

# Modelling the Effects of Delayed Haptic and Visual Feedback in a Collaborative Virtual Environment

CAROLINE JAY  
MASHHUDA GLENCROSS  
and  
ROGER HUBBOLD  
University of Manchester<sup>1</sup>

---

Collaborative virtual environments (CVEs) enable two or more people, separated in the real world, to share the same virtual 'space'. They can be used for many purposes, from teleconferencing to training people to perform assembly tasks. Unfortunately, the effectiveness of CVEs is compromised by one major problem: the delay that exists in the networks linking users together. Whilst we have a good understanding, especially in the visual modality, of how users are affected by delayed feedback from their own actions, little research has systematically examined how users are affected by delayed feedback from other people, particularly in environments that support haptic (force) feedback. The current study addresses this issue by quantifying how increasing levels of latency affect visual and haptic feedback in a collaborative target acquisition task. Our results demonstrate that haptic feedback in particular is very sensitive to low levels of delay. Whilst latency affects visual feedback from 50 msec, it impacts on haptic task performance 25 ms earlier, and causes haptic measures of performance deterioration to rise far more steeply than visual. The 'Impact-Perceive-Adapt' model of user performance, which considers the interaction between performance measures, perception of latency and the breakdown of the perception of immediate causality, is proposed as an explanation for the observed pattern of performance.

Categories and Subject Descriptors: H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Theory and methods*; *Haptic I/O*.

General Terms: Human Factors, Measurement, Performance

Additional Key Words and Phrases: Haptics, latency, virtual environments, distributed collaboration

---

## 1. INTRODUCTION

Collaborative Virtual Environments (CVEs) offer significant potential for geographically distant participants to work towards achieving a shared objective in a variety of diverse application areas including telesurgery, computer aided design and training. However, meaningful collaboration in virtual environments can be seriously limited by the characteristics of the network used for sharing data between different sites. If users are geographically remote from one another, then data travelling between them may be subject to considerable latency.

Improvements to network design can reduce latency, but not sufficiently for many long distance Internet links. A fundamental restriction on the speed of data transfer is imposed by the time it takes the electrical signal to propagate through an appropriate transfer medium. Add to this bandwidth constraints and processing time at either end, and the minimum theoretical latency we could expect for data travelling between the United Kingdom and the East coast of the United States, for example, is at best eighty, and more usually over a hundred milliseconds.

---

<sup>1</sup> This research was supported by the UK Engineering and Physical Sciences Research Council, grant number GR/523087/01.

Authors' address: Advanced Interfaces Group, School of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL; email caroline.jay@manchester.ac.uk.

In this paper we present a detailed study of the effects of network communication induced latency in haptic and visual feedback from remote participants. The study considers the effects of between 25 and 400 msec of additional end-to-end latency (levels much lower than those considered in previous studies) on a task which relies on continuous haptic and visual exchange between participants to acquire a target. We analyze the effects quantitatively in terms of error rates, task completion times and users' perceptions of the difficulty of the task. To our knowledge no similarly comprehensive study has been conducted; the most relevant literature, which includes research conducted in single user environments, is discussed below. An important result of our study is the 'Impact-Perceive-Adapt' model of user behaviour, which describes how the perception of latency, and the breakdown in the perception of immediate causality, cause movement times to increase in direct proportion to latency, but error rates to rise in a non-linear fashion.

### 1.1 Delayed feedback in single-user virtual environments

Psychologists have long been investigating the effects of delaying sensory feedback using handwriting analyzers, oscilloscopes and master-slave robotic arms (Kalmus, Fry & Denes, 1960; Smith, McCrary & Smith, 1960; Kao, 1977; Sheridan & Ferrell, 1963; Ferrell, 1966). These studies tell us that a latency in visual feedback increases both the amount of time needed to complete a task and the number of errors made. Analysis of handwriting, for instance, shows that when participants wrote on a telescriber (a precursor to the facsimile used to transmit handwriting over a network), a delay between participants moving the pen and seeing the resulting marks caused letters not only to be formed more slowly, but also to vary more in shape (Kalmus, Fry & Denes, 1960; Smith, McCrary & Smith, 1960). We also learn that the relationship between the level of latency and the deterioration in performance appears to be linear: writing speed decreased proportionally with the amount of delay (Kalmus, Fry & Denes, 1960). This was also the case in a positioning task conducted on a master-slave operator (a system in which the movements of a robotic arm at one site correspond directly to actions initiated through an input device at another site). The time to position an object increased as a function of the amount of latency in the system, as did the number of positioning errors (Sheridan & Ferrell, 1963).

In a computer-based environment, Mackenzie & Ware (1993) conducted the first and arguably the most important quantitative research looking at the effects of visual latency, in a study where participants completed a Fitts' Law target acquisition task using a desktop display and mouse. Participants had to move the mouse from a starting point to a target with a latency from moving the mouse to seeing the cursor move on the screen of between 25 msec and 225 msec. The study produced two particularly interesting results. Firstly, there was a linear relationship between the amount of latency and the time participants took to move to the target. Secondly, the effects of delay depended upon task difficulty: the harder the task, the greater the detriment caused by latency. They incorporated their findings in a predictive model of performance, which factored the performance deterioration caused by an increase in latency into the Fitts' law model that described how movement times increased according to the Index of Difficulty (ID). At 0 latency, movement times (MT) are best predicted by model 1 (Fitts, 1964), where  $C_1$  and  $C_2$  are experimentally determined constants. However, when latency is introduced, it interacts with ID and the two have a multiplicative effect on movement time, as shown in model 2.

1.  $MT = C_1 + C_2 ID$
2.  $MT = C_1 + (C_2 + C_3 LATENCY) ID$

This study is perhaps the best known analysis of the effects of visual latency, and has frequently been cited as the defining model of how latency affects performance in distributed virtual environments. Other models of performance deterioration in the face of delayed visual feedback have been proposed by Ware & Balakrishnan (1994), who have extended Mackenzie & Ware's (1993) model to account for 3D reaching movements, and Day, Holt & Russell (1999), who have explained the effects of much longer delays (between 2 and 6 seconds) in a remote driving experiment in terms of working memory disruption.

As yet, there is no literature that offers a similar explanation of the effects of latency for the haptic domain. An early study by Ferrell (1966), showed that, as with visual feedback, delayed haptic feedback in a positioning task results in an increase in performance time and errors, but the intervals of delay (0.3, 1, 3 seconds) were not sufficiently small that we can learn about how the user might respond to subtle increases in latency. A study of latency in telesurgery (Ottensmeyer et al. 2000) demonstrated that surgeons are more sensitive to a latency in haptic feedback than visual when performing a laparoscopy task. However, the levels of latency examined (600 or 1200 msec) were again too large, and too few in number, to provide a model of performance.

Previous work by the current authors monitored the effects of delayed haptic feedback in a more systematic manner, using a Fitts' target acquisition task (Jay & Hubbold 2004), but failed to find an effect of latency on haptic performance. This occurred because the feedback in the study (a 'buzzing' over the target) was discrete, rather than continuous (like visual feedback), making it almost superfluous, as users could rely solely on the visual feedback. Whenever the visual and haptic feedback were out of synch, and hence conflicted, users ignored the haptic feedback, which occurred only briefly, choosing to rely instead on the continuous visual feedback. A second study used a tapping task where haptic feedback was crucial: participants relied on it to ensure they did not penetrate the surface of the targets (Jay & Hubbold 2005). However, latency still had far less effect in the haptic channel than in the visual. Whilst a delay of 69 msec disrupted visual feedback sufficiently to significantly increase movement times, a delay in haptic feedback did not affect performance until it reached 187 msec. To profile the effects of latency in the haptic channel, it is therefore important to ensure it is a) continuous, and b) essential to the task, factors that were both taken into account when designing the current experiment.

## 1.2 Delayed feedback in CVEs

There have been many studies observing visual and haptic collaboration in CVEs (e.g. Noma & Miyasato, 1997; Park & Kenyon, 1999; Basdogan et al. 2000; Sallnas et al. 2000; Hubbold 2002, Hikichi et al. 2002, Kim et al. 2004; Sallnas & Zhai, 2003, Gunn et al., 2005). However, none of these has systematically assessed in detail the effects of latency on task performance. The experiments that have addressed this issue are reported below. In most cases, they provide only qualitative results, or consider only a few levels of latency. In no cases are there quantitative measures stringent enough to provide a robust understanding of the effects of end-to-end latency in the visual or haptic channels.

A study of a virtual ball game in which players had to 'hit' the ball into the opposing player's goal resulted in some interesting observations of the effects of delayed visual feedback: the game was playable with a delay of 150 msec, but became progressively harder after this, becoming almost impossible after 500 msec (Vaghi, Greenhalgh & Benson, 1999). The players' dialogue revealed the discrepancies caused by the latency (the divergent views of the world led to arguments about the location of the ball, for

example), but the quantitative measure, the number of goals scored, was inadequate, as players scored so few even with the baseline level of zero delay.

