback delay in their random-delay condition was shorter (i.e., *M* = 960 ms, *SD* = 140 ms) than the constant delay (i.e., 1000 ms). If true, this argument indicates that the effect of mean feedback delay on performance was greater than the effect of feedback-delay variability, since the effect of a standard deviation of 140 ms was negligible compared to the effect of changing mean feedback-delay by 40 ms.

Because of their different motivations and the way they manipulated delay variability, the results from these studies provide limited information about the nature of the relationship between feedback-delay variability and individual motor performance. More research is necessary to better understand how network-like variable delays affect motor performance. Using feedback-delay distributions more similar to the ones teleoperators face when controlling a robotic arm over the Internet, yet keeping experimental control, may help understand how real-network delays affect performance.

Impact of Variable Feedback Delays on Collaborative Performance

A few studies interested in collaborative motor performance have explored the impact of delay variability across a network (e.g., the Internet). In this context, participants needed to coordinate their movements with their delayed virtual partners in order to perform a specific task. Although the nature of collaborative tasks is different from individual tasks, some of the findings from these studies provide useful information about the effect of variability of transmission delays on human motor performance.

Park and Kenyon (1999) compared the performance of pairs of participants interacting through different types of networks, creating combinations of latency (i.e., mean delay) and jitter (i.e., delay variability) and the average and maximum delays for each network condition were measured during experimental sessions. For example, two of the networks created a constant 200-ms delay (*Constant-Long*) and a variable delay with an average between 150 and 300 ms and a maximum delay of approximately 2 s (*Variable-Long*). During the collaborative task, each participant in the VE controlled either a ring or a "path" (i.e., the path was a stick-like object with one of several shapes: straight, curving, etc.). The collaborative task consisted of transferring the ring, controlled by the first participant, along the corresponding path, controlled by the second participant, as quickly and accurately as possible. Each contact between path and ring was considered an error and both MT and number of errors were recorded as dependent measures.

Park and Kenyon found a significant effect of mean delay and of delay variability on collaborative performance. In their post-hoc analysis, they showed that MTs for the Variable-Long condition were significantly longer than MTs for the Constant-Long condition. However, these findings were obtained in a collaborative task and the authors did not systematically vary mean delay and delay variability (i.e., it was not possible to compare delays of same mean magnitude but different variability or of the same variability but different mean magnitude).

Gutwin (2001) was also interested in the effect of network delays on collaborative performance. He explored the effects of latency (i.e., mean delay) and jitter (i.e., delay variability) on performance in two different tasks: a *prediction task* and a *coordination task*. In the prediction task, the participants needed to determine as quickly as possible the location where the other participant was pointing, whereas in the coordination task they were required to pick objects from a central area and transport them to their respective target areas. Although the participants picked and transported the objects individually in the coordination task, they were not allowed to grab the same object from the central area at the same time. When they both tried to pick the same object from the central area, it was considered an error (i.e., some collaboration was required). Gutwin found that delay variability had an impact on both MT (for *SD* ≥ 600 ms) and error rate (for all levels of delay variability) in the prediction task, but only on error rates in the coordination task. Mean delay also had an effect on error rate in the coordination task (for *M* ≥ 240 ms).

Although Gutwin did vary both mean delay and delay variability and was able to compare the effects of each one of them on performance, only feedback from the partner's movements was delayed. Therefore, Gutwin's study provided no direct information about how delaying feedback resulting from the participants' own cursor movements affects performance. However, some of the phenomena described in his study are relevant to individual performance. For example, if we assume that operators use a move-and-wait strategy when performing visual target-acquisition tasks with feedback delays (i.e., they perform an OL movement, wait for feedback, perform the next OL movement, and so forth) as previous research has suggested (e.g., Ferrell, 1965), then determining when the visual display is reflecting the end of their OL movement is important. It is reasonable to expect the use of this move-and-wait strategy to be hindered if a cursor pause cannot be reliably interpreted as the end of operator's movement (i.e., the cursor may also stop due to a freeze resulting from feedback-delay variability). Even though findings from collaborative studies are informative, more research exploring the effect of variable feedback delay affecting the participant's own movements is necessary to understand how feedback-delay variability impacts performance in individual visual target-acquisition tasks.

Strategies Adopted by Human Operators to Deal with Feedback Delays

Given the limited research exploring delay variability, it is not surprising that the strategies operators use to deal with feedback delays described in this section are mostly concerned with constant delays. However, in the second subsection, we attempt to describe how these strategies would apply to the case of variable feedback delays.

Strategies to Deal With Constant Feedback Delays

As Smith et al. (1960) described, performing in the presence of feedback delays can be difficult and frustrating even for very simple movement tasks. When presented with this type of situation, humans tend to adopt new strategies to overcome this difficulty. In agreement with previous less systematic observations (e.g., van Bergeijk & David, 1959; Crossman & Goodeve, 1963/83), Ferrell (1965) showed how participants performing an individual movement task with 2 degrees of freedom plus grasp in the presence of constant feedback delays tended to adopt a move-and-wait strategy for feedback delays of 300 ms or more. In general, these results have been confirmed in the literature (e.g., Black, 1971, found similar results for a 6-degree-of-freedom task). Only the magnitude of the feedback delay at which operators transition from a relatively continuous movement to a move-and-wait strategy has been disputed (e.g., Hoffmann, 1992, proposed that the strategy changes when the human operator is faced with feedback delays of 700 ms or more, instead of the 300-ms feedback delay proposed by Ferrell), but not the fact that operators naturally adopt this strategy in the presence of feedback delays.

Another strategy operators use to deal with feedback delays is anticipating future events, so that they need to rely less on feedback. For example, if a person is asked to track a predictable signal, such as a sine wave, at first the person will tend to follow the feedback signal but, with some practice, the person will realize the repetitive pattern and will begin to anticipate the future position of the target without the need to wait for the feedback (Jagacinski & Flach, 2003). When this happens, the person starts using an internal model of the target trajectory to guide the movements, instead of the slower visual feedback. In visual target-acquisition tasks, something similar can happen if, for example, the participant knows in advance exactly where the next target is going to appear (i.e., if the position of future targets is predictable). Given this tendency of humans to anticipate future events, researchers often use randomly appearing stimuli in both visual-tracking and visual-target-acquisition experiments to confound such predictive strategies (Sheridan & Ferrell, 1974).

Hanly and Broderick (2005) reviewed some strategies adopted by telesurgeons to deal with feedback delays. These strategies included slowing down (i.e., lowering the gain, in Control-theory terms), anticipating (i.e., lead), and delegating high-precision tasks to the in-site surgeon. This last strategy obviously depends on the availability of a local human (or intelligent robotic) agent.

Strategies to Deal With Variable Feedback Delays

The first two strategies described above (i.e., move-and-wait and anticipation) have been found in research exploring human motor performance in the presence of constant feedback delays. However, it is reasonable to expect that operators will attempt to adopt similar strategies when feedback delays are variable, even if the low predictability of variable feedback delays may make the use of these strategies more difficult and frustrating for the operator. For example, using a move-and-wait strategy may be more difficult in the presence of variable feedback delays because the waiting period is continuously changing. In addition, an anticipatory strategy may be difficult in these circumstances because feedback arrives asynchronously, providing a temporally warped view of the movement of the target or the actions of the effector (e.g., a smooth movement of the operator often appears as a sequence of cursor jumps). However, no empirical research has directly addressed the strategies adopted by human operators in the presence of variable feedback delays.

One study that provided some information about how human operators adapt to variable feedback delays came from the collaborative-performance literature. Vaghi, Greenhalgh, and Benford (1999) were interested in studying how humans adapt to variable feedback delays (in their case, gradually increasing in magnitude, not randomly changing like in network delays) in collaborative tasks. They used a virtual ball game (soccer-like) in which avatars (controlled by the participants) were used to hit a ball into the opponent's goal as their collaborative (actually opponent) task. The movements of each participant's avatar were limited to the half of the field corresponding to the participant. Within each 23-min trial, feedback delay was gradually increased from 0 ms to 999 ms (there were 9 different levels of delay). Vaghi et al. observed how their participants changed their coping strategies as delay increased. They showed that participants took some time to notice the delay change but, once they did, they tended to adapt to it naturally. The coping strategies adopted by their participants included slowing down (i.e., lower gain) and anticipating the future position of the ball based on its movement trajectory (i.e., lead). This study led to the idea that providing the participants with information about the current feedback-delay characteristics may facilitate performance by reducing the time needed to notice feedback-delay changes.

Design Techniques to Ameliorate the Effect of Feedback Delays

Many techniques have been developed and implemented to help users deal with feedback delays, mostly for the case of constant feedback delays. Nevertheless, these techniques are often applicable to variable feedback delays. In addition, a few techniques have been developed exclusively to deal with variable delays. However, these have often been developed to deal with variable haptic-feedback delays and it is unclear how they apply to visual-feedback delays. This aspect is important because, even though the importance of haptic feedback for telerobotic performance has been recognized, many telerobotic applications (e.g., telesurgery, Hanly & Broderick, 2005) provide only visual feedback.

Design Techniques to Aid With Constant Feedback Delays

Probably two of the most important aids are *predictive displays* and *supervisory control*. Predictive displays provide a visual indication of the expected effector position and they are often superimposed on the actual display (e.g., Lane et al., 2002, used a translucent image of the effector for the predictive display and a solid image for the last known position). In order to make accurate predictions, it is necessary to have a dynamic model of the robotic arm and the environment. However, even with imperfect models, predictive displays can improve performance substantially.

Sheridan (1992) proposed the use of supervisory control in order to deal with transmission delays. Supervisory control gives more autonomy to the telerobot at the remote site (see Figure 1 on page 11), which is able to understand higher order commands (i.e., ‘pick up the glass’ instead of ‘move to the left’) given by the operator. Therefore, the operator does not directly control the detailed movement of the telerobot, but instead the telerobot receives higher order operator commands and translates them into actions. The use of supervisory control minimizes the effect of transmission delay by giving more control to the telerobot. This strategy resembles the one telesurgeons use when they delegate part of their workload to the in-site surgeon (e.g., Hanly & Broderick, 2005). If the task is very simple and predictable, full automation could also be used and the feedback-delay problem could be completely eliminated.

Other techniques that have been used to ameliorate the effect of constant feedback delays on performance include gain adaptation and level-of-detail (LOD) management. The gain-adaptation approach (e.g., Sheik-Nainar et al., 2005) consisted of adjusting the gain of the operator-effector system as a function of the feedback-delay conditions. That is, in the presence of a longer feedback delay, the gain of the system was lowered and the same operator movement resulted in a slower movement of the effector. This approach was based on the Control-theory idea that, as feedback delay increases, a reduction of system gain is necessary to maintain the system stability. For example, even in the presence of very long feedback delays, if the operator moves slowly enough, there may be few errors. As expected, Sheik-Nainar et al. found that participants performing the task in the presence of feedback delays took longer to complete the task, but made significantly less errors, in the gain-adaptation condition when compared to the no-adaptation condition.

The LOD-management approach (e.g., Watson et al., 1998), instead of assisting the operator to perform the task in the presence of feedback delays, attempts to reduce the feedback delay experienced by the operator in order to preserve performance. Specifically, LOD reduces the quality or complexity of the visual image on the display when feedback delays are considered long enough to affect performance. Because less information needs to be transmitted through the communication channel when LOD is reduced, feedback delay can be reduced as well and, as a consequence, performance can be maintained. However, there is a limit to how much LOD can be reduced before performance is affected; that is, there is a tradeoff between LOD and feedback delay. The nature of this trade-off seems likely to be extremely task dependent and, therefore, the LOD-management approach may need to be tuned for each specific task and subtask.

Design Techniques to Aid With Variable Feedback Delays

Although designed to aid with constant delays, we expect predictive displays and supervisory control to be also helpful in the presence of variable feedback delays. For the specific case of variable delays, some researchers (e.g., Kosuge & Murayama, 1998) interested in haptic-feedback delays have proposed approaches to eliminate (actually, minimize) delay variability at the expense of increasing mean delay by using a *buffering technique*. When using this technique, a variable delay is transformed into a (nearly) constant delay by determining the criterion delay value under which 95% of the possible delay magnitudes fall. For example, in the case of a uniformly distributed variable delay with a mean of 400 ms and a range of 400 ms (i.e., 200-600 ms), 95% of the delay values are expected to be under 590 ms (i.e., between 200 ms and 590 ms). Once this criterion delay value has been determined, feedback information that arrives in less than 590 ms is put into a buffer and not displayed until 590 ms have elapsed so that, 95% of the time, the delay experienced by the operator is constant and equal to 590 ms. This approach assumes that variable delays are so disruptive that dramatically increasing the overall delay (e.g., from M = 400 ms to $M \cong$ 590 ms) would improve performance. Kosuge and Murayama's technique was developed to deal with haptic-feedback-delay variability and its applicability to visual-feedback delays is unknown.