A study by Allison et al. (2004) examining delayed haptic feedback showed that end-to-end latency caused significantly longer completion times and higher error rates in a co-operative positioning task. However, the study considered only 2 levels of additional latency (100 or 200 msec). Alhalabi et al. (2003) asked users to ‘shake hands’ in several predefined locations in a CVE, under varying levels of latency. In this case, the quantitative measure, completion time, was not affected until latency reached 400 msec, indicating that the task was not particularly sensitive to delay, and the application of the results was limited.

### 1.3 A collaborative target acquisition task with increasing latency


The current study was designed to systematically quantify how CVE users respond when data specifying their collaborator’s actions are delayed. The principal goals were to understand how performance deteriorated as a function of latency and to quantify users’ perceptions of task difficulty and disruption under increasing amounts of latency. The objective was to produce models of these data, providing a method of predetermining the effects of latency in CVEs, information that can be used for understanding the limitations of CVEs and managing the effects of latency.

An important aspect of the study was to gain a greater understanding of latency in the haptic channel, as very little is currently known about this. Virtual Interpersonal Touch is used infrequently in CVEs, but its effectiveness at conveying emotion means it is becoming increasingly popular, particularly in games (Bailenson et al., 2006; Bailenson & Yee, 2007).


The task was designed so that haptic feedback was essential. To understand the effects of latency on visual and haptic feedback individually, each modality must provide different, but equally important, types of information. To achieve this, the current study used a ‘collaborative target acquisition’ task. Users sat opposite each other in a very simple environment (see Figures 1 and 2). They had to reach forward (i.e., extend or move back a haptic device in order to achieve a change in the Z dimension) to touch the other person, and then move to a target together without losing contact. Their vision told them the location of their collaborator and the target in the XY plane. The environment had very poor visual depth cues: a wire box indicated the boundaries of the environment (see Figure 2), and perspective projection caused objects to become slightly bigger or smaller as they moved towards and away from the user. As such, users had to rely continuously on haptic feedback to know they were in contact with their collaborator. This set-up allowed us to differentiate between the aspects of performance that were affected by latency in the visual channel, and those that were affected by latency in the haptic channel. Visual latency was the primary influence on movement times and *aiming* errors (where the user ‘fell off’ his or her collaborator), as users relied on visual feedback to know the position of the objects in the XY plane. On the other hand, *penetration* and *separation* errors (where the user passed through or failed to remain touching the surface of his or her collaborator respectively), were likely to be due almost entirely to latency in the haptic channel, as users relied on haptic feedback to know they were in contact.

Performance data were recorded during both the *initiation period* (when participants first made contact) and the *movement period* (when participants were moving towards the target), allowing us to compare the effects of latency on two different types of movement. To monitor users’ conscious perceptions of the effects of latency, they were asked how

difficult they considered the task as a whole, and also to determine separately the level to which visual and haptic feedback felt disrupted.



*Figure 1: Moving collaboratively from the start circle in the center to the active target. Participants are facing each other in the virtual environment: the first screen shows the view of one participant; the second screen shows the view of the other. In each case the local user is represented by the small sphere and the remote user by the larger black circle. Alternative target positions are in each of the corners.*



*Figure 2: The equipment used for the experiment. Movements in the XY plane correspond directly to the left-right or up-down movements of the user's hand. Movements in the Z direction are achieved by moving the arm backwards and forwards.*

## 2. METHOD

### 2.1 Apparatus

The virtual environment was rendered using the Maverik++ SDK (Glencross et al. 2005) on a 2GHz PC with 512 MB RAM at one peer, and on a 3.2GHz PC with 1GB RAM at the other. Force feedback was provided by two FCS HapticMASTERS, which rendered geometry and forces at a rate of 2.5KHz. A HITACHI CP-X320 and a Sharp NoteVision 5 LCD projector displayed the environment on two 180cm x 250cm screens. The PCs and the haptic servers were set up on a dedicated LAN using a net-lynx mini-switch. Both PCs were running Linux and the netem kernel module (<http://linux-net.osdl.org/index.php/Netem>) was used to introduce the end-to-end latency. The mean frame rate of 72fps meant data was sent by one peer and received by the other on average every 14 msec.

### 2.2 Participants

Thirty participants (4 female and 26 male) between the ages of 20 and 30 took part in the study in pairs. Participants had not completed the task before and were naïve as to the aims of the experiment. This criterion was used for selecting participants as a primary aim of the study was to understand how people respond to latency when they are not expecting it. For the same reason, a confederate was not considered appropriate, and participants completed the experiment in pairs assigned at random. Each person was paid £10 for completing the experiment.

## 2.3 Virtual Environment

The very simple environment used for the task is shown in Figures 1 and 2. Using the HapticMASTER, participants each controlled an object. The object corresponding to the local end effector (the virtual representation of the user's movements) appeared as an orange sphere, which is the default representation of the end effector when using the HapticMASTER API. The remote end effector was represented by a larger black circle – the colour contrasted with the orange sphere, and the size made it difficult, though not impossible, to maintain contact. The participants were asked to assume they were facing each other in the virtual world. They were able to 'touch' each other by positioning the end effector over the black circle and moving forwards until they could feel its surface, which was rendered haptically to feel like a solid, slightly pliant circular plane 5cm in diameter. Collisions between the plane and the end effector were detected by the haptic server, and rendered using admittance control<sup>2</sup>. The discrepancy between the sizes of the local and remote object representations was necessary because the HapticMASTER, like most haptic devices, only supports point contact: if the remote object were actual size, it would be impossible for the users to remain touching. Pilot studies indicated that the 5cm diameter was appropriate, as it was possible for the users to maintain contact, but it was still difficult enough to require considerable concentration. Active collaboration between the two users was ensured by minimizing friction between the sphere and circle: both users had to participate equally to maintain contact, as they were not able to drag, or be dragged, by the other user. The lack of depth cues meant that users had to rely on the haptic feedback to know that they were touching their collaborator.

In an effort to limit external distractions, the environment was projected onto a screen in front of the participant, such that the visual rendering of the target and remote object was 8 times larger than the haptic rendering.

## 2.4 Network Topology

The experiment utilized a peer-to-peer design (see Figure 3), where the environment was simulated fully at each site, and locally initiated updates regarding the position and force on the end effector were applied immediately, at a rate of 2.5kHz in the haptic simulation, and every 14 msec (the frame rate of 7 msec + the mean projector refresh rate of 7 msec) on average in the visual simulation. These data were sampled every time around the graphical rendering loop and transmitted to the other environment, where the remote object was moved to the specified position and the force was applied on the end effector (if the objects were touching) as soon as an update was received. The remote object's position and force data in subsequent frames were rendered according to the last received update. This was crucial as in most cases updates were sent at a much slower rate than that at which either the graphics or particularly the haptics were rendered. If the update was applied only in the frame it arrived, there would be gaps in the rendering of the remote object, making the environment very unstable.

As there was no additional latency applied to local updates, any performance deterioration could be attributed entirely to the delay in updates from the other user. 10 levels of latency were assessed in the experiment: 0 msec; 25 msec; 50 msec; 75 msec; 100 msec; 150 msec; 200 msec; 250 msec; 300 msec; 400 msec. There was a focus on low levels of latency in particular, as these will persist in even the fastest wide area networks.

---

<sup>2</sup> A full description of the HapticMASTER's collision detection algorithm can be found at <http://www.fcs-cs.com/robotics/technology/#impvsadm>

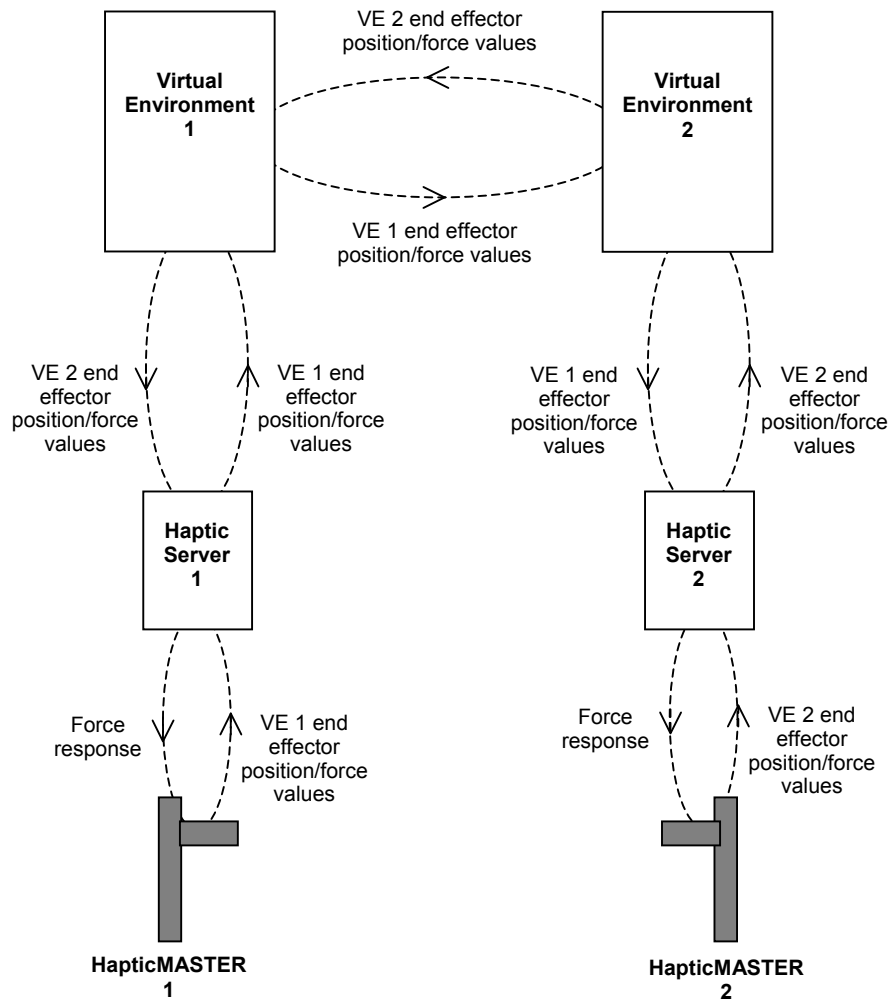


Figure 3: The network topology



## 2.5 Procedure

Each participant sat 3m from the projection screen used to display the environment with the HapticMASTER in front of him or her. Although both participants were in the same room (the projection screens were side by side), a barrier down the center of the room prevented them from seeing the other person or display. Participants were asked not speak to each other during the experiment.

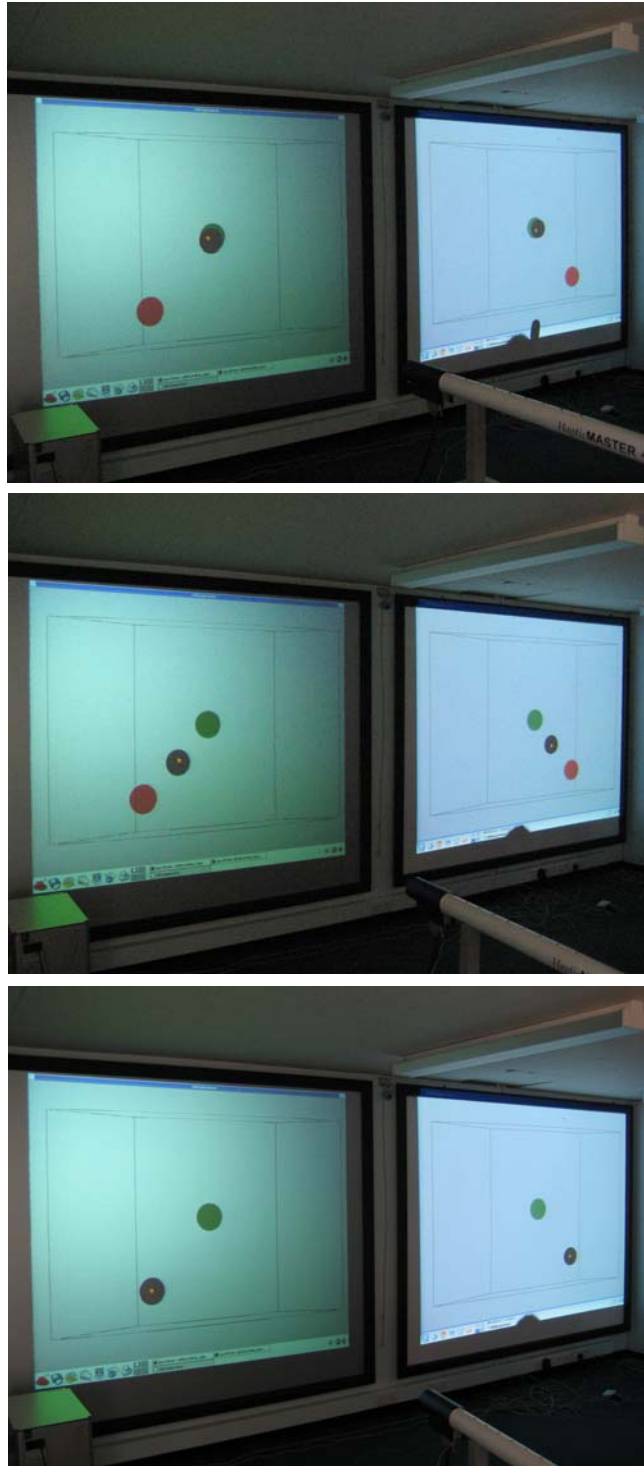
After they had read the information sheet, the participants were shown how to use the HapticMASTER and perform the task. The session began when participants positioned their end effectors in the green start circle in the center of the environment and initiated contact with each other. After two seconds, a red target circle appeared at the back of the environment in the corner 21cm from the start (in the haptic modality). The participants were asked to move towards it as quickly as possible, whilst ensuring that they maintained contact with each other. If they lost contact at any point, they were asked to re-establish it before continuing with the task. When they were over the target in the XY plane it disappeared, signifying the end of the trial. At this point, the collaborators broke contact and returned to the green start circle in their own time. They were then asked to repeat the process for the remaining eleven trials: initiating contact, waiting for the target to appear and then moving towards it together. The target remained the same distance from the start and the same size throughout the experiment, but the corner in which it appeared varied at random. Figure 4 shows two participants moving from the start to the target.

Participants completed one practice session of 12 trials with no added end-to-end latency, followed by 12 experimental sessions – one for each level of latency, ordered at random. To minimize practice effects, participants then performed the experimental sessions again in the reverse order. During the experiment, the lights were lowered to ensure the participants could see the display properly.

## 2.6 Performance measures

Movement times – the interval between the target appearing and the participants reaching it – were logged for every trial. The positions of the local and remote objects were stored in a buffer in every frame during both the initiation period (from the moment the participants established contact over the start until the target appeared two seconds later) and the movement period (the moment from the target appeared until the participants acquired it). The data were written to a file at the end of each trial, during the period when participants were returning to the start. The initiation and trial data were classed differently as users were required to make a different type of movement in each period, and there is a possibility that latency may affect each type of movement differently. In the initiation period, the users reached towards each other along the Z axis and stopped when contact occurred. In the movement period, the users had to maintain contact with each other in the Z direction, but in addition had to move in the XY plane towards the target.

Error rates were calculated by comparing the position data of the end effectors. To avoid an error, the local end effector had to be within 2.5cm of the remote end effector in the XY plane, and within 2.5mm of the remote end effector along the Z axis. At any point in time, the positions of the objects will be slightly different in each environment. All the data was thus recorded in both peers, and mean values of each measure were used for statistical analysis.



*Figure 4: Performing a trial. The barrier between the screens has been removed for clarity.*

Errors were classified in three ways, as shown in Figure 5. An *aiming* error occurred when participants failed to position their objects correctly in the XY plane, so the end effector was to the side of the remote object, rather than directly over it. As participants used vision to do this, a rise in aiming errors could reasonably be attributed to an increased disruption of visual feedback. If participants aimed correctly, they may still have made a *penetration* or *separation* error. A penetration error occurred when the participants moved towards each other too quickly, causing the remote object to be rendered around, or on the wrong side of the local end effector (the objects ‘passed through’ each other during the period between updates). To the participant, this gave the effect of penetrating the remote object. In a separation error, the participant failed to stay in touch with the surface of the remote object. As participants relied on touch to tell them if they have made contact, a rise in penetration and separation errors were attributed to an increased disruption of haptic feedback.

As latency may increase the time to reach the target, it would be inappropriate to measure the increase in errors per trial – the longer a participant spends in a particular condition, the greater the opportunity for making mistakes. Instead, the following measures were used for the error data: percentage of time spent in each type of error; mean duration of an error (the time between losing and regaining contact); frequency (number of errors per second); and magnitude (the distance by which the two objects became separated).

## 2.7 Questionnaire

After each session, participants were asked to rate on a scale of 1 to 10 how difficult they found the task, how disrupted they found the visual feedback and how disrupted they found the haptic feedback. They were asked to do this with regard to both *initiating* contact (at the start of a trial) and *maintaining* contact (when moving towards the target). The practice session, which did not contain any additional latency, acted as a baseline for their ratings, and thus had a score of 1. A score of 10 was appropriate if the user considered the task completely impossible.




Figure 5: The three categories of error. The remote user's current position is shown in black and the previous position is shown in grey. The arrows indicate the remote user's direction of movement.

## 2.8 Hypothesis

The experiment recorded movement times, error rates and users' perceptions of task difficulty and feedback disruption over 10 levels of latency. These measures allowed us to assess whether latency had different effects on visual and haptic feedback, whether its consequences varied according to the type of movement (initiating or maintaining contact), and whether it affected users' perceptions of latency in the same manner as their physical responses. Previous target acquisition studies have shown that delaying feedback from a user's own actions causes performance to deteriorate in a linear fashion: delaying feedback from another user's actions may well have a similar effect. The hypothesis was thus that all measures of performance, whether they applied to haptic or visual feedback, initiating or maintaining contact, actual task performance or users' questionnaire ratings, would increase in direct proportion to the amount of latency.