Gutwin, Benford, Dyck, Fraser, Vaghi, and Greenhalgh (2004) used collaborative tasks similar to those used by Gutwin (2001) and studied how providing information about the current feedback-delay conditions through display *decorators* affected collaborative performance. In their first experiment, for example, participants had to determine the position at which the partner had stopped moving (i.e., prediction task). Delay jitter (i.e., variability) was introduced between the moment the partner moved and the moment that information was displayed to the participant by freezing the cursor in 5% of the frames for an interval of time corresponding to the jitter condition (i.e., 0, 800, 1100, or 1400 ms). As soon as this interval of time had elapsed, the cursor jumped to its current position. In the decorator condition, the color of the cursor changed as a function of the time since the last cursor-position update was received. That is, a white cursor indicated that the cursor position was up to date, a cursor with darker shades of grey indicated longer time since last update, and a black cursor indicated that no update had been received for 1000 ms. Providing information about the current feedback-delay conditions (e.g., by changing the color of the cursor) helped users adopt the corresponding strategies faster and, as a consequence, improved their performance when compared to the no-decorator condition. Gutwin et al.'s study used collaborative tasks and the delay only affected feedback resulting from the partner's actions. The effectiveness of this display strategy to ameliorate the effect of variable feedback delays affecting the operator's actions while working individually is unclear. However, using a display cue (i.e., decorator) that informs the operator about the feedback-delay conditions seems like a reasonable idea if network characteristics (i.e., mean delay and delay variability) can be identified in an instant-to-instant basis.

LOD-management and gain-adaptation techniques could also be used to help operators perform motor tasks in the presence of variable feedback delays when network characteristics can be identified in an instant-to-instant basis. Sheik-Nainar et al. (2005) found that gain adaptation was similarly beneficial for performance (i.e., shorter completion times and less errors) in the presence of random feedback delays as it was in the presence of constant delays.

Goals of the Present Thesis

The main goal of this thesis is to study how variable feedback delays (e.g., Internet delays) affect performance in visual target-acquisition tasks similar to those performed in teleoperation situations. Previous studies exploring how variable delays affect collaborative performance (e.g., Park & Kenyon, 1999) and how variable haptic-feedback delays affect individual performance (e.g., Kosuge & Murayama, 1998) have found a negative relationship between delay variability and motor performance (i.e., higher levels of feedback-delay variability resulting in worse motor performance). In the context of visual-feedback delays and individual performance, Lane et al. (2002) stated that because shorter variable feedback delays were found to be more detrimental to performance than longer constant feedback delays, a strategy that holds information for a period of time to guarantee a constant feedback delay could be used to improve performance. However, Lane et al. did not actually present any empirical data to support this claim. Furthermore, most studies exploring the effect of feedback-delay variability on individual motor performance suggested that the effect of variable feedback delays might be very small (e.g., Watson et al., 1998) or even negligible (e.g., Sheik-Nainar et al., 2005) compared to the effect of mean feedback delay.

In this thesis, we provide empirical data to determine the effect of feedback-delay variability on visual target-acquisition performance, as well as how this effect compares to the effect of mean feedback delay. Improving our understanding of these topics is essential for the development of models to describe human performance in the presence of variable feedback delays, as well as for the development of new techniques to aid operators in dealing with this type of delays (and evaluation of the current ones).

In the present study, participants performed a visual target-acquisition task, based on Fitts' (1954) original work, similar to the one described by Mateo et al. (2005). This task was chosen because it was representative of some of the components of a real teleoperation task (e.g., moving a robotic arm to a desired position). In addition, the virtual environment used by Mateo et al. allowed the task to be performed in 3D space, therefore making it more similar to real teleoperation tasks and it allowed the experimenter to manipulate relevant variables, such as task difficulty and feedback-delay characteristics. *Target size*, *distance between targets*, *mean feedback delay*, and *feedback-delay variability* were manipulated in the present study. Given that the motivation for this research was to understand the effect of real network delays on motor performance, feedback-delay variability was manipulated simulating Internet delays and feedback delays were kept in the range of mean-delay and delay-variability values reported for Internet connections in the literature (e.g., Gutwin, 2001).

Dependent variables included both *movement time* (i.e., *MT*, time elapsed from the appearance of a target to the moment the participant successfully acquired it) and *number of errors* (i.e., number of times the participant attempted to unsuccessfully acquire a target by pressing the button while outside the target area). Movement times were further subdivided into *reaction time* (*RT*); time to complete the initial distance-covering, open-loop (OL) part of the movement (MT_{OL}); and time to complete the later "homing in," closed-loop (CL) part of the movement (MT_{CL}).

Predictions

We examined a total of eight predictions in this thesis. The first four predictions describe the expected main effects for each of the four independent variables and the last four describe how we expected a subset of the effects of the independent variables to compare to or interact with each other.

Prediction 1: Effect of Target Size

We predicted smaller target sizes to result in longer MT_{CL} and more errors than bigger target sizes, but we expected changes of target size to have no effect on RT or MT_{OL}. This prediction was based on the assumption that target size would be particularly important during the CL part of the movement but would not affect either motor planning (e.g., Fitts & Peterson, 1964) or the initial OL part of the movement (e.g., So & Chung, 2002). Participants were only expected to press the response button during the CL part of the movement (i.e., when close to the target).

Prediction 2: Effect of Distance Between Targets

We predicted longer distances between targets to result in longer MT_{OL}, but we expected changes in distance to have no effect on RT, MT_{CL}, or number of errors. The rationale behind this prediction was that the initial OL movement would cover most of the distance and, therefore, during the CL part of the movement distance would have no effect on performance. Predictions 1 and 2 were based on So and Chung's (2002) and Woodworth's (1899) findings.

Prediction 3: Effect of Mean Feedback Delay

We expected longer mean feedback delays to result in longer MT_{CL} and more errors than shorter mean feedback delays, but we expected changes in mean feedback delay to have little or no effect on RT or MT_{OL}. Given the nature of the manipulation, we expected changes in mean feedback delay to affect only parts of the movement dependent on feedback (i.e., MT_{CL} and number of errors).

Prediction 4: Effect of Feedback-Delay Variability

We expected greater feedback-delay variability to result in longer MT_{CL} and more errors than lower delay variability, but we expected changes in delay variability to have little or no effect on RT or MT_{OL}. The rationale behind this prediction was identical to the rationale behind prediction 3: given that delay variability is a feedback manipulation, only dependent measures sensitive to feedback should be affected.

Prediction 5: Differential Effects of Mean Delay and Delay Variability

We expected the effect of increasing (e.g., doubling) mean feedback delay to have a greater negative effect on performance (i.e., longer MT_{CL} and more errors) than an equivalent increase (e.g., doubling) of feedback-delay variability. This prediction was based on the results of previous studies (e.g., Sheik-Nainar et al., 2005; Watson et al., 1998) and our unpublished pilot study.

Prediction 6: Interaction Between Mean Delay and Delay Variability

We expected the effect of feedback-delay variability on MT_{CL} and number of errors to be greater for longer mean feedback delays than for shorter ones. The rationale behind this prediction was that, as mean delay makes the task more challenging to perform, the same levels of variability that were easy to handle with short or no feedback delays would become more difficult to handle.

Prediction 7: Interaction Between Target Size and Mean Delay

We expected the effect of target size on MT_{CL} and on number of errors to be greater for longer mean feedback delays than for shorter ones. This prediction is consistent with So and Chung's (2002) results and it is based on the idea that, as feedback conditions worsen (i.e., longer mean delay), acquiring smaller targets becomes more difficult than under more favorable feedback conditions.

Prediction 8: Interaction Between Target Size and Delay Variability

The effects of target size on MT_{CL} and on number of errors were expected to be greater for higher levels of delay variability than for lower ones. Although we are not aware of any study exploring this relationship before, the rationale behind this prediction is identical to the rationale behind prediction 7, given that both mean delay and delay variability are manipulating feedback-delay conditions.

II. METHOD

Participants

There were 6 paid participants (3 males and 3 females) from the subject panel at AFRL/HECB, Wright-Patterson Air Force Base. The participants ranged in age from 19 to 25 years old. All of them were right-handed and had normal color vision. Out of the 6 participants, 3 were familiar with the equipment and had some experience with similar target-acquisition tasks. The remaining 3 had no previous experience with the equipment or task.

Apparatus

During this experiment, the participants were immersed in a three-dimensional (3D) virtual environment (VE) created using two computers (i.e., a server and a client) connected through a local network. Participants wore head-mounted display (HMD) through which the workspace was presented to them in 3D and used a PHANTOM® haptic device for input to the VE.

Computers

The server computer was a dual-processor Alienware (AMD Athalon 2800 MP) and the client computer was a dual-processor Dell Precision 530 workstation (2.2 GHz IntelXeon™ processor). The server contained the necessary trial-by-trial information (i.e., order of trial presentation, target size and location, and frame-by-frame delays) and transferred it to the client during the experiment. The participant interacted exclusively with the client, which was connected to both the input (PHANTOM®) and output (head-mounted display) devices. The position of the PHANTOM® device at every frame was transferred from the client to the server and stored in the latter. No data were stored in the client computer. Both computers used Windows XP Professional. The 3D VE was generated with a program created in C++ using the Vega Prime API graphics (MultiGen-Paradigm).

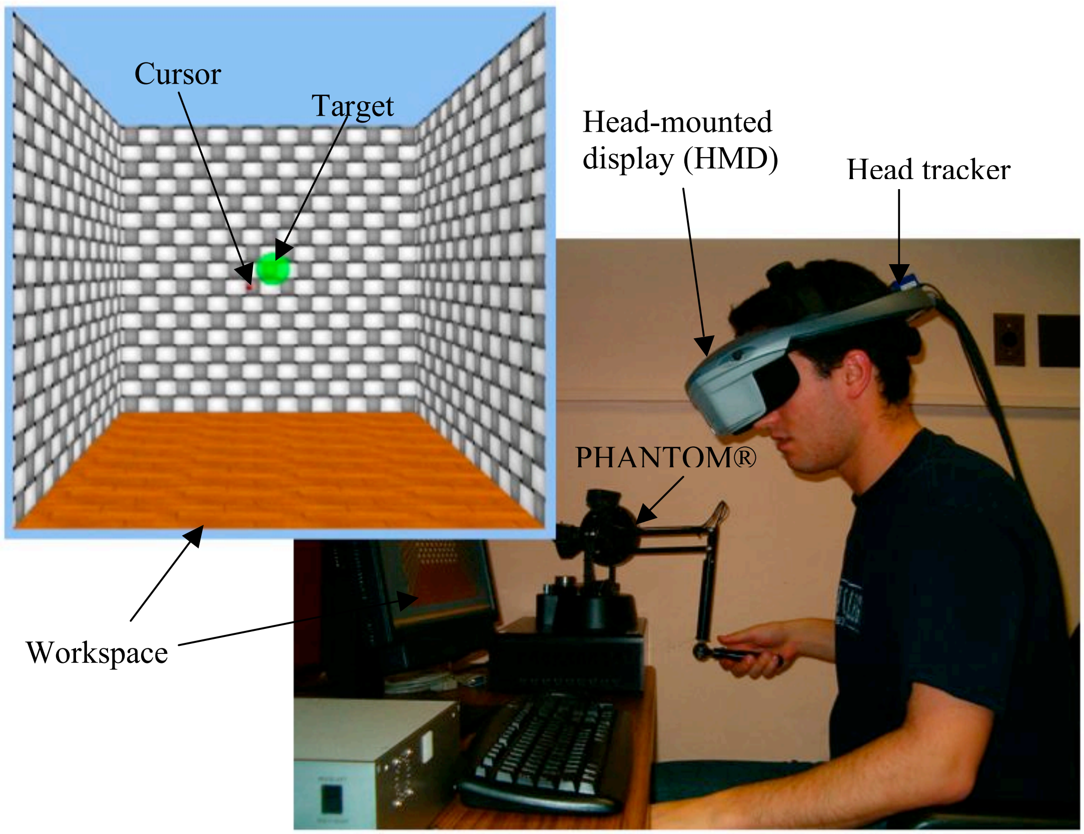

Figure 4. Each participant sat in front of a computer, wearing a head-mounted display (HMD) through which the 3D workspace was displayed. The 3DOF head tracker attached to the back of the HMD. The 3D image displayed through the HMD in stereo was also displayed on the monitor (in 2D) for the experimenter's use. A view of this workspace is shown on the top left corner (see text for details). The PHANTOM® was placed on the right side of the participant.

Head-Mounted Display (HMD)

The 3D VE was displayed to the participants through a Visor SX HMD (1280 x 1024 pixels, 60 Hz stereo), from NVIS, Inc. A head tracker (Intersense) with 3 degrees of freedom (3DOF) was attached to the HMD so that the VE could be updated to match the participant's head orientation (i.e., head, pitch, and roll), allowing participants to use motion-parallax cues. The HMD was connected directly to the client workstation and a 2D version of the image being displayed to the participant was available to the experimenter on the monitor of the client workstation (see Figure 4).