## 3. RESULTS

### 3.1 Qualitative Observations of User Performance

Latency had a negative effect on both visual and haptic feedback. The manner in which latency affected participants' movements is shown in Figures 6 and 7. Figure 6 shows a pair of participants moving in the XY plane to a target in the top right under each level of latency. Up to the 100 msec level, the participants had no trouble moving together in a straight line towards the target. After this point, the increase in the gap between updates made it more difficult to remain synchronized. It is interesting to observe, however, that latency does not appear to have a linear effect on aiming performance. Whilst movement at 400 msec of latency is clearly more disrupted than movement at 0 msec, performance is not obviously worse than at 150 msec of latency.

Figure 7 shows the same pair of participants moving in the Z direction. Latency had a much more serious effect on haptic feedback, and to show this in detail, the graphs show only the first 100 frames of the movement period. When there is no delay, participants are able to maintain close contact with each other, but by 50 msec, separation errors have become a problem. At 100 msec, penetration errors are also occurring regularly. Behaviour is far more disrupted by the latency in haptic feedback, and movements are far less predictable. It is possible to see the growing gap between updates and observe how users respond to it. However, it is difficult to tell through observation alone the extent to which different levels of latency impair haptic interaction.

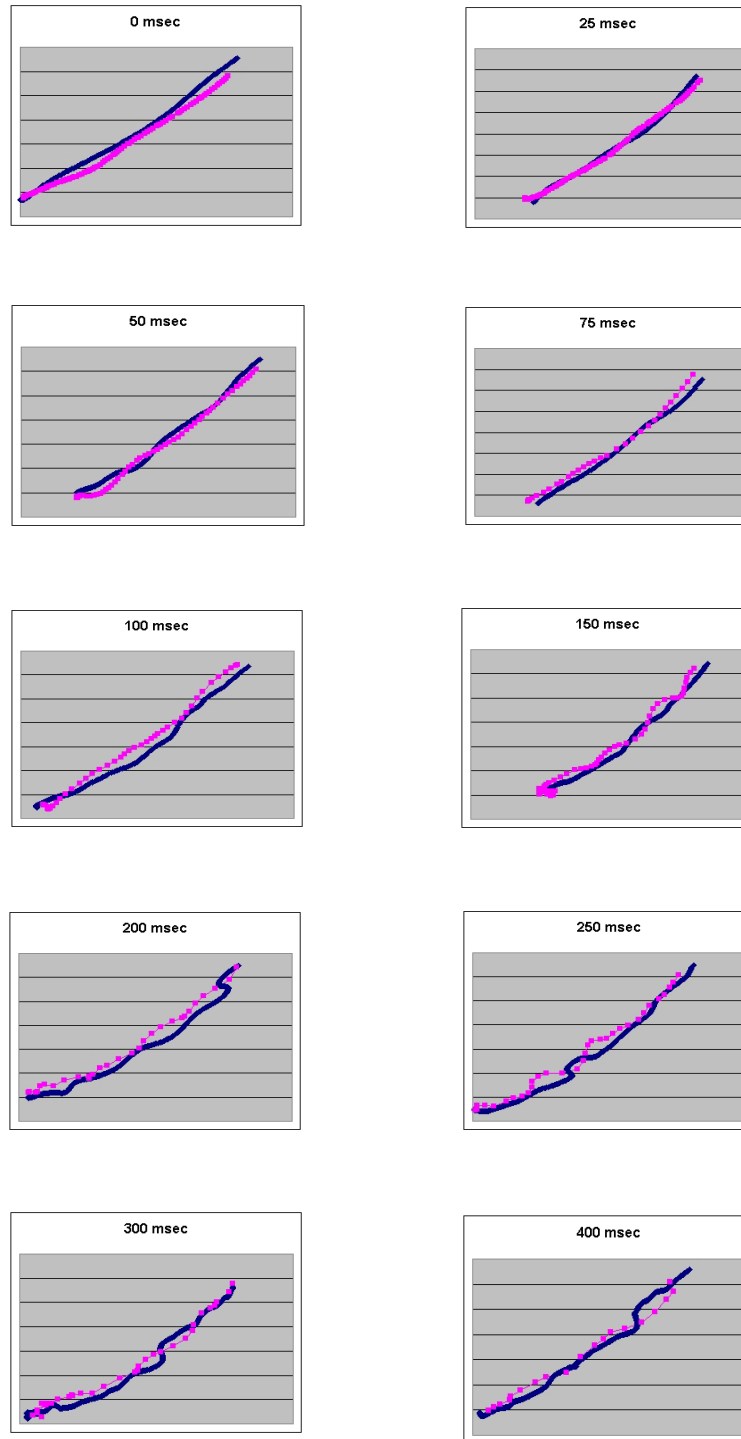
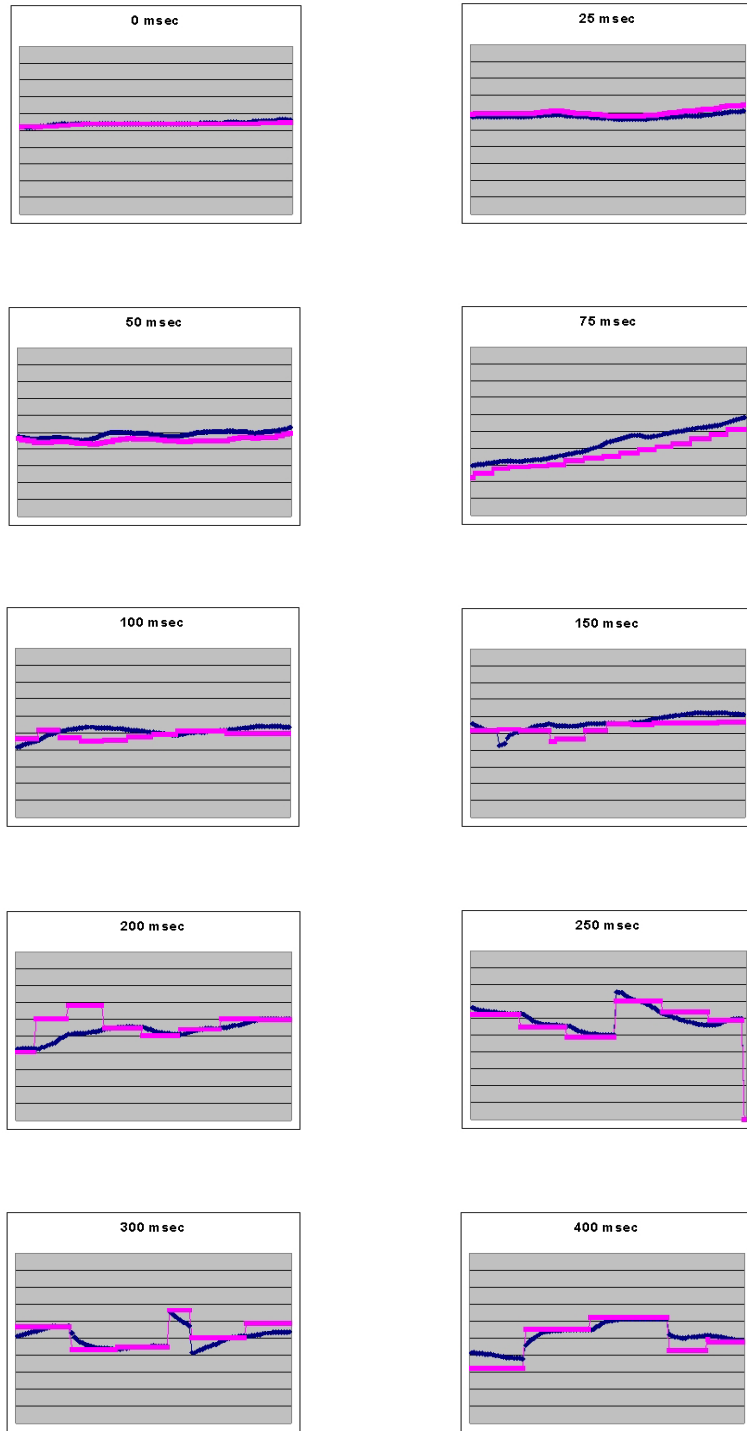


Figure 6: End effector movements in the XY plane from the start to a target in the top right becoming increasing unstable due to delayed visual feedback.



*Figure 7: End effector movements in the Z direction over the first 100 frames of a trial becoming increasingly unstable due to delayed haptic feedback*

### 3.2 Quantitative Measures of User Performance

The dependent variable, additional end-to-end latency, had ten levels ranging from 0 to 400 msec as described in Section 2.4, and applied to data travelling from both sites. If we term the standard latency between sites  $t$ , and the additional end-to-end latency  $x$ , site 1 received updates  $t + x$  msec after they had occurred at site 2, and site 2 received updates  $t + x$  msec after they had occurred at site 1. As participants completed the experiment in pairs, and each person's performance was dependent on his or her partner's, the pair was used as the sampling unit, and all the performance measures reported below are the mean of the values obtained by each member of the pair.

In every measure of performance, latency produced a detrimental effect – slowing movement down, increasing errors and causing higher ratings of difficulty and disruption. However, contrary to the hypothesis, the measures did not necessarily increase in direct proportion to latency, nor were they affected by latency in a uniform manner. While the increase in movement times and questionnaire ratings corresponded closely to the level of latency, error rates had a non-linear relationship with delay, which varied according to modality and movement type.

The results are divided into three sections, each considering a particular type of data in detail. The first section looks at movement times, the principle measure of users' responses to visual delay. The second section deals with the error data, considering the percentage, frequency, duration and magnitude of separation and penetration errors (which vary according to haptic delay) and aiming errors (which vary according to visual delay) during both the initiation and movement periods. The third section examines the questionnaire data, which indicate how delay affects users' perceptions of task difficulty and feedback disruption.

Graphs showing how each measure changes with latency are provided in each section, along with the results of ANOVAs used to determine the levels at which latency significantly altered performance. The data were also fitted with linear, quadratic and cubic regression models. The results of all the regression analysis can be found in Appendix A, but if a model provided a particularly good fit ( $R^2 > 0.9$ ), it is discussed alongside the rest of the results in this section.

The following subsections consider all of the quantitative measures, and contain numerous detailed observations, which have been included to provide a complete picture of the effects of latency in this situation. The concluding section provides a more concise summary of the data, and outlines the most important results and their implications.

#### 3.2.1 Movement Times

The mean target acquisition time under each condition is shown in Figure 8. A repeated measures ANOVA shows a significant main effect of latency on movement time ( $F_{9,126} = 48.752, p < 0.001$ ), revealed by pairwise comparisons to start at 50 msec.

Regression models<sup>3</sup> for the movement time data are shown in Table 1. Although simple linear regression explains 98% of the variance, a slightly better fit is provided by a quadratic model. However, the advantage is only small.

---

<sup>3</sup> For brevity we refer to Latency as LAG in the models.

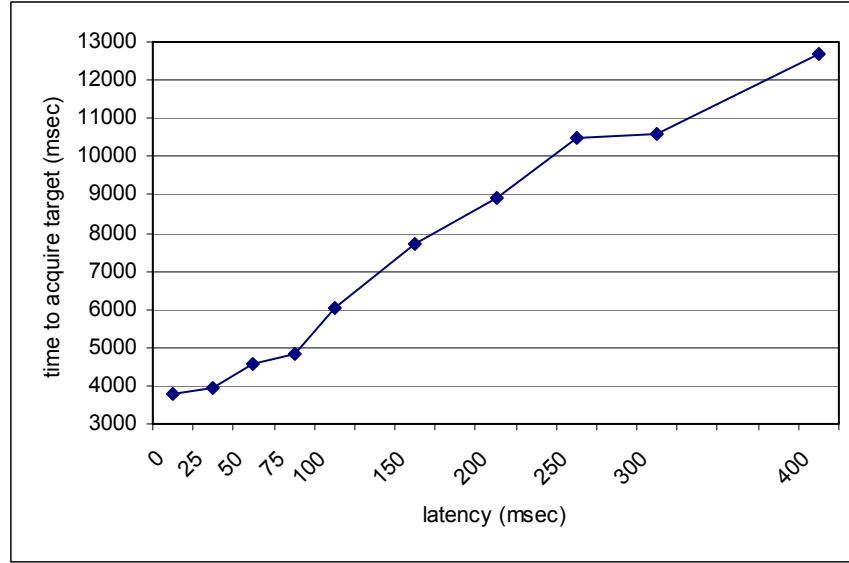


Figure 8: The mean time in msec to move from the start to the target.

Model	R <sup>2</sup>
MT = 3616.33 + 24.1657LAG	0.98
MT = 3224.72 + 31.7348LAG - 0.0199 LAG <sup>2</sup>	0.99

Table 1: Regression models for the movement time data (MT).

### 3.2.2 Errors

The positions of the local and remote objects in every frame were recorded for each trial from the moment that the participants first touched over the start, to the moment that they reached the target. If the two objects lost contact at any point during this period, it was regarded as an error. The first section below considers the errors that occurred in the initiation period – the two seconds between the participants first making contact and the target appearing. The following section then discusses the error data for the movement period – the time between the target appearing and the participants acquiring it. The same scales are used for the graphs in both sections, to illustrate the differences between the two periods.

#### *Initiation Period*

The first thing to note is that virtually no aiming errors are made during the initiation period, regardless of the level of latency. This is unsurprising, as there is very little movement in the XY plane during this period – participants are asked simply to make contact and remain over the start circle. Aiming errors are therefore not considered further in this section.

Figure 9 shows the percentage of errors in the initiation period as a function of latency. There is a slight rise in the percentage of penetration errors ( $F_{9,126} = 3.80$ ,  $p < 0.001$ ); however, pairwise comparisons show this is not significant until 200 msec. By contrast, the percentage of separation errors rises steeply from 25 msec, peaks at 100 msec and



then slowly tapers off, until it is significantly smaller again at 400 msec ( $F_{9,126} = 7.25$ ,  $p < 0.001$ ). Separation errors are far more common than penetration errors, and exert the greatest influence on the overall error pattern. A cubic regression model provides the best fit for the overall percentage error and the penetration percentage error (see Table 2).

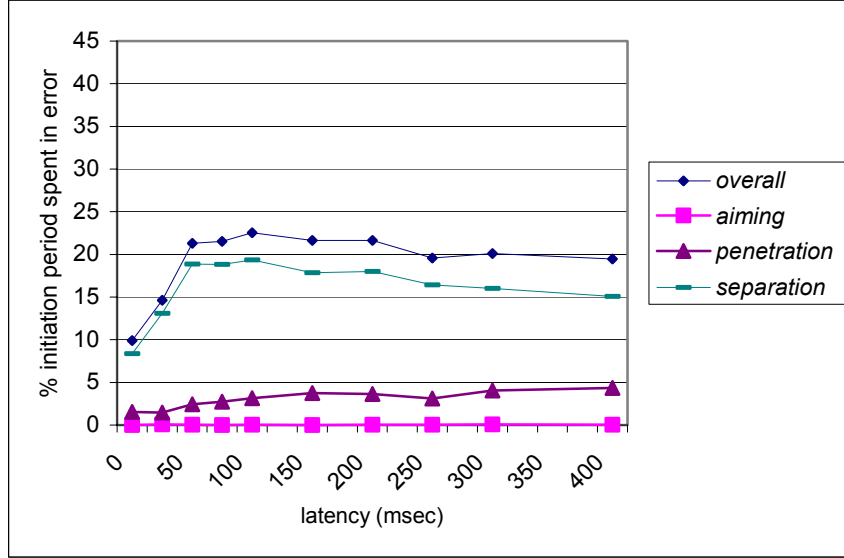


Figure 9: Percentage of initiation period spent in error.

Model	R <sup>2</sup>
O PER = 10.8412 + 0.2056LAG - 0.001LAG <sup>2</sup> + 0.0000014LAG <sup>3</sup>	0.90
P PER = 1.2591 + 0.0279LAG - 0.0001LAG <sup>2</sup> + 0.00000015LAG <sup>3</sup>	0.92

Table 2: Regression results for the percentage of time spent in error during the initiation period. O\_PER = overall; P\_PER = penetration.

The pattern of error frequency is also different for separation and penetration errors (see Figure 10). Whilst separation errors peak at 75 msec and then start to fall ( $F_{9,126} = 5.20$ ,  $p < 0.001$ ), penetration errors rise at a slower rate, but continue to increase in frequency as latency increases ( $F_{9,126} = 11.48$ ,  $p < 0.001$ ). Table 3 shows the regression results for the overall frequency of errors and the frequency of penetration errors during the initiation period. In both cases, cubic models offer a good fit, but these models fail to capture the peak at 75 msec, and subsequent drop at 100 msec, indicating that they are not a completely reliable explanation for the data.

There is no main effect of latency on the duration of errors during the initiation period (see Figure 11), although there is a main effect of type: penetration errors have a mean duration of 57 msec; and separation errors are on average more than twice as long at 128 msec ( $F_{2,28} = 7.92$ ,  $p < 0.001$ ). The magnitude of errors (see Figure 12) is not affected by latency either, but again, penetration and separation errors are significantly greater than aiming errors ( $F_{2,28} = 6.59$ ,  $p < 0.005$ ).




Figure 10: Mean error frequency/second during the initiation period

Model	R <sup>2</sup>
O_FREQ = 1.0289 + 0.0157LAG - 0.00007LAG <sup>2</sup> + 0.000000083LAG <sup>3</sup>	0.93
P_FREQ = 0.2513 - 0.0038LAG - 0.000006LAG <sup>2</sup>	0.93
P_FREQ = 0.1990 + 0.0059LAG - 0.00002LAG <sup>2</sup> + 0.000000023LAG <sup>3</sup>	0.96

Table 3: Regression results for error frequency/sec during the initiation period. O\_FREQ = overall, P\_FREQ = penetration.