PHANTOM® Device

Participants interacted with the VE by means of a 6DOF PHANTOM® Premium 1.5 haptic device (SensAble Technologies, Inc). The 6DOF PHANTOM® device allowed the participants to move the cursor in the three dimensions of the virtual environment (x, y, and z) and press a button (see Figure 5) to select a target. Haptic feedback was not provided (i.e., the PHANTOM® was used merely as a 3D mouse) and the remaining 3DOF of the PHANTOM® (i.e., head, pitch, and roll) were not used or recorded. The only reason to use a 6DOF PHANTOM® device instead of its 3DOF counterpart was that the former has a button (see Figure 5), which was used as response input, that the latter lacks.

The PHANTOM® was controlled with the Ghost API (SensAble Technologies, Inc.) and was only connected to the client workstation. At the beginning of each block of trials, the PHANTOM® was reset at its starting position (i.e., center of the workspace) using wooden jig built specifically for that purpose (see Figure 5).

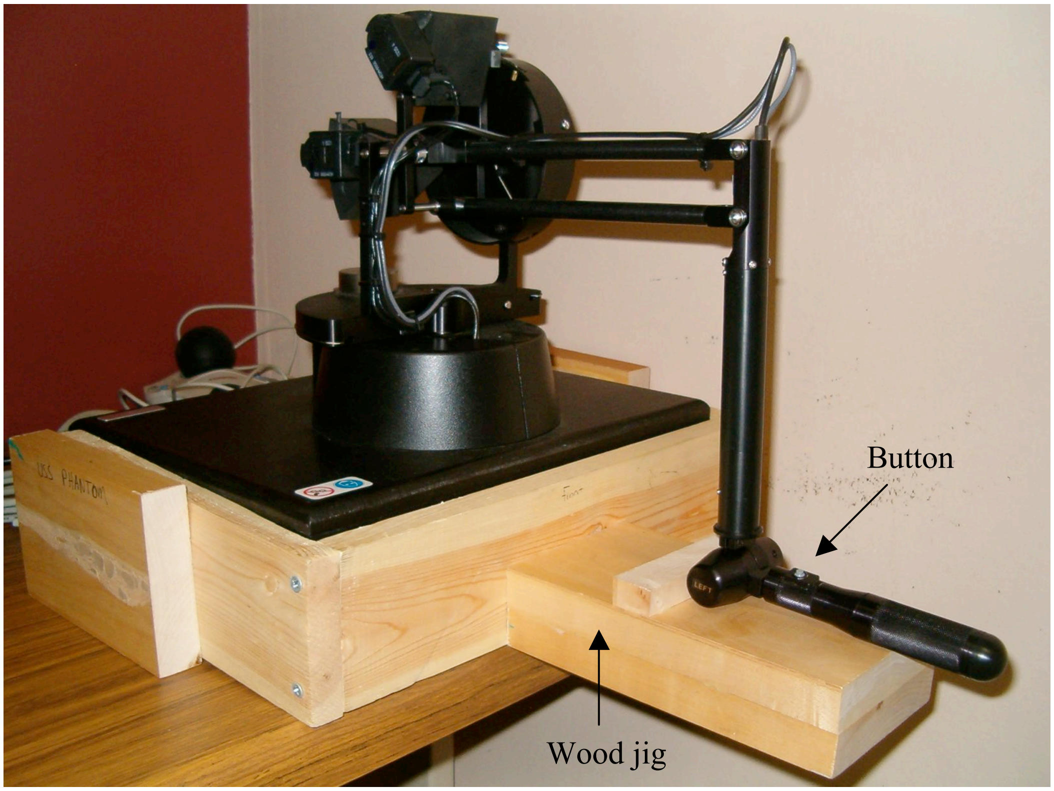

Figure 5. The PHANTOM® device had a button on its handle that participants pressed to select a target. A custom-made wood jig was used to ensure that the reference position of the PHANTOM® (i.e. center of the workspace) was identical across trials. Before the block started, this wood jig was removed.

Virtual Environment (VE)

The VE depicted a 3D cubic workspace of dimensions 0.66 x 0.66 x 0.66 units (the virtual appearance of 1 unit was approximately 1 m in length). The workspace had a room-like appearance with white and dark grey walls and brown hardwood-like floor (see the top left quadrant of Figure 4). These surface-information cues were added to facilitate the perception of depth within the workspace. The top of the cube (ceiling) and the side (wall) facing the participant were transparent to allow the participant to see the interior of the workspace.

Targets appeared sequentially one at a time within the workspace in 1 of 17 possible locations. One possible location was the *origin* (i.e., the center of the workspace) and the remaining 16 locations were located along (invisible) line segments from the origin to each of the 8 corners of the cubic workspace: 8 of these locations were 0.115 units away from the origin and the other 8 were 0.415 units away from the origin. Targets were green semi-translucent spheres (as shown in Figure 4). We used targets with spherical shape, as opposed to the cylindrical targets used by Mateo et al. (2005), in order to avoid some of the issues associated with computing target size for different approach angles (MacKenzie & Buxton, 1992) and because we did not need to accommodate a collaborative task, as did Mateo et al.

A red spherical cursor, 0.01 units in diameter, was always present in the workspace. Its movements were restricted to the interior of the workspace and controlled by the participant through the PHANTOM® device. When the cursor was inside the target area, the color of the target changed from green to red, matching the color of the cursor that penetrated it. The position of the PHANTOM® device was sampled once per frame (frame rate = 30 Hz) and the VE had a minimum feedback delay of 32 ms before any intentional delay was added to the system (due to the time between the moment the PHANTOM® position was read and the moment it was displayed). Therefore, all feedback delays intentionally added to the VE were added to this inherent delay. For example, if a constant 93-ms feedback delay were added to the system, the cursor in the VE would move 125 ms (93 ms + 32 ms) after the corresponding movement was sampled from the PHANTOM®. The feedback-delay values reported in the Design and Procedure section correspond to the feedback delay experienced by the participants and, therefore, they already include the inherent delay of 32 ms. No delay was added to the head-tracking system.

Initially, we considered creating real network delays by physically separating the server and the client computers. Although there are obvious advantages of this approach (i.e., external validity), adopting it would limit our ability to control and manipulate the feedback-delay conditions (e.g., mean feedback-delay and feedback-delay variability). Given that our main goal was understanding how mean-delay and delay-variability levels affected performance and how these variables interacted with each other and with other variables (i.e., target size and distance between targets), we emphasized internal over external validity. For this reason, we used simulated feedback delays that allowed us to manipulate mean delay and delay variability independently, even if mean-delay and delay-variability levels do not change independently of each other in real networks.

Weibull distributions were used to simulate network-delay distributions. Because the values from Weibull distributions are always positive, the shape of the feedback-delay distribution had to change with variability manipulations. For example, the top panel of Figure 6 shows how, as delay variability increases from a standard deviation of 10 ms to a standard deviation of 80 ms, the Weibull distribution (always with M = 93 ms) becomes more positively skewed. There were two main advantages of using Weibull (instead of, e.g., Gaussian) distributions to create variable feedback delays: first, negative values are not possible in Weibull distributions (the same way negative transmission delays are not possible in telerobotic situations) and, second, the changes in distribution shape are similar to changes of delay distributions occurring naturally in real networks, including the Internet (e.g., Phillips & Hernandez, 2004).

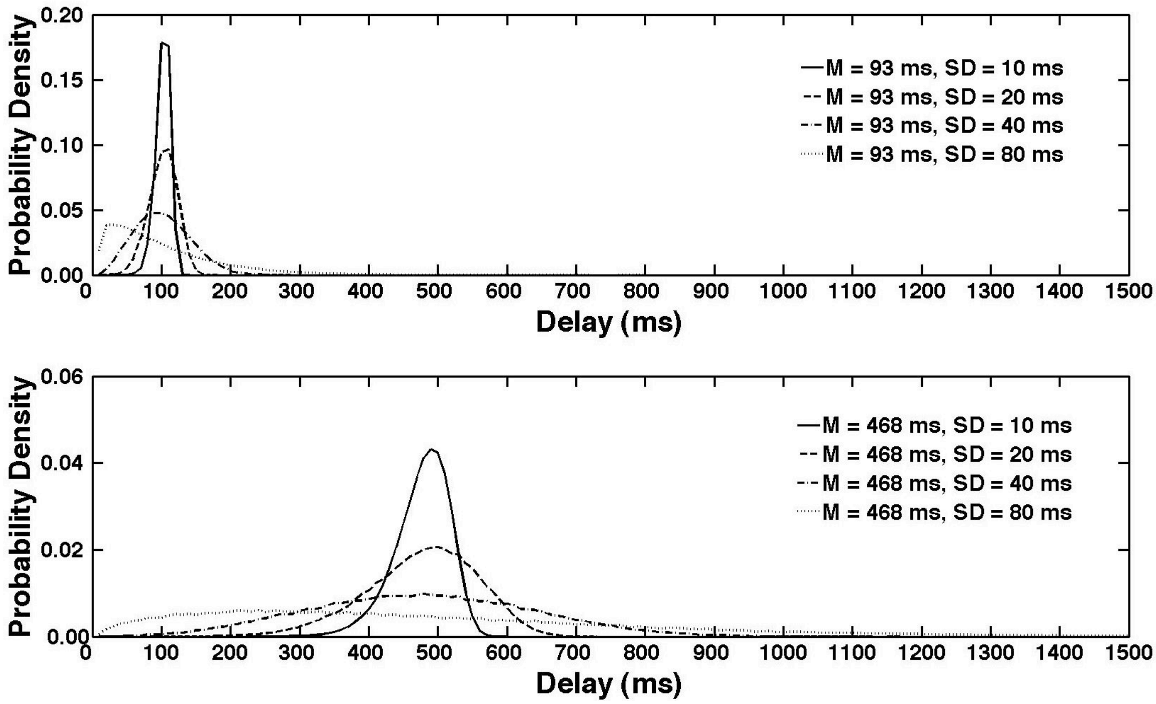

Figure 6. Changes in the shape of Weibull probability density functions as the mean and standard deviation of the feedback-delay distribution were manipulated. Notice that, in our study, distributions were shifted 32 ms due to the inherent delay in the VE; in this case, this resulted in a mean of 125 ms and 500 ms, respectively, and a minimum delay of 32 ms (instead of 0 ms).

In order to simulate feedback delays resulting from real transmission delays, a value (i.e., *delay*) was selected from a Weibull distribution on each frame. This value was used to look up, in a buffer, the position of the PHANTOM® *delay* ms earlier. However, the PHANTOM® position was sampled discretely 30 times per second (e.g., at 33.33 ms, 66.67 ms, 100.00 ms, 133.33 ms, 166.67 ms, etc.), whereas *delay* was continuous (e.g., it could have a value at 40.50 ms). In order to address this issue, the position of the PHANTOM® *delay* ms earlier was estimated by linearly interpolating from the recorded positions surrounding *delay* ms earlier. For example, if *delay* was 40.50 ms and we knew that the position of the PHANTOM® 33.33 ms earlier was 1 and its position 66.67 ms earlier was 2, we expected the position 40.50 ms earlier to be somewhere between 1 and 2.

If we use the approach described above with no further constraints, every time *delay* for the current frame is more than 33.33 ms longer than *delay* for the previous frame, the cursor would be displayed in a position that is “older” than the position previously displayed. In other words, the cursor would have to “jump back” in time. In order to avoid this, we added an additional constraint: if *delay* pointed to a time earlier than the current cursor position, the position of the cursor did not change (i.e., it froze). For example, let us assume that *delay* was 100 ms for frame A and 200 ms for frame B (i.e., the next frame, 33.33 ms later). In frame B, instead of displaying the position of the PHANTOM® 66.67 ms before the position displayed in frame A (i.e., “jumping back” in time), the cursor stayed in the position displayed in frame A (i.e., it froze) until *delay* in a subsequent frame pointed to a more recent position.

Button presses were recorded in the frame in which they occurred, independent of visual feedback. That is, the visual feedback informing participants about the success (or lack of success) of the button press was delayed, but button presses occurred in real time. Therefore, if a participant pressed the button during a cursor freeze and this button press was successful, the consequences of the successful button press (in this case, the beginning of a new trial, as explained in the Task subsection below) were not displayed until the *delay* corresponding to the frame in which the button press occurred had elapsed. This button press would be recorded even if the visual information associated with that frame were never displayed.

As a consequence of freezing the cursor when delays were too long, the feedback delays actually experienced by the user corresponded to a truncated Weibull distribution (see Figure 7). The mean and standard-deviation values presented in this thesis correspond to the experienced (output) delays, not to the input Weibull delays. Nevertheless, as Figure 7 shows, the distribution shape of the output delays is very similar to the shape of another Weibull distribution with the same mean and standard deviation as the output delay distribution.