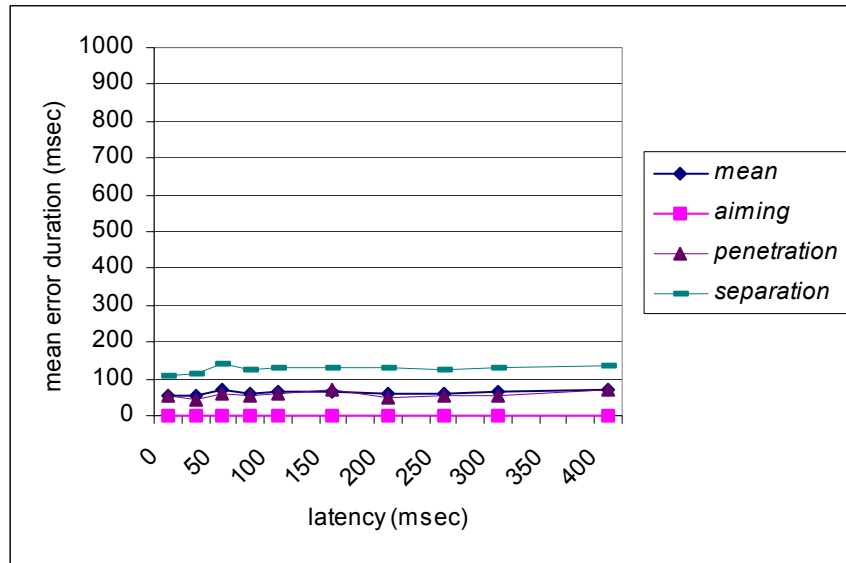


Figure 11: Mean error duration during the initiation period.




Figure 12: Mean error magnitude during the initiation period

#### *Movement Period*

Figure 13 displays the mean percentage of time spent in each type of error during the movement period. There is a main effect of latency on the overall percentage error ( $F_{9,126} = 14.864$ ,  $p < 0.001$ ), which starts at 25 msec. However, there is also a main effect of type ( $F_{2,28} = 11.532$ ,  $p < 0.005$ ), and an interaction between latency and type ( $F_{18,252} = 2.872$ ,  $p < 0.001$ ), indicating that, again, the effects of latency are not uniform for each type of error.

Pairwise comparisons show that the significant increase in the error percentage occurs from 25 msec only for separation errors ( $F_{9,126} = 7.979$ ,  $p < 0.001$ ). The percentage does not increase significantly until 100 msec for aiming errors ( $F_{9,126} = 4.768$ ,  $p < 0.001$ ) and penetration errors ( $F_{9,126} = 3.373$ ,  $p < 0.001$ ). Regression does not adequately explain the effects of latency on any individual error type, but it does provide a strong cubic model fitting the overall percentage of error during the trial period (see Table 4).

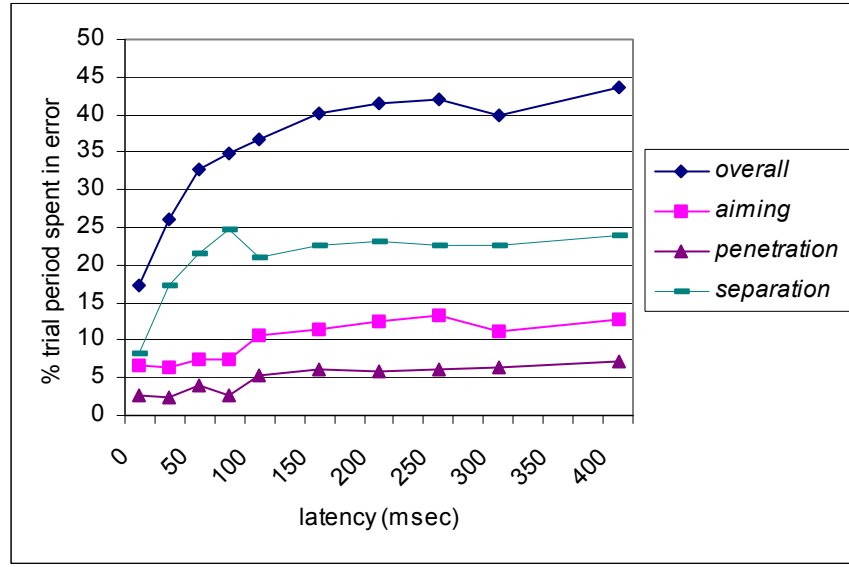


Figure 13: Percentage of movement period spent in error

Model	R <sup>2</sup>
O_PER = 19.0438 + 0.2854LAG - 0.0012LAG <sup>2</sup> + 0.0000016 <sup>3</sup>	0.99

Table 4: Regression results for the percentage of time spent in error during the movement period. O\_PER = overall.

Figure 14 shows the error frequency results from the movement period. If errors are considered as a whole, there is a main effect of latency on error frequency, starting at 25 msec ( $F_{9,126} = 7.503$ ,  $p < 0.001$ ). However, at 100 msec, in contrast to the percentage of errors, the number of errors per second starts to decline. At 400 msec of additional latency the frequency is not significantly different to that at 0 msec. As before, there is a main effect of type ( $F_{2,28} = 12.924$ ,  $p < 0.001$ ), and an interaction between latency and type ( $F_{18,252} = 5.646$ ,  $p < 0.001$ ).

Separation errors are again the most common, and thus have the greatest influence on the overall pattern. The main effect of latency starts at 25 msec, peaks at 75 msec and then starts to fall again, until at 400 msec the frequency of separation errors is no greater than it is at 0 msec ( $F_{9,126} = 6.837$ ,  $p < 0.001$ ). Penetration errors follow a similar pattern ( $F_{9,126} = 4.645$ ,  $p < 0.05$ ), although the slope is shallower and the main effect is not significant until 50 msec. Aiming errors, on the other hand, increase in frequency steadily as latency increases, leveling off at 150 msec ( $F_{9,126} = 2.301$ ,  $p < 0.05$ ). This relatively shallow slope is the only one to be adequately explained by a regression model (see Table 5).

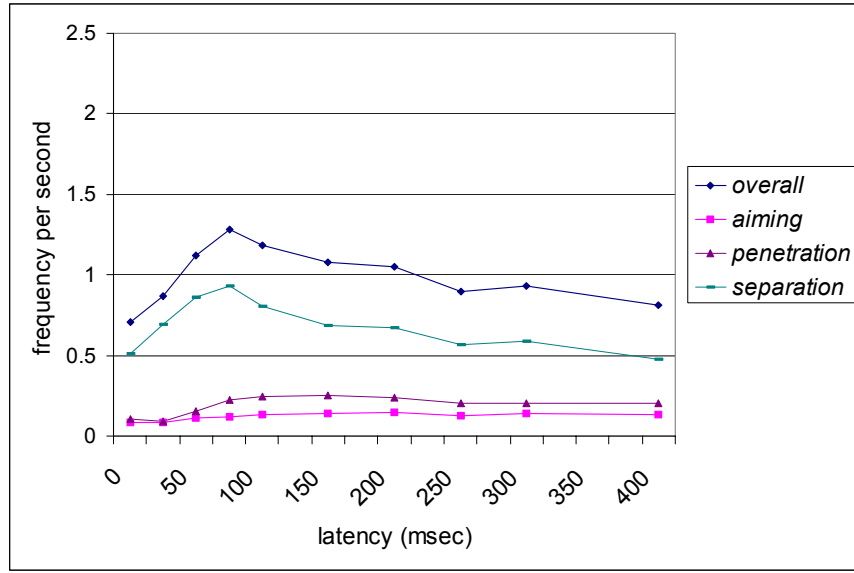


Figure 14: Mean error frequency/second during the movement period

Model	R <sup>2</sup>
A FREQ = 0.0899 + 0.0004LAG - 0.0000008LAG <sup>2</sup>	0.90
A FREQ = 0.0827 + 0.0007LAG - 0.000003LAG <sup>2</sup> + 0.0000000032LAG <sup>3</sup>	0.96

Table 5: Regression results for error frequency/sec during the movement period.  
A\_FREQ = aiming.

Figure 15 shows the mean duration of movement period errors as a function of delay. In contrast to the initiation period, there is a main effect of latency on duration ( $F_{9,126} = 14.665$ ,  $p < 0.001$ ) as well as a main effect of type ( $F_{2,28} = 41.53$ ,  $p < 0.001$ ). Pairwise comparisons show that separation errors start to increase significantly at 25 msec ( $F_{9,126} = 17.641$ ,  $p < 0.001$ ). Despite the fact that they are by far the most common, aiming errors do not increase significantly in duration until 150 msec ( $F_{9,126} = 3.743$ ,  $p < 0.001$ ), like penetration errors, which are the least common ( $F_{9,126} = 14.626$ ,  $p < 0.001$ ). Table 6 displays the results of the error duration regression. Separation error duration is explained by a linear model, whilst for the overall mean and penetration error duration, the quadratic and cubic models have a slight advantage. However, the difference between these and the linear models is only small.

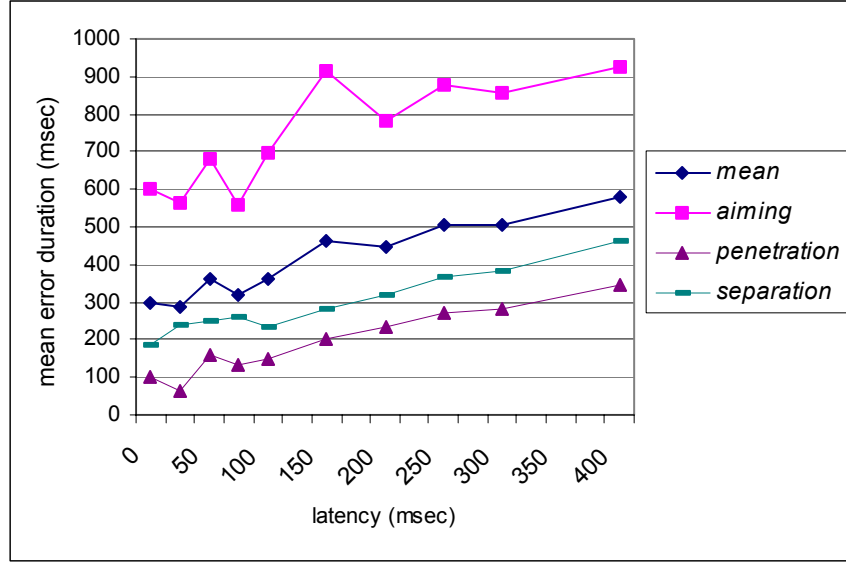


Figure 15: Mean error duration during the movement period

Model	R <sup>2</sup>
M DUR = 303.459 + 0.7558LAG	0.94
M DUR = 283.843 + 1.135LAG - 0.001LAG <sup>2</sup>	0.96
P DUR = 96.2488 + 0.6457LAG	0.94
P DUR = 89.6966 + 0.7723LAG - 0.0003LAG <sup>2</sup>	0.94
P DUR = 92.4978 + 0.6616LAG - 0.0004LAG <sup>2</sup> - 0.000001LAG <sup>3</sup>	0.95
S DUR = 198.993 + 0.6949LAG	0.96

Table 6: Regression results for error duration during the movement period.  $M\_DUR$  = mean;  $P\_DUR$  = penetration;  $S\_DUR$  = separation.

Figure 16 shows the mean magnitude of each type of error. There is a significant main effect of latency ( $F_{9,126} = 13.367$ ,  $p < 0.001$ ) starting at 50 msec, a significant main effect of type ( $F_{2,28} = 19.233$ ,  $p < 0.001$ ) and an interaction between latency and type ( $F_{18,252} = 3.973$ ,  $p < 0.001$ ). Although penetration errors are affected by latency from 50 msec ( $F_{9,126} = 12.649$ ,  $p < 0.001$ ) and separation errors increase significantly when latency is only 25 msec ( $F_{9,126} = 17.09$ ,  $p < 0.001$ ), delay does not have a significant effect on the magnitude of aiming errors. Table 7 shows regression models for the overall mean, penetration and separation error magnitude. A linear model predicts how separation error magnitude rises with latency. A quadratic model provides the best fit for the overall mean error magnitude, whilst a cubic model provides the best fit for penetration error data, but only by a small margin.




Figure 16: Mean error magnitudes during the movement period

Model	R <sup>2</sup>
M MAG = 0.2974 + 0.0019LAG - 0.000002LAG <sup>2</sup>	0.93
P MAG = 0.171 + 0.0035LAG - 0.000004LAG <sup>2</sup>	0.92
P MAG = 0.134 + 0.0049LAG - 0.000001LAG <sup>2</sup> - 0.000000016LAG <sup>3</sup>	0.93
S MAG = 0.1179 + 0.001LAG	0.99

Table 7: Regression results for error magnitude during the movement period. *M\_MAG* = mean; *P\_MAG* = penetration; *S\_MAG* = separation.

### 3.2.3 Questionnaire data

Figures 17 and 18 show the participants' mean questionnaire ratings. Again, the value entered into the analysis was the mean of the pair's ratings. There is a main effect of latency, starting at 50 msec ( $F_{9,126} = 2.23$ ,  $p < 0.001$ ). However, there is no effect of question type: ratings for difficulty and the disruption of visual and haptic feedback in both the initiation and trial phases all rise at the same rate as latency increases.

Regression models for the questionnaire data are shown in Table 8. Cubic or quadratic models provide the best fit, but in most cases the advantage they have over a linear model is very small.




Figure 17: Mean difficulty and disruption ratings for initiating contact

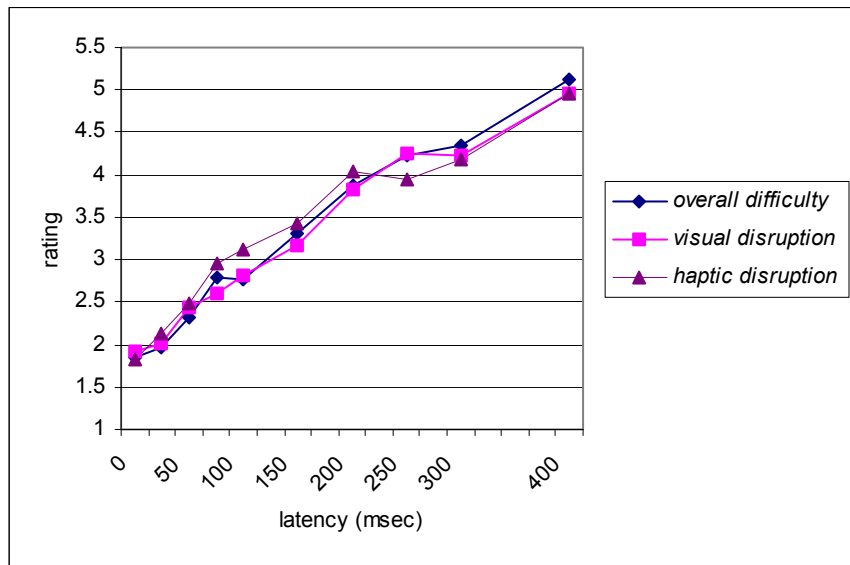


Figure 18: Mean difficulty and disruption ratings for maintaining contact.



Model	R <sup>2</sup>
ID = 1.8575 + 0.0081LAG	.98
ID = 0.7046 + 0.0111LAG - 0.000008LAG <sup>2</sup>	.99
IV = 1.9238 + 0.0073LAG	.97
IV = 1.7889 + 0.0099LAG - 0.000007LAG <sup>2</sup>	.98
IH = 1.9779 + 0.0074LAG	.96
IH = 1.7522 + 0.0118LAG - 0.00001LAG <sup>2</sup>	.99
MD = 1.9598 + 0.0084LAG	.98
MD = 1.7930 + 0.0116LAG - 0.000008LAG <sup>2</sup>	.99
MV = 2.0006 + 0.0079LAG	.98
MV = 1.8459 + 0.0109LAG - 0.000008LAG <sup>2</sup>	.99
MH = 2.1692 + 0.0073LAG	.94
MH = 1.9469 + 0.0116LAG - 0.00001LAG <sup>2</sup>	.97
MH = 1.7633 + 0.0189LAG - 0.00006LAG <sup>2</sup> + 0.000000081LAG <sup>3</sup>	.99

Table 8: Regression results for difficulty and disruption ratings. *I* = initiating contact; *M* = maintaining contact; *D* = overall difficulty; *V* = visual disruption; *H* = haptic disruption.

### 3.3 Summary

The data showed that the relationship between latency and the visual and haptic performance measures was complex and multifaceted. Whilst movement times and questionnaire ratings increased from 50 msec of latency, and in direct proportion to delay, error rates increased from 25 msec of latency, tended to be non-linear, and varied according to both modality and the type of movement that the user performed. For example, aiming errors did not occur during the initiation period, when users simply remained stationary over the start, but did occur during the movement period, and increased significantly with latency. In contrast, haptic errors occurred during both the types of movement, and were far more common than visual errors. The main effects of type showed that some errors occurred more frequently than others. The percentage of separation errors was by far the highest, indicating that in this task maintaining contact with the surface of an object in the Z direction is a particular problem. There were also interactions between the different error types, indicating that latency did not affect each aspect of performance in exactly the same way.

Despite these variations in the precise effects of latency, it is possible to see some broad trends in the data. The key results can be summarized as follows:

- Questionnaire ratings rose in direct proportion to latency from 50 msec, indicating that users accurately perceived the increase in latency.
- Movement times also increased from 50 msec, closely fitting a linear model.
- In contrast, the error percentage rose steeply from 25 msec, but stagnated after 100 msec.

So, the question is, why did movement times and ratings of feedback disruption and task difficulty increase with latency as hypothesized, but error rates follow a non-linear pattern? The next section attempts to answer this question, proposing the Impact-Perceive-Adapt model of performance as an explanation for the pattern of user behaviour observed here.

#### 4. THE 'IMPACT-PERCEIVE-ADAPT' MODEL OF PERFORMANCE

Whilst movement times and questionnaire ratings increase in direct proportion to latency, error rates do not: the ability to touch and maintain contact deteriorates rapidly as latency rises to start with, but declines far more slowly with latencies over 100 msec.

The relatively static error rates that eventually occur can be attributed to a speed/accuracy trade-off, first documented in target acquisition tasks by Fitts (1954). As the Index of Difficulty of the task increases, people slow their movements down to avoid an increase in errors. Here we manipulate latency, rather than the width and distance of the target, but the effect is the same: as latency increases, people slow their movements to avoid increasing the error rate.