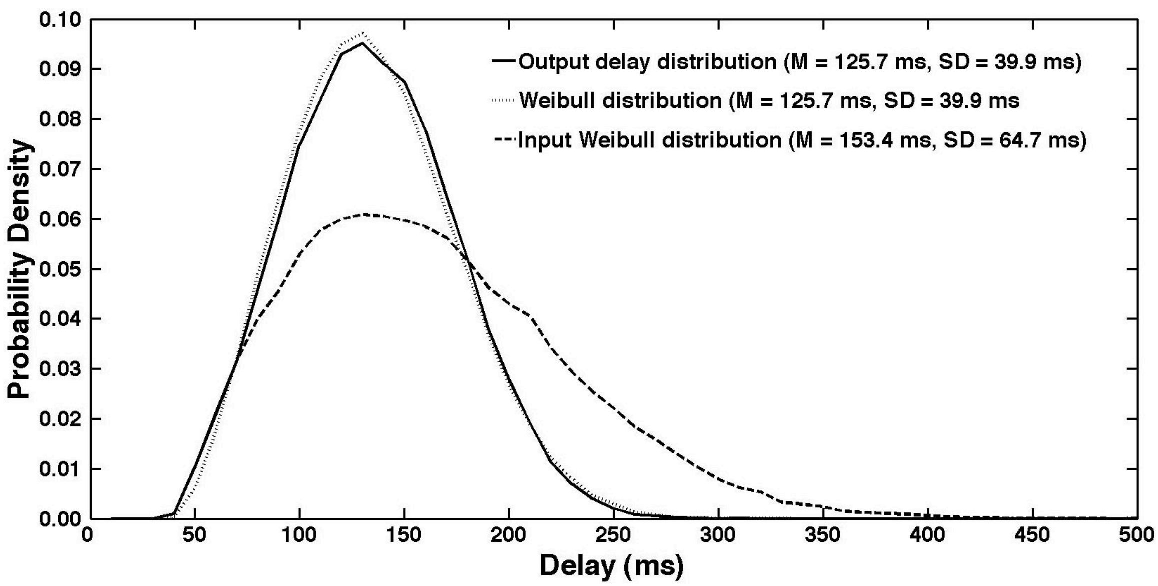

Figure 7. Probability density function of the input Weibull distribution (dashed), the output delay distribution (solid), and another Weibull distribution with the same mean and standard deviation as the output delay distribution (dashed and dotted). This figure illustrates how the computer truncated the original delay distribution by eliminating longest feedback delays. Nevertheless, the resulting (output) delay-distribution shape was still Weibull-like.

Task

Participants were instructed to move the cursor as quickly as possible to the spherical target present in the workspace and press the button while the cursor was within the target area. A successful acquisition was computed when the button was pressed while the cursor was within the target boundaries in real time, independently of what the visual display showed. When a target was successfully acquired, the trial ended. Pressing the button while outside the target area was recorded as an error.

The visual feedback informing the participant about the end of a trial (i.e., disappearance of the target from its location and reappearance at one of the other locations) did not occur immediately after the person successfully acquired the target (i.e., the end of the trial). Rather, this visual feedback occurred *delay* seconds after the end of the trial and marked the beginning of the next trial. The data recorded between the successful button press and the disappearance/reappearance of the target did not belong to either trial and were ignored in our analyses.

Design and Procedure

The 4 independent variables manipulated in the experiment were: *target size* (0.03 and 0.06 units), *distance between targets* (0.115 and .415 units), *mean feedback delay* (*M* =125 and 500 ms), and *feedback-delay variability* (*SD* = 8, 16, 32, and 64 % of the mean). The actual standard-deviation values corresponding to these percentages of the mean were: *SD* = 10, 20, 40, and 80 ms, for the 125-ms mean delay, and *SD* = 40, 80, 160, and 320 ms, for the 500-ms mean delay. In addition, we intended to run a zero-variability condition (i.e., *SD* = 0 ms) but, due to experimenter error, this condition had a longer mean delay than all other conditions. As a consequence, the zero-variability condition was not analyzed in the present study, but it is mentioned in the design description below. Because we expected target size and distance between targets to affect the dependent variables differently under different feedback-delay conditions, they were considered as two separate independent variables in our statistical analyses. However, for comparison purposes, we computed the indexes of difficulty resulting from combining our two sizes and two distances using MacKenzie's (1992) modification, $\mathrm{ID} = \log_2\left(\frac{D}{S}+1\right)$. These indexes of difficulty were: 1.54, 2.27, 2.98, and 3.89 bits.

This 2 x 2 x 2 x 5 design resulted in a total of 40 *conditions* (i.e., combinations of target size, distance between targets, mean feedback delay, and feedback-delay variability). Within each block of trials, mean feedback delay and feedback-delay variability were kept constant and only target size and distance between targets were varied across trials. Each block consisted of a total of 37 trials: 1 *positioning trial*, 4 *practice trials*, and 32 *test trials*, in this order. The target always appeared at the origin at the beginning of each block and the purpose of the initial positioning trial was to make sure that the cursor was near the origin when the first practice trial began. On every even-numbered trial, the target appeared at 1 of the 16 peripheral locations (i.e., requiring a movement *from the origin* to that location). On every odd-numbered trial, the target appeared at the origin (i.e., requiring a movement from the location where the previous trial ended back *to the origin*). Therefore, in two consecutive even and odd trials (e.g., the 2nd and 3rd trials) the participant was required to perform a back-and-forth movement from the origin to a peripheral location and back to the origin. For this thesis, we were only interested in the (less predictable by the participant) origin-to-periphery trials and, therefore, only data from the even trials (i.e., movements from the origin) were analyzed.

The purpose of the 4 consecutive practice trials (i.e., 2 from the origin and 2 to the origin) was to acclimate the participants with the particular condition. In the first practice trial (i.e., 2nd trial of the block and, therefore, an even-numbered trial), the target location was randomly selected from the 16 possible peripheral locations and the target size was randomly selected from the two possible sizes. The same distance and size were kept for the second practice trial back to the center. The target size in the third and fourth practice trials (i.e., 4th and 5th trials of the block) was the same, but different from the target size in the first two practice trials (i.e., 2nd and 3rd trials of the block). The target location for the third practice trial was randomly selected from the 8 possible locations that corresponded to the distance not used for the first two practice trials. After these 4 practice trials, the participant completed 32 test trials using each of the 16 possible peripheral locations once, in random order. After the 37th trial, the workspace disappeared from the HMD, signaling the end of the block. Data from practice trials were not included in the data analysis.

Participants completed a total of 40 blocks grouped in four sequences. Each 10-block sequence contained 1 block for each combination of mean delay and delay variability and the 10 blocks within each sequence were completed in random order (without replacement). In every session, the participants completed a set of 3 blocks, took a 5-min break, and then completed a set of 2 blocks. A 1.5-min break was provided between each of the blocks within each set. Each of the 8 different 5-block sessions took approximately 30 min (they ranged from 20 to 45 min). Participants were allowed to run a maximum of 2 sessions within the same day, as long as they took a minimum break of 45 min between sessions. (Only 1 participant completed two sessions in the same day once.) Each participant took approximately 4 hr to complete the whole experiment (i.e., 8 sessions).

Movement time (MT) was measured from the moment a target appeared in the workspace to the moment the participant acquired it. Therefore, higher values of the MT reflect worse performance. In addition, recording of frame-by-frame data allowed us to divide this total MT into RT, MT_{OL}, and MT_{CL} (as explained in the next section). *Errors* were computed every time the participant pressed the button while outside the target area.

Data Analysis: Dividing Total MT into RT, MT_{OL}, and MT_{CL}

The task participants performed in this experiment can be conceptualized as having three components: a *pre-movement* portion of the trials (i.e., the participant sees the position of the target and plans the initiation of the movement), an *open-loop* (OL) part of the movement (i.e., the participant performs an initial fast distance-covering movement to reach the proximity of the target without using feedback), and a *closed-loop* (CL) part (i.e., the participant performs numerous slower short submovements in the vicinity of the target to "home in" to it, relying heavily on feedback).

Following this conceptual framework, frame-by-frame data were used to divide the total MT for each trial into reaction time (RT, the time elapsed from the appearance of the stimulus to the moment the participant began to move toward the target), movement time to complete the initial OL part of the movement (MT_{OL}), and movement time to complete the CL part of the movement (MT_{CL}).

Although dividing the total MT into RT, MT_{OL}, and MT_{CL} may seem straightforward at first, some of the characteristics of our experimental design made this division somewhat cumbersome. First, participants in our experiment did not stop completely between the end of one trial and the beginning of the next. Rather, the movements occurred serially and the appearance of the next target (i.e., the beginning of the next trial) was determined only by the successful acquisition of the target (i.e., the end of the previous trial). This difficulty would have been eliminated had we used a similar approach to the one used by MacKenzie and Ware (1993). That is, instructing participants to move to the starting position and wait there (without moving) until the next target appears. However, MacKenzie and Ware’s 1D task with a computer monitor and a mouse, instead of a 3D virtual environment with a 3D input device (i.e., PHANTOM®), made it easier for participants to stay still in the starting position (e.g., they could rest their arm). In our experiment, the starting position was in the center of the 3D workspace and the only way to use a similar procedure would have been to restrain their movements while in the starting position (until the target appeared). In order to deal with these difficulties, we developed an automated procedure to systematically identify these break points without the need to hand code each trial. After some trial and error, we developed a procedure that used the kinematic properties of the movement (e.g., velocity and acceleration) and provided, in general, adequate segmentation of data. However, the three dimensionality of the virtual environment made the use of these kinematic properties a little more difficult. For example, if a person performs a fast movement away from the target after it appears, it is unclear whether this indicates that the participant has initiated an intentional movement trying to get to the target and is simply moving the wrong way or that the participant has not yet seen the target.

We initially considered using a simple target-boundary-crossing criterion to determine the break between the end of the OL movement and the beginning of the CL movement. That is, the movement would be considered OL until the person crossed the target boundary for the first time (and CL, thereafter). However, participants almost never crossed the target boundary within what appeared to be the initial OL movement (especially for long-distance trials) and, therefore, this criterion would have included CL movements in the OL part of the movement, overestimating the duration of the latter.

After some exploratory analyses and extensive visual inspection of individual trials, we developed the procedure described in detail below.

Identifying the End of RT and the Beginning of MT_{OL}

Due to our experimental design, participants were not waiting without moving at the beginning of each trial (i.e., during the "pre-movement" part). In contrast, they were attempting to acquire the previous target (and therefore, moving) until they detected the new target and began the intentional movement to the target. So, in order to identify the beginning of the OL movement, we located a period of acceleration toward the target. Frame-by-frame position data were transformed into distance-to-target (DTT) data and velocity toward the target (VTT) was computed from these DTT data and smoothed by averaging with a moving window of five samples (i.e., low-pass filtered). Then, acceleration toward the target (ATT) was computed from the smoothed VTT data. The initial acceleration peak was often (although not always) the largest ATT in the trial, but it always happened when the participant was (relatively) far from the target. Therefore, ATT data were weighted as a function of distance to the target (i.e., farther weighted more than closer), as follows (where *i* is the frame number):

$$\mathrm{ATTw}_i = \begin{cases} \mathrm{ATT}_i * \dfrac{\mathrm{DTT}_i}{\max(\mathrm{DTT})} & \text{if} \quad \dfrac{\mathrm{DTT}_i}{\max(\mathrm{DTT})} > \dfrac{1}{2} \\[2ex] 0 & \text{if} \quad \dfrac{\mathrm{DTT}_i}{\max(\mathrm{DTT})} \leq \dfrac{1}{2} \end{cases}$$

Then, the largest peak in ATTw was identified and frames before this peak acceleration were examined to find the beginning of this acceleration toward the target. The first frame when ATTw no longer increased (i.e., slope of ATTw was zero) was used to signal the beginning of the intentional movement toward the target: the end of the RT and the beginning of the MT_{OL}.

Identifying the End of MT_{OL} and the Beginning of MT_{CL}

The initial OL movement was expected to be fast and long compared to subsequent CL submovements; as soon as the participant started using feedback, slower and shorter movements were expected. Once we had identified the beginning of the initial OL movement toward the target, we used direction-independent kinematic properties of the movement to determine the end of the OL part of the movement and the beginning of the subsequent slower CL submovements. Three-dimensional speed (SPEED) data were computed from the frame-by-frame position data and were smoothed by averaging with a moving window of five samples. The resulting smoothed SPEED data were used to compute 3D acceleration (ACCEL). In order to identify the peak SPEED associated with the initial OL movement, we weighted the smoothed SPEED data by the relative distance to the target (again, farther was weighted more than closer) as follows (where *i* is the frame number):

$$\mathrm{SPEEDw_i} = \begin{cases} \mathrm{SPEED_i} * \dfrac{\mathrm{DTT_i}}{\max(\mathrm{DTT})} & \text{if} \quad \dfrac{\mathrm{DTT_i}}{\max(\mathrm{DTT})} > \dfrac{2}{3} \\ 0 & \text{if} \quad \dfrac{\mathrm{DTT_i}}{\max(\mathrm{DTT})} \leq \dfrac{2}{3} \end{cases}$$

Once the first peak of SPEEDw was identified, indicating the maximum speed of the initial OL movement, we proceeded to find the frame when this initial OL movement ended and feedback started to be used. We used the ACCEL function to search forward (from the fame of the SPEEDw peak) for frames where ACCEL crosses zero (i.e., when the participant stopped decelerating). In order to use feedback for control, the participant needed to move relatively slow, so we picked the first zero-acceleration frame after peak SPEEDw when the SPEED was less than half of the maximum SPEED: $\frac{SPEED_i}{\max(SPEED)} < \frac{1}{2}$. The frame fulfilling all of these criteria was selected as the break point between MT_{OL} and MT_{CL}.

III. RESULTS

The effects of target size, distance between targets, mean feedback delay, and feedback-delay variability were examined in four separate repeated-measures ANOVAs. Specifically, one 2 x 2 x 2 x 4 ANOVA was run for each dependent variable (i.e., RT, MT_{OL}, MT_{CL}, and number of errors). The four levels of feedback-delay variability included in the analysis were standard deviations of 8%, 16%, 32%, and 64% of the mean. As mentioned in the Design and Procedure section, the zero-variability condition was not included in the analysis because, due to implementation error, the mean feedback delay in this condition was always longer than the mean feedback delay in all other delay-variability conditions and would have contaminated the results. Only results showing statistical significance ($\alpha = .05$) are reported in the text below.

Reaction Time (RT)

Only mean feedback delay had a significant effect on RT, $F(1, 5) = 46.89$, $p < .01$. Specifically, RTs were longer in the long-mean-delay condition ($M = 0.369$ s, $SD = 0.093$ s) than in the short-mean-delay condition ($M = 0.295$ s, $SD = 0.053$ s).

Movement Time for the Open-Loop Part of the Movement (MT_{OL})

Both distance between targets, $F(1, 5) = 131.95$, $p < .01$, and mean feedback delay, $F(1, 5) = 11.37$, $p < .05$, had significant effects on MT_{OL}. Specifically, MT_{OL} was longer in the long-distance condition ($M = 1.122$ s, $SD = 0.127$ s) than in the short-distance condition ($M = 0.802$ s, $SD = 0.103$ s) and longer in the long-mean-delay condition ($M = 0.998$ s, $SD = 0.204$ s) than in the short-mean-delay condition ($M = 0.925$ s, $SD = 0.185$ s).

Movement Time for the Closed-Loop Part of the Movement (MT_{CL})

All independent variables: target size, $F(1, 5) = 198.44$, $p < .01$, distance between targets, $F(1, 5) = 131.00$, $p < .01$, mean feedback delay, $F(1, 5) = 236.91$, $p < .01$, and feedback-delay variability, $F(3, 15) = 7.60$, $p < .01$, showed a main effect on MT_{CL}. Smaller sizes ($M = 5.636$ s, $SD = 2.798$ s), longer distances ($M = 5.297$ s, $SD = 2.849$ s), and longer mean delays ($M = 5.978$ s, $SD = 2.591$ s) resulted in longer MT_{CL} than bigger sizes ($M = 2.933$ s, $SD = 1.591$ s), shorter distances ($M = 3.272$ s, $SD = 1.966$ s), and shorter mean delays ($M = 2.591$ s, $SD = 1.241$ s), respectively. In the case of feedback-delay variability, unplanned pairwise comparisons suggested that MT_{CL} in the 64% condition ($M = 4.724$ s, $SD = 2.865$ s) was greater than in the 16% condition ($M = 4.041$ s, $SD = 2.532$ s), $t(5) = 3.65$, $p < .05$, and in the 8% condition ($M = 4.015$ s, $SD = 2.430$ s), $t(5) = 6.22$, $p < .01$, but not significantly different from MT_{CL} in the 32% condition ($M = 4.358$ s, $SD = 2.747$ s), $t(5) = 2.15$, $p = .08$. Actually, MT_{CL} in the 32% condition was not different from MT_{CL} in any of the other three feedback-delay-variability conditions. See Figure 8 for an illustration of how feedback-delay variability affected MT_{CL}.

Mean feedback delay interacted with both target size, $F(1, 5) = 39.67$, $p < .01$, and distance between targets, $F(1, 5) = 26.89$, $p < .01$. Specifically, the effects of both size and distance on MT_{CL} were greater for longer mean feedback-delays than for shorter mean feedback-delays (see Figures 9 and 10). There was also a significant interaction between target size and distance between targets, $F(1, 5) = 48.52$, $p < .01$. Specifically, the effect of size was greater for long-distance trials than for short-distance trials (see Figure 11).

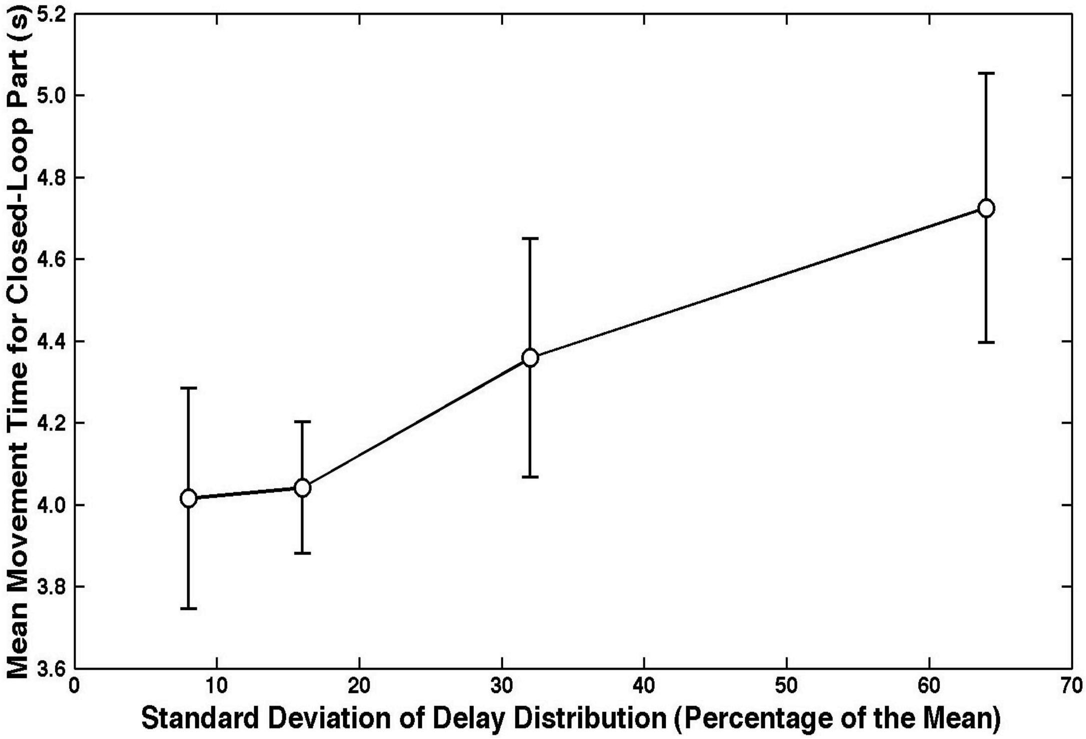

Figure 8. Mean MT_{CL} as a function of feedback-delay variability (*SD* = 8%, 16%, 32%, & 64% of the mean). Error bars show standard errors of the mean.

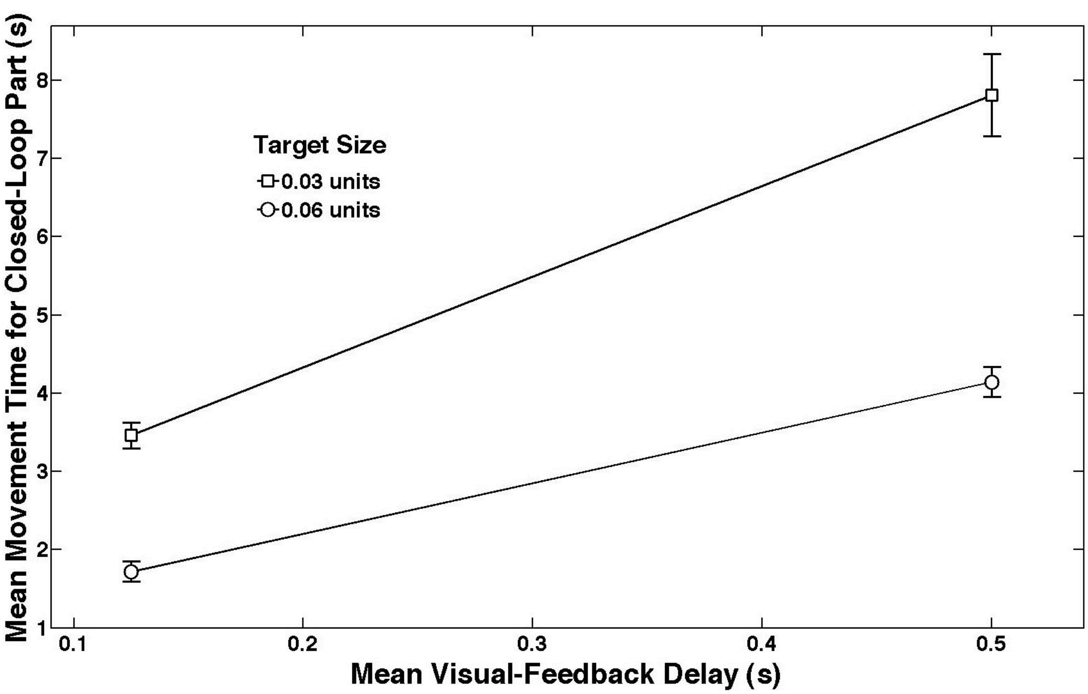

Figure 9. Mean MT_{CL} as a function of mean feedback delay (*M* = 125 ms & 500 ms) for both levels of target size (0.03 & 0.06 units). Error bars show standard errors of the mean.

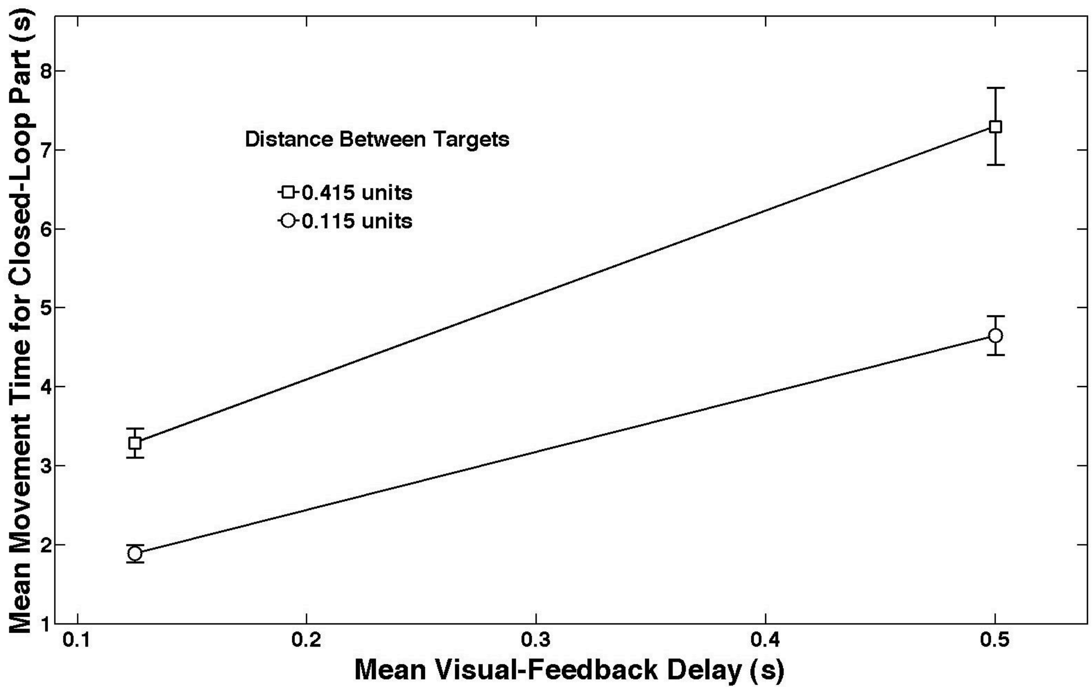

Figure 10. Mean MT_{CL} as a function of mean feedback delay (M = 125 ms & 500 ms) for both levels of distance between targets (0.115 & 0.415 units). Error bars show standard errors of the mean.

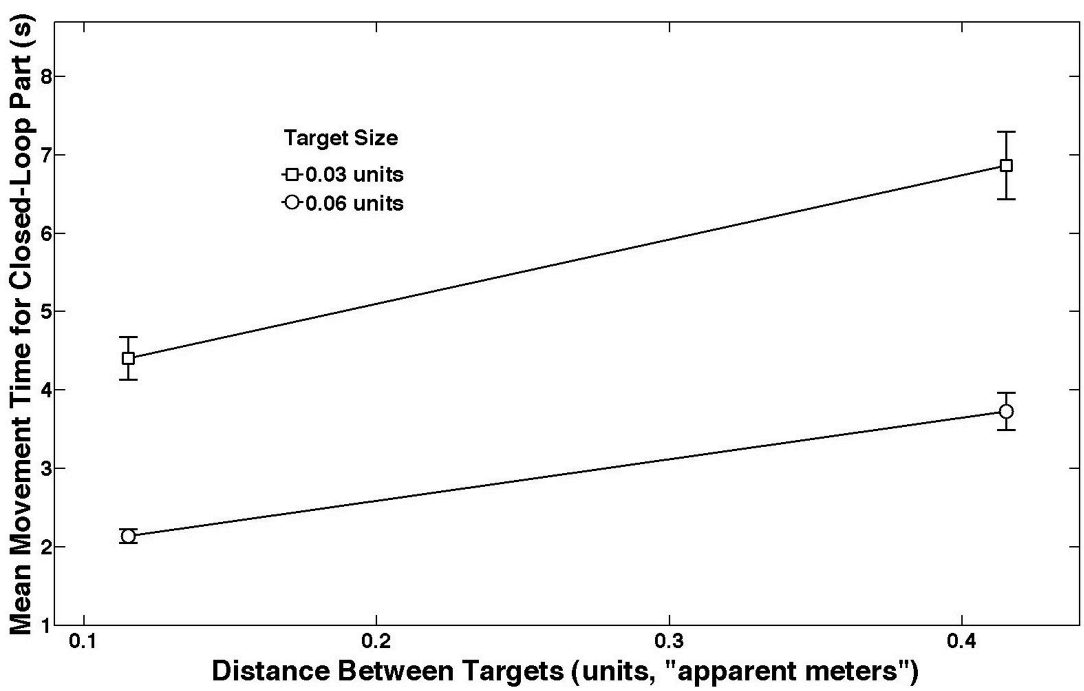

Figure 11. Mean MT_{CL} as a function of distance between targets (0.115 & 0.415 units) for both levels of target size (0.03 & 0.06 units). Error bars show standard errors of the mean.

Number of Errors

Both target size, $F(1, 5) = 83.52$, $p < .01$, and mean feedback delay, $F(1, 5) = 64.29$, $p < .01$, showed a main effect on number of errors. That is, participants committed more errors on trials with smaller target sizes ($M = 0.55$ errors, $SD = 0.26$ errors) than in those with bigger target sizes ($M = 0.16$ errors, $SD = 0.12$ errors) and on trials with longer mean delays ($M = 0.46$ errors, $SD = 0.31$ errors) when compared to those with shorter mean delays ($M = 0.25$ errors, $SD = 0.21$ errors). In addition, the interaction between target size and mean feedback delay was significant, $F(1, 5) = 29.92$, $p < .01$. Specifically, the effect of target size on number of errors was greater for longer mean feedback delays than for shorter mean feedback delays (see Figure 12).

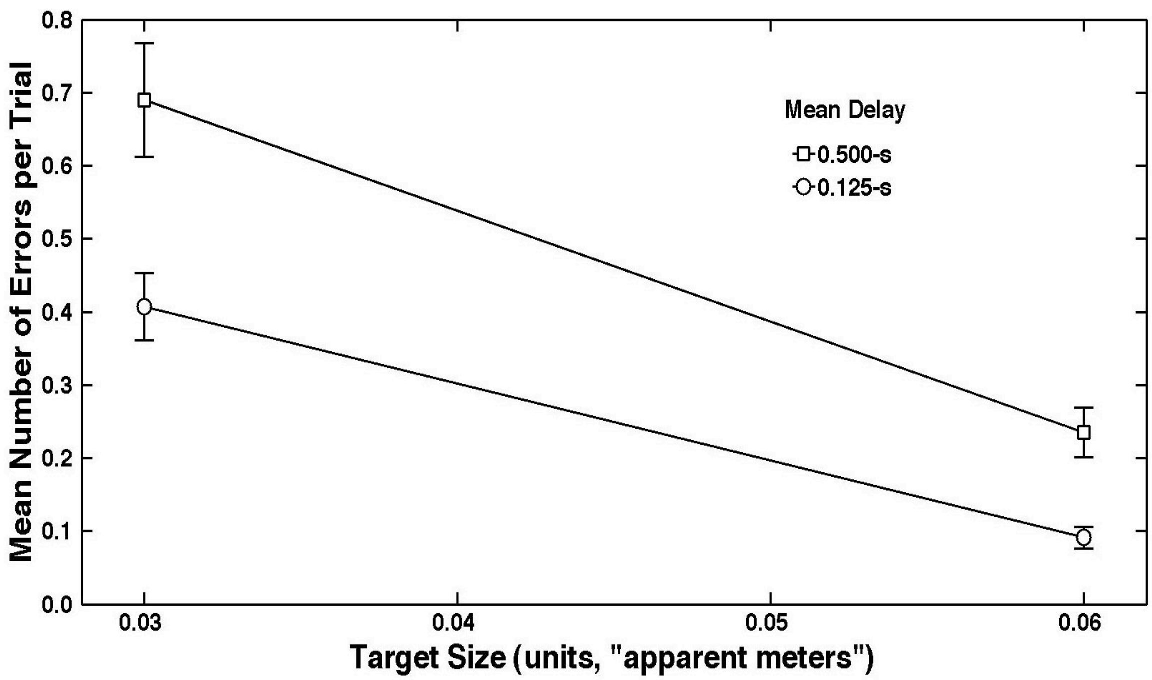

Figure 12. Mean number of errors per trial as a function of target size (0.03 & 0.06 units) for both levels of mean feedback delay (M = 125 & 500 ms). Error bars show standard errors of the mean.

IV. DISCUSSION

The main ideas behind Woodworth's (1899) model, which divided aimed movements into two parts (i.e., initial impulse and current control), are still widely accepted and are the basis for many recent models of human movement (see Elliott, Helsen, & Chua, 2001, for a review). The initial-impulse part of Woodworth's model consists of a rapid OL movement to get to the vicinity of the target, while the current-control part consists of CL movements to "home in" on the target. Our first two predictions were based on this conceptual framework and stated that target size would only affect the CL part of the movement (during which the target boundary is expected to be crossed), while distance between targets would only affect the OL part of the movement (during which most of the distance is expected to be covered). Neither target size nor distance between targets was expected to affect RT (e.g., Fitts & Peterson, 1964).

As expected by our first prediction, smaller target sizes resulted in longer MT_{CL} and more errors than bigger target sizes and changes in target size did not affect RT or MT_{OL}. Our second prediction stated that distance between targets would only affect MT_{OL}, but not RT, MT_{CL}, or number of errors. This second prediction was only partially supported because, in addition to MT_{OL}, distance between targets also affected MT_{CL}. We attribute this unexpected result to our (erroneous) assumption that the initial OL part of the movement would always cover most of the distance and bring the cursor into the vicinity of the target (independent of how far the targets were from each other). That is, we expected the distance left to cover during the CL part of the movement to be small and comparable in short- and long-distance trials. Even though, as expected, the initial OL movement was longer for long-distance trials than for short-distance trials, post-hoc inspection of the data suggested that, at the end of the initial OL movement, the cursor was farther from the target in long-distance trials ($M = 0.138$ units, $SD = 0.031$ units) than in short-distance trials ($M = 0.052$ units, $SD = 0.010$ units), $t(5) = 9.21$, $p < 0.001$ (*one-tailed*). Given that the distance to cover during what we considered the CL part of the movement was consistently longer in long-distance trials than in short-distance trials, it is not surprising that distance between targets had a significant effect on MT_{CL}.

It is important to keep in mind that our task was performed in a 3D workspace in which participants performed 3D oblique movements. Previous research (e.g., Phillips & Triggs, 2001) suggested that 2D diagonal movements (i.e., involving X and Y dimensions) were more problematic to participants than vertical or horizontal movements. By extension, it is reasonable to expect that 3D oblique movements (i.e., involving X, Y, and Z dimensions) might be more difficult than movements involving any one of those dimensions (e.g., directly to the right or directly up). The added difficulty of 3D oblique movements may be affecting the ability of participants to reach the vicinity of the target with the initial OL movement in long-distance trials. However, we intentionally chose these more complicated movements to make them more similar to movements in real teleoperation situation. It is possible that, with more extensive practice (once they become extremely proficient at programming the initial OL movement for long-distance trials), participants may be able to reach the vicinity of the target with the initial OL movement and the effect of distance between targets on MT_{CL} may be reduced or even disappear. Although this possibility is consistent with the motor-learning literature reviewed by Elliot et al. (2001), which suggested that participants improve the accuracy of their initial impulse with practice, further research is necessary to test the validity of this claim.

Given that both target size and distance between targets affected MT_{CL}, we cannot conclude from our analysis that target size was the main task-difficulty parameter affecting the CL part of the movement. However, previous unpublished research in our lab suggested that plotting the results as a function of index of difficulty, $ID = \log_2\left(\frac{D}{S}+1\right)$, can reveal important information about the relative effects of target size and distance between targets on MT.

Given that index of difficulty assumes that size and distance have an equivalent effect on MT (i.e., halving the size results in the same increase of movement time as doubling the distance), a linear fit to the MT data when plotted as a function of index of difficulty would support this assumption. Deviations from this linear fit would provide information about the relative impact of size and distance on MT. Because we only used two levels of size and two levels of distance in our study, changes from the smallest index of difficulty (i.e., ID = 1.54 bits) to the second smallest (i.e., ID = 2.27 bits) and from the second largest index of difficulty (i.e., ID = 2.98 bits) to the largest (i.e., ID = 3.89 bits) corresponded exclusively to changes in target size (i.e., distance between targets was kept constant). Therefore, when the MT increase resulting from a reduction in size (from ID = 1.54 to ID = 2.27 bits and from ID = 2.98 to ID = 3.89 bits) is small in comparison to the MT increase resulting from an increase in distance (from ID = 1.54 to ID = 2.98 bits and from ID = 2.27 bits to ID = 3.89 bits), it indicates the effect on MT is distance-dominated. In contrast, when the MT increase resulting from an increase in distance is small in comparison to the MT_{CL} increase resulting from a reduction in size, it indicates the effect on MT_{CL} is size-dominated.

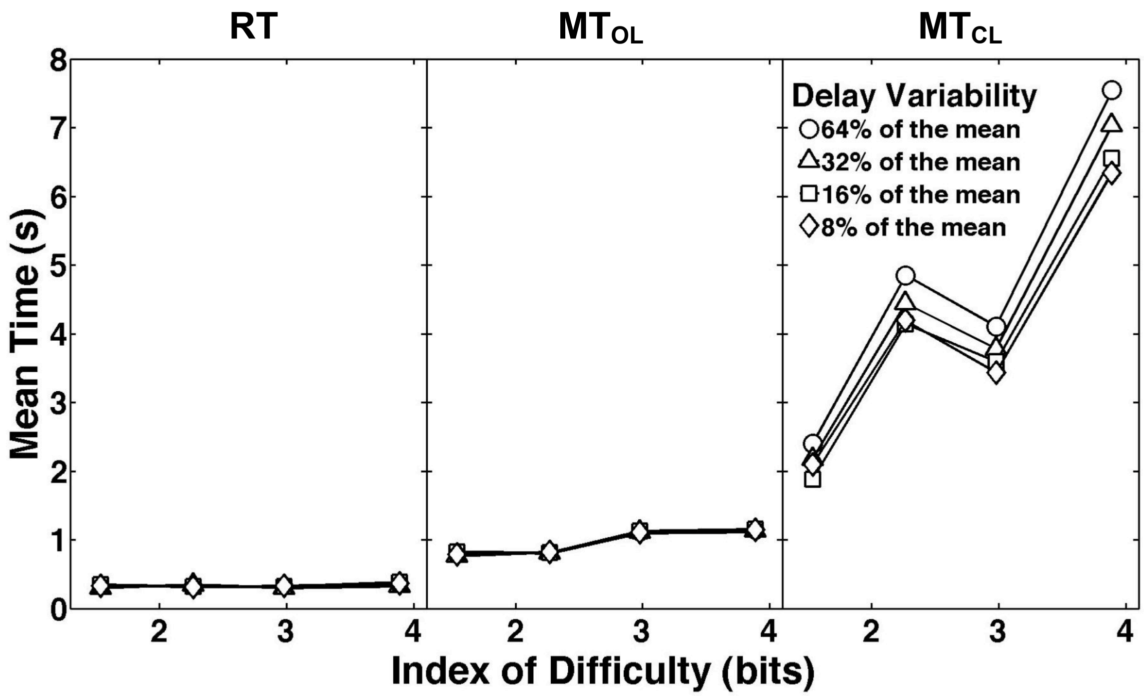

Figure 13. Mean RT (left panel), MT_{OL} (center panel), and MT_{CL} (right panel) as a function of index of difficulty (ID = 1.54, 2.27, 2.98, & 3.89 bits) for all levels of feedback-delay variability (*SD* = 64%, 32%, 16%, & 8% of the mean). The four indexes of difficulty resulted from the combination of: Big size – Short distance (ID = 1.54 bits), Small size – Short distance (ID = 2.27 bits), Big size – Long distance (ID = 2.98 bits), and Small size – Long distance (ID = 3.89 bits).

Looking at the center panel of Figure 13, the steeper lines connecting equal-size data points when compared to equal-distance data points suggest a distance-dominated effect on MT_{OL}. This result is consistent with the result of our ANOVAs. Looking at the right panel of Figure 13, it is obvious that the pattern in the data is quite different, showing steeper slopes in lines connecting equal-distance data points than in lines connecting equal-size data points and, thus, suggesting a size-dominated effect on MT_{CL}. The ANOVAs did not allow us to conclude that this was the case, but looking at Figure 13 clearly suggests that there was a change from a distance-dominated MT_{OL} (center panel) to a size-dominated MT_{CL} (right panel), as shown by the different deviations from linearity in each panel.

In terms of feedback-delay manipulations, we hypothesized (i.e., predictions 3 and 4) that both mean feedback delay and feedback-delay variability would affect only the CL part of the movement (i.e., MT_{CL} and number of errors), but not the OL part (i.e., RT and MT_{OL}). These predictions were only partially supported by our results. Although mean feedback delay affected MT_{CL} and number of errors, as expected, it also showed significant (and unexpected) effects on both RT and MT_{OL}. In the case of feedback-delay variability, we found an effect on MT_{CL} and no effect on RT and MT_{OL}, as expected. However, we did not find the effect of feedback-delay variability on number of errors that we expected.

It is unclear why mean feedback delay affected the OL part of the movement. The increase in MT_{OL} with longer mean feedback delays could have resulted from participants moving slower overall with increasing delays, from participants carrying out longer OL movements with increasing delays, or from a combination of both. The first possibility (i.e., slower movements) is based on the idea that, given that the feedback-delay conditions are kept constant within blocks, participants may have moved slower (i.e., lowered their gain, in Control-theory terms) throughout long-feedback-delay blocks to adapt to the feedback-delay conditions, even when they were not using visual feedback. A post-hoc examination revealed that the peak speeds (during the OL part of the movement) in trials with long mean feedback delays (M = 0.326 units/frame, SD = 0.039 units/frame) were indeed slower than the peak speeds in trials with short mean feedback delays (M = 0.367 units/frame, SD = 0.030 units/frame), $t(5) = 3.90$, $p < 0.01$ (*one-tailed*). This post-hoc examination is consistent with the first explanation: if the first OL movement is slow in trials with longer feedback delays, greater MT_{OL} are expected for longer feedback delays even if no feedback is used during this part of the movement. Another post-hoc analysis was conducted to investigate whether the distance covered during the OL part of the movement may have affected the longer MT_{OL} in blocks with longer mean feedback delays. Our analysis showed that the length of the initial OL movement was longer in short-delay trials (M = 1.65 units, SD = 0.23 units) than in long-delay trials (M = 1.40 units, SD = 0.21 units), $t(5) = 4.80$, $p < 0.01$ (*two-tailed*). This result is incompatible with the possibility that longer MT_{OL} in long-delay trials were the result of longer OL movements. Therefore, these post-hoc examinations are consistent with the idea that participants moved slower overall during longer feedback-delay blocks. It may be interesting to perform more research to further explore the validity of this explanation.

In terms of feedback-delay variability, our analyses did not allow us to compare the four different levels of feedback-delay variability to a zero-variability condition. We could only compare the four levels of feedback-delay variability to each other. These unplanned pairwise comparisons suggest that only the largest delay-variability level (i.e., SD = 64% of the mean delay) differed, in terms of their effect on MT_{CL}, from the smallest delay-variability level (i.e., 8% of the mean delay). Unfortunately, no conclusions can be drawn from our results about which level of feedback-delay variability, if any, affected MT_{CL} above and beyond the effect of mean feedback delay. However, if we assume that MT_{CL} in the zero-variability condition was equal to or shorter than MT_{CL} in the smallest delay-variability level (i.e., if we assume that delay variability did not help performance), then our results suggest that at least the largest delay-variability condition resulted in longer MT_{CL} than a constant feedback delay of the same mean. In spite of the rational appeal of this claim, future research needs to test its validity.

We also hypothesized (i.e., prediction 5) that an increase in mean feedback delay would result in greater increases of MT_{CL} and number of errors than an equivalent increase of feedback-delay variability. This prediction was supported by our results. Actually, a four-fold increase of mean feedback delay (from M = 125 to 500 ms) resulted in a greater average increase of MT_{CL} (ΔM = 3.387 s) than an eight-fold increase of feedback-delay variability from a standard deviation of 8% of the mean to a standard deviation of 64% of the mean (ΔM = 0.709 s), $t(5) = 16.34$, $p < .00001$. Consistent with our expectations, the same four-fold increase of mean feedback delay resulted in a greater increase of number of errors (ΔM = 0.21 errors) than an eight-fold increase of feedback-delay variability (ΔM = -0.01 errors), $t(5) = 4.20$, $p < .01$.

Given that the effects of mean feedback delay are much greater than the effects of feedback-delay variability, it seems unwise to employ strategies that minimize feedback-delay variability at the expense of increasing mean feedback delay. That is, strategies such as the buffering technique proposed by Kosuge and Murayama (1998) to deal with haptic-feedback delays should not be adopted when dealing exclusively with visual-feedback delays. This study did not explore haptic-feedback delays and, therefore, our results do not directly apply to techniques used to ameliorate the effect of variable haptic-feedback delays. In addition, our findings may not apply to tasks that differ significantly from the visual target-acquisition task used in this experiment. For example, the effect of feedback-delay variability on tasks that depend more on closed-loop feedback, such as visual tracking, may be greater than the effect found with our visual target-acquisition task.

The results from our analyses also suggest that the difference between the effects of mean feedback delay and feedback-delay variability may not be only quantitative but also qualitative in nature. That is, not only does increasing mean feedback delay result in longer MTs and more errors than feedback-delay variability (i.e., greater effect: quantitative), but mean feedback delay may be affecting measures that feedback-delay variability does not affect (i.e., different pattern of effects: qualitative). As shown in our results, mean feedback delay affected all dependent measures independent of whether or not we had predicted them to depend on feedback, whereas feedback-delay variability affected only one of the two dependent measures expected to depend heavily on feedback (i.e., MT_{CL}). One way to explore whether the difference between the effect of mean feedback delay and feedback-delay variability is purely quantitative would be to use higher levels of feedback-delay variability to explore how large feedback-delay variability affects dependent measures. If extremely high feedback-delay-variability levels only affect MT_{CL}, it is possible that the difference between the effects of these two variables is qualitative in nature. If using higher levels of feedback-delay variability results in increases in all dependent measures, it is possible that the difference between the effects of mean feedback delay and feedback-delay variability is merely quantitative.

Prediction 6 stated that an interaction between mean feedback delay and feedback-delay variability was expected, but our results did not support this prediction. That is, neither the effect of feedback-delay variability on MT_{CL} nor its effect on number of errors was affected by changes in mean feedback delay. We tested this prediction using delay-variability levels in units of percentage of the mean (e.g., *SD* = 32% of 125 ms and *SD* = 32% of 500 ms) instead of absolute magnitude. In spite of the fact that this choice would have made our prediction more likely to be supported, we did not find supportive results. Rather, our results strongly suggest that our prediction was wrong and that the effect of feedback-delay variability on performance did not change as a function of mean feedback delay. The lack of interaction between mean feedback delay and feedback-delay variability, when standard deviations are measured in units of percentage of the mean feedback delay, is consistent with Watson et al.'s (1998) recommendation to use percentage-of-the-mean units rather than absolute-magnitude units to manipulate feedback-delay variability. In addition, this lack of interaction suggests that, if a cutoff feedback-delay-variability level is to be identified above which performance begins to suffer, manipulating feedback-delay-variability in units of percentage of the mean feedback delay would be appropriate.

Studies using Fitts' law in the presence of feedback delays (e.g., MacKenzie & Ware, 1993) assume that the relative impact of target size and distance between targets on MT is the same across feedback delays. However, some studies (e.g., So & Chung, 2002) have suggested that, with increasing mean feedback delay, the effect of target size on MT may become greater than the effect of distance between targets. The rationale behind So and Chung's hypothesis (based on Woodworth's model) was that increasing feedback delay would affect the CL part of the movement and, as a result, the effect of target size (but not the effect of distance) on MT would also increase. So and Chung's results supported their hypothesis and our seventh prediction built on their results. In addition, our division of MT into MT_{OL} and MT_{CL} improves our ability to test the theoretical predictions.

Our results showed that longer mean feedback delays resulted in a greater effect of target size on MT_{CL} and number of errors than shorter mean feedback delays, supporting prediction 7. The rationale behind this prediction (as described in the previous paragraph) also assumed that the interaction between distance and mean feedback delay would not be significant. In the case of number of errors, our results fulfilled this condition, supporting both our rationale and So and Chung's (2002) findings (even though they did not directly explore number of errors).

However, in the case of MT_{CL} the interaction between distance and mean feedback delay was significant, showing a greater effect of distance between targets for longer mean feedback delays than for shorter ones. Our results are inconsistent with the idea that just the effect of target size (but not the effect of distance between targets) on MT_{CL} is affected by feedback-delay manipulations. We attribute this unexpected result to the same source as the unexpected main effect of distance between targets on MT_{CL} (i.e., our criterion to divide OL and CL parts of the movement). That is, given that the distance left to the target at the end of the initial OL movement was greater in long-distance trials than in short-distance trials, it is less surprising that it took longer to cover this distance in the presence of longer feedback delays.

Notice, in the right panel of Figure 14, how the slope changes between first-second and third-fourth data points (i.e., change in target size) and between the first-third and second-fourth data points (i.e., change in distance between targets) as a function of mean feedback delay. In spite of having only partially supportive statistical results, the trend in

Figure 14 seems consistent with our prediction (i.e., the slope linking different-size points is steeper for longer mean feedback delays than for shorter ones).

Prediction 8 followed the same rationale as the previous prediction, given that both mean feedback delay and feedback-delay variability are feedback manipulations. Thus, we expected the effect of target size on MT_{CL} and on number of errors to be greater for higher levels of feedback-delay variability than for lower ones. Our results did not support this prediction, suggesting that feedback-delay variability does not interact with the effect of target size on MT_{CL} and number of errors.

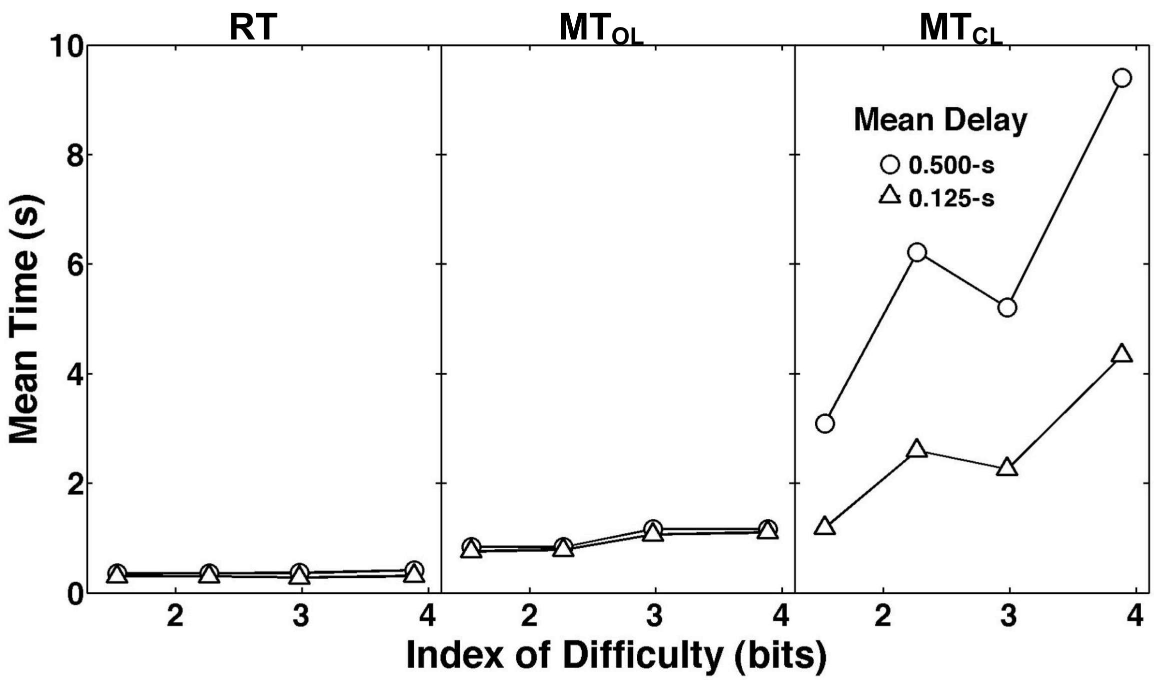

Figure 14. Mean RT, MT_{OL}, and MT_{CL} as a function of index of difficulty (ID = 1.54, 2.27, 2.98, & 3.89 bits) for both levels of mean feedback delay (M = 0.500 ms & 0.125 ms). The four indexes of difficulty resulted from the combination of: Big size – Short distance (ID = 1.54 bits), Small size – Short distance (ID = 2.27 bits), Big size – Long distance (ID = 2.98 bits), and Small size – Long distance (ID = 3.89 bits).

Summary of Results in Relation to Our Predictions

Overall, we found at least partial support for most of our predictions. As expected, target size affected only the CL part of the movement (i.e., MT_{CL} and number of errors) and distance between targets affected the OL part of the movement (i.e., MT_{OL}). However, our results also showed an unexpected effect of distance on MT_{CL}. This effect is attributed to our conceptualization of the movement as a single OL movement that brings the cursor into the vicinity of the target followed by a series of CL submovements to acquire it. This description may be accurate in short-distance trials but, in long-distance trials, what we considered the initial OL movement usually ends far away from the target and, arguably, a series of these distance-covering OL movements (instead of a single one) are required to reach the vicinity of the target. In our study, every movement after the initial one is considered CL and, therefore, distance had a significant effect on MT_{CL}.

Mean feedback delay affected all dependent variables, even those we had not expected to rely on feedback *a priori*. One possibility is that participants compensated for longer mean feedback delays by lowering their gain (i.e., moving slower) and this strategy may have affected other parts of the movement in subsequent trials. As predicted, feedback-delay variability affected one of the two dependent variables that were expected to rely on feedback (i.e., MT_{CL}) and did not affect the dependent variable that was expected to be independent of feedback (i.e., MT_{OL}). However, we did not find an effect of feedback-delay variability on the other dependent variable expected to rely on feedback (i.e., number of errors).

As predicted, the effect of mean feedback delays was much greater than the effect of feedback-delay variability on MT_{CL} and number of errors. The fact that mean feedback delay also affected dependent measures that were supposedly independent from feedback raises the possibility that the effects of mean feedback delay and feedback-delay variability may differ qualitatively, as well as quantitatively.

As predicted, the effect of target size on MT_{CL} and number of errors was greater for longer mean feedback delays. However, the effect of distance on MT_{CL} was also greater for longer feedback delays, making it more difficult to conclude that the relative impact of target size on MT_{CL} (compared to the impact of distance) becomes greater with increasing mean feedback delay. We attribute this unexpected result to the same source as our unexpected effect of distance on MT_{CL}. That is, we assumed that a single OL movement would always bring the cursor to the vicinity of the target when, in fact, this was not always the case.

Predictions addressing interactions between feedback-delay variability and mean feedback delay, as well as between feedback-delay variability and size, were not supported. The fact that feedback-delay variability did not interact with mean feedback delay suggests that it may be possible to find a level of feedback-delay variability (i.e., a certain percentage of mean feedback delay), independent of mean feedback delay, at which performance begins to deteriorate. However, more research is necessary to determine if this claim is true and to identify what these threshold levels are.

Other Issues to Consider When Studying Variable Feedback Delays

Although these data provide an initial picture of the effect of feedback-delay variability on human motor performance, many questions remain unanswered. One of the reasons for this is that feedback-delay variability is a complicated variable and other factors (in addition to the magnitude of the standard deviation) may affect how it impacts performance. These factors include: the skewness of the delay distribution, the frequency of variation, the level of predictability of the variation, and the number of cursor freezes resulting from the delay variability.

Watson et al. (1998) mentioned that the skewness of the delay distribution might have an effect on how much feedback-delay variability affected performance. Specifically, they suggested that positively skewed distributions might be less detrimental to performance than negatively skewed distributions. In our study, we chose feedback-delay distributions that were positively skewed or approximately symmetrical (i.e., Weibull) for practical reasons. That is, we were interested in network delays that tend to have these properties. However, how distribution skewness affects motor performance should be studied further in order to better understand the impact of feedback-delay variability on performance. For example, a study comparing the effects of symmetrical, positively skewed, and negatively skewed feedback-delay distributions (with the same mean and standard deviation) on performance could provide useful information.

In our study, feedback delay was varied on every frame (i.e., the frequency of variation was high: 30 times per second). However, feedback delay may be constrained to vary less frequently (i.e., show a low-pass characteristic) in some applications. Informal observations during our pilot study suggested that, when the frequency of variation is very low, performance is not affected much (if at all) compared to a constant feedback delay. This finding is not surprising, given that as frequency of variation decreases, the condition becomes closer to a constant feedback delay.

Another factor to consider when discussing the effects of feedback-delay variability on performance is the level of predictability of the feedback delays. For example, different results may be expected if feedback delay is manipulated in an orderly manner (e.g., from shorter to longer, Vaghi et al., 1999), in a periodically varying manner (e.g., following a sinusoidal pattern), or in a randomly varying manner (e.g., using a real or simulated network). Although we are not aware of any study directly exploring these issues in this context, it seems reasonable to expect that less predictability would result in greater performance decrements than more predictability.

One last factor that may affect the effect of feedback-delay variability is how often it results in a cursor freeze. If cursor freezes have an impact on performance, then it is possible that number of freezes may be a better description of feedback-delay variability (in terms of its effect on human performance) than the standard deviation of the distribution of feedback delays. Future research should explore this possibility in order to identify the best way to quantify feedback-delay variability when interested in its effects on motor performance.

V. CONCLUSIONS

Feedback-delay variability can have a detrimental effect on performance, even if its effects are much weaker than those of mean feedback delay. More research is necessary to determine exactly when feedback-delay variability starts to affect performance, but our data clearly suggest that any technique that increases mean feedback delay in order to reduce variability (e.g., Lane et al., 2002) is likely to be counterproductive. In addition, our results suggest that target size is a critical variable when performing motor tasks in the presence of feedback delays and that objects that may need to be manipulated remotely should be designed accordingly (i.e., no small switches). In situations where the size of the objects to be manipulated cannot be changed (e.g., telesurgery), research should be performed to develop a display-and-control system that facilitates performance. Finally, future research is needed to explore how other variables (e.g., frequency of variation, skewness of the delay distribution) may affect the impact of feedback-delay variability on performance.

REFERENCES

Anderson, R. J., & Spong, M. (1988). Bilateral control of teleoperators with time delay. *Proceedings of the 27th Conference on Decision and Control, USA*, 167-173.

Black, J. H. (1971). *Factorial study of remote manipulation with transmission delay*. Unpublished master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

Bergeijk, W. A. van, & David, E. E. (1959). Delayed handwriting. *Perceptual and Motor Skills, 9*, 347-357.

Chopra, N., Spong, M. W., Hirche, S., & Buss, M. (2003). Bilateral teleoperation over the Internet: The time varying delay problem. *Proceedings of the American Control Conference, USA, 1*, 155-160.

Craik, K. J. W. (1947). Theory of the human operator in control systems. I. The operator as an engineering system. *British Journal of Psychology, 38*, 56-61.

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.

Elliott, D., Helsen, W. F., & Chua, R. (2001). A century later: Woodworth's (1899) two-component model of goal-directed aiming. *Psychological Bulletin, 127*, 342-357.

Ferrell, W. R. (1965). Remote manipulation with transmission delay. *IEEE Transactions on Human Factors in Electronics, 6*, 24-31.

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

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

Goldberg, K., Gentner, S., Sutter, C., Wiegley, J., & Farzin, B. (2002). The Mercury project: A feasibility study for online robots. In K. Goldberg & R. Siegwart (Eds.), *Beyond webcams. An introduction to online robots* (pp. 17-36). Cambridge, MA: The MIT Press.

Goldberg, K., & Siegwart, R. (Eds.) (2002). *Beyond webcams. An introduction to online robots*. Cambridge, MA: The MIT Press.

Gutwin, C. (2001). Effects of network delay on groupwork in shared workspaces. *Proceedings of the Seventh European Conference on Computer Supported Cooperative Work (ECSCW'01), Germany*, 299-318.

Gutwin, C., Benford, S., Dyck, J., Fraser, M., Vaghi, I., & Greenhalgh, C. (2004). Revealing delay in collaborative environments. *Proceedings of ACM Conference on Computer-Human Interaction (CHI'04), Austria*, 503-510.

Hanly, E. J., & Broderick, T. J. (2005). Telerobotic surgery. *Operative Techniques in General Surgery, 7*, 170-181.

Hoffmann, E. R. (1992). Fitts' law with transmission delay. *Ergonomics, 35*, 37-48.

Hu, H., Yu, L., Tsui, P. W., & Zhou, Q. (2001). Internet-based robotic systems for teleoperation. *Assembly Automation, 21*, 143-151.

Jagacinski, R. J. (1977). A qualitative look at feedback control theory as a style of describing behavior. *Human Factors, 19*, 331-347.

Jagacinski, R. J., & Flach, J. M. (2003). *Control theory for humans. Quantitative approaches to modeling performance*. Mahwah, NJ: Lawrence Erlbaum Associates.

Jay, C., & Hubbold, R. (2005). Delayed visual and haptic feedback in a reciprocal tapping task. *Proceedings of the First Joint Eurohaptics Conference and Symposium* *on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Italy*, 655-656.

Kosuge, K., & Murayama, H. (1998). Teleoperation via computer network. *Electrical Engineering in Japan, 124*, 49-56.

Lane, J. C., Carignan, C. R., & Akin, D. L. (2000). Time delay and communication bandwidth limitation on telerobotic control. *Proceedings of Mobile Robots XV and Telemanipulative and Telepresence Technologies VII, USA*, 405-419.

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*, 113-128.

Lee, B. S. (1950). Effects of delayed speech feedback. *Journal of the Acoustical Society of America, 22*, 824-826.

MacKenzie, I. S., & Buxton, W. (1992). Extending Fitts' law to two-dimensional tasks. *Proceedings of the ACM Conference in Human Factors in Computing Systems – CHI '92, USA*, 219-226.

MacKenzie, I. S., & Ware, C. (1993). Lag as a determinant of human performance in interactive systems. *INTERCHI'93, USA*, 488-493.

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

Mateo, J. C., Manning, J. T., Cowgill, J. L., Simpson, B. D., Moore, T. J., Weisenberger, J. H., & Gilkey, R. H. (2005). Evaluation of a collaborative movement task in a distributed three-dimensional virtual environment. *Proceedings of the HFES 49th Annual Meeting, USA*, 1578-1582.

Oboe, R., & Fiorini, P. (1997). Issues on internet-based teleoperation. *Proceedings of the Fifth IFAC Symposium on Robot Control, France*, 611-617.

Park, K. S., & Kenyon, R. V. (1999). Effects of network characteristics on human performance in a collaborative virtual environment. *Proceedings of IEEE VR'99, USA*, 104-111.

Phillips, I. W., & Hernandez, J. A. (2004). *Internet performance, measuring, and modelling*. Retrieved May 4, 2005, from: http://www.ripe.net/ripe/meetings/ripe-49/presentations/ripe49-tt-weibull.pdf

Phillips, J. G., & Triggs, T. J. (2001). Characteristics of cursor trajectories controlled by the computer mouse. *Ergonomics, 44*, 527-536.

Shannon, C. E., & Weaver, W. (1963). *The mathematical theory of communication*. Urbana, IL: University of Illinois Press. (Original work published 1948)

Sheik-Nainar, M. A., Kaber, D. B., & Chow, M. (2005). Control gain adaptation in virtual reality mediated human-telerobot interaction. *Human Factors and Ergonomics in Manufacturing, 15*, 259-274.

Sheridan, T. B., & Ferrell, W. R. (1974). *Man-machine systems: Information, control, and decision models of human performance*. Cambridge, MA: The MIT Press.

Sheridan, T. B. (1992). *Telerobotics, automation, and human supervisory control*. Cambridge, MA: The MIT Press.

Smith, W. M., McCrary, J. W., & Smith, K. U. (1960). Delayed visual feedback and behavior. *Science, 132* (3433), 1013-1014.

So, R. H. Y., & Chung, G. K. M. (2002). Combined and interacting effects of hand and head movement delays on discrete manual performance in a virtual environment. *Ergonomics, 45*, 105-123.

Vaghi, I., Greenhalgh, C., & Benford, S. (1999). Coping with inconsistency due to network delays in collaborative virtual environments. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, UK*, 42-49.

Watson, B., Spaulding, V., Walker, N., & Ribarsky, W. (1997). Evaluation of the effects of frame time variation on VR task performance. *Proceedings of the IEEE Virtual Reality Annual Symposium, USA*, 38-44.

Watson, B., Walker, N., Ribarsky, W., & Spaulding, V. (1998). Effects of variation in system responsiveness on user performance in virtual environments. *Human Factors, 40*, 403-414.

Welford, A. T. (1960). The measurement of sensory-motor performance: Survey and reappraisal of twelve years' progress. *Ergonomics, 3*, 189-230.

Woodworth, R. S. (1899). The accuracy of voluntary movements. *Psychological Review, 3*(2).