This phenomenon provides a good explanation for the stagnation of error rates, but it does not explain why they rise so rapidly to start with. To understand this, it is pertinent to consider how people perceive latency, and the effect it has on their performance. If we consider the questionnaire ratings, we can see that users do not rate the task as more difficult, or feedback as more disrupted, until latency reaches 50 msec. If people do not perceive any difficulties caused by latency, they are presumably unaware that it is affecting their performance, and do not change their behaviour.

Users become aware of latency at 50 msec, and move more slowly, although they do not slow down enough to halt the rise in errors at this point. Figure 8 shows that the increase in movement times has a relatively shallow gradient for levels of latency below 100 msec. Observation of the error percentage (Figures 8 and 12) shows that it is also below 100 msec that the increase in error rates is at its greatest.

So, although users are aware of latency affecting their performance, when it is less than 100 msec they fail to slow their movements down enough to stop an increase in errors. Why is this? A possible explanation is that although people are aware of the latency, it remains relatively undistruptive, as the gap between updates is not sufficient to cause the breakdown of the *perception of immediate causality* (Card, 1983). At low levels of latency (up to 100 msec), when their collaborator's movements still appear smooth, users alter their behaviour little, and haptic error rates rise rapidly. As soon as feedback becomes obviously disrupted, however, they start to use error-limiting strategies, such as slowing their movement down, in order to reduce errors. Users start to increase their ratings of task difficulty and feedback disruption from 50 msec, and as such this level must be regarded as the point at which latency-induced artifacts became salient. However, as error rates continue to rise after this point, it may actually be 100 msec, regarded as the upper threshold for the breakdown of the perception of immediate causality (Card, 1983), which is really significant when it comes to altering users' strategies and behaviour.

This observation suggests three latency thresholds at which user performance changes:

- At the *impact threshold (25 msec)*, errors increase significantly, but users are yet to become aware of the latency.
- At the *perception threshold (50 msec)*, users perceive latency-induced artifacts and start to change their behaviour, but as the movements of the remote collaborator still appear smooth, the delay is not disruptive enough to slow participants' movements sufficiently to halt the rise in error rates.
- At the *adaptation threshold (100 msec)*, the time between updates is sufficient to cause a breakdown in the perception of immediate causality, so the remote

collaborator's movements appear jerky and disjointed. From this point onwards, people slow their movements in direct proportion to latency to stop the error rate from rising any further.

Although this model can be applied to both types of movement investigated here, error rates are slightly different in each case: during the initiation period, the error percentage rises rapidly to 50 msec, peaks at 100 msec, and then remains static; during the movement period, the rate continues to rise after 100 msec, but much more slowly than it did initially. Error duration and magnitude remain static in the initiation period. In the movement period, magnitude and duration increase in approximately direct proportion to latency: the added challenge of moving in another dimension makes it increasingly difficult to regain contact as latency grows. It therefore appears that in both cases, people adapt by minimizing error frequency – the number of mistakes they make. The different percentage error rates for each type of interaction occurs to a large extent because magnitude and duration, over which users may have less control, do not increase with latency when people are initiating contact, but significantly increase with latency when they are moving towards the target.

## 5. DISCUSSION

Collaborative Virtual Environments have the potential to greatly enhance long-distance communication by enabling people to interact visually and haptically in a shared virtual space. However, if we are ever to design CVEs that are truly effective, we must fully understand the issues caused by end-to-end network latency. In the work described here, we advance this knowledge with regard to the delayed sensory information travelling between users.

The aim of the experiment was to quantify performance deterioration and the user's perception of task difficulty in the face of increasingly delayed visual and haptic feedback from the remote user. The hypothesis predicted that all the experimental measures, whether they applied to haptic or visual feedback, initiating or maintaining contact, actual task performance or users' questionnaire ratings, would increase in direct proportion to the amount of latency.

The results showed that the consequences of latency were far more complicated than this, varying according to both modality and movement type. Participants were clearly able to perceive the effects of delay, as they rated task difficulty and feedback disruption consistently higher with every increment in the level of latency above 50 msec. However, only in the case of movement times did the questionnaire results bear any direct relation to performance deterioration. The rise in error rates started at 25 msec (before users perceived the increase in latency), and slowed considerably after 100 msec, despite the fact that participants continued to experience a rise in task difficulty and feedback disruption.

The Impact-Perceive-Adapt model of performance provides an explanation for this pattern of behaviour. It identifies three key thresholds at which user performance changes in response to latency. At the *impact threshold*, errors have started to increase, but users are unaware of this fact, so do not take any action to stop it. At the *perception threshold*, people can perceive latency-induced artifacts and start to slow down a little (or use another, as yet unknown error reducing strategy), but the effects of delay are not sufficient to induce them to slow down enough to halt the rise in errors. At the *adaptation threshold*, the perception of immediate causality breaks down, making the remote collaborator's movements appear discontinuous, and from this point onwards, people

slow their movements in direct proportion to latency to stop the error rate from rising any further. In the task reported here, the thresholds are at 25 msec, 50 msec and 100 msec respectively. All three may vary according to the type of interaction users are engaged in, although if the breakdown of the perception of immediate causality underlies the adaptation threshold, it should occur between 50 and 100 msec (Card, 1983).

For both types of movement assessed here, this model corresponds broadly to the frequency with which people make errors. However, it is important to note that the precise manner in which latency impacts on performance is very complex. For example, there is an interaction between the types of error: as separation errors drop, penetration errors start to rise. The rise in error magnitude and duration during the movement period means that the overall error rate takes a different form to the one in the initiation period, where these measures remained static, and this in turn may affect error frequency.

To understand the intricacies of how latency affects user performance in CVEs requires further investigation. In the current task, low levels of latency were particularly detrimental to haptic performance, causing a large number of separation errors. However, this may not hold for all types of task or interaction. Further work examining the effects of the cross modal interaction would be interesting, as would investigating the relationship between latency and task difficulty or type. How would users respond if task difficulty were varied, for example, or if users were to manipulate an object, rather than touching each other directly? In particular, it would be useful to investigate whether the Impact-Perceive-Adapt explanation of user behaviour has a broader application, beyond describing the kind of interaction documented here.

Alhalabi, M., Horiguchi, S. & Kunifuji, S. (2003). An experimental study on the effects of network delay in cooperative shared haptic environment. *Computers & Graphics* 27, 205-213.

Allison, R. S., Zacher, J. E., Wang, D. & Shu, J. (2004). Effects of delay on a collaborative motor task with telehaptic and televisual feedback. In *Proceedings of the ACM SIGGRAPH International Conference on VRCAI*, 375-381.

Bailenson, J.N., Yee, N., Brave, S., Merget, D., & Koslow, D. (2006, in press). Virtual Interpersonal Touch: Expressing and Recognizing Emotions through Haptic Devices. *Human-Computer Interaction*.

Bailenson, J. & Yee, N. (2007 in press). Virtual interpersonal touch: Haptic interaction and copresence in collaborative virtual environments. *International Journal of Multimedia Tools and Applications*

Basdogan, C., Ho, C., Srinivasan, M. & Slater, M. (2000). An experimental study on the role of touch in shared virtual environments. *ACM Transactions on Computer-Human Interaction*, 7, 4, 443-460.

Card, S., Moran, T. and Newell, A. *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum Associates, 1983.

Day, P. N., Holt, P. O. & Russell G. T. (1999). Modelling the effects of delayed visual feedback in real-time operator control loops: A cognitive perspective. In *Proceedings of the XVIII European Annual Conference on Human Decision Making and Manual Control*, 70-79.

Ferrell, W. R. (1966). Delayed force feedback. *Human Factors*, 8, 449-455.

Fitts, P. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47, 6, 381-391.

- Glencross, M., Hubbard, H. & Lyons, B. (2005). Dynamic primitive caching for haptic rendering of large-scale models. In *Proceedings of the WorldHAPTICS First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 517-518.
- Gunn, C., Hutchins, M. Stevenson, D. Adcock, M. & Youngblood, P. (2005). Using collaborative haptics in remote surgical training. In *Proceedings of the WorldHAPTICS First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 481-482.
- Hikichi, K., Morino, H., Sezali, K. & Yasuda, Y. (2002). The evaluation of delay jitter for haptics collaboration over the internet. In *Global Telecommunications Conference (GLOBECOM)*, 1492-1496
- Hubbold (2002). Collaborative stretcher carrying: a case study. In *Proceedings of the 8<sup>th</sup> Eurographics Workshop on Virtual Environments*, 7-12.
- Jay, C. & Hubbard, R. (2004). Does performance in the face of delayed sensory feedback change with practice? In *Proceedings of Eurohaptics Conference*, 530-533.
- Jay, C. & Hubbard, R. (2005). Delayed visual and haptic feedback in a reciprocal tapping task. In *Proceedings of the WorldHAPTICS First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 655-656.
- Kalmus, H., Fry, D. & Denes, P. (1960). Effects of delayed visual control on writing, drawing and tracing. *Language & Speech*, 3, 96-108.
- Kao, H. (1977). Effects of feedback and intermittency of feedback in compensatory tracking. *Perceptual and Motor Skills*, 44, pp. 1079-1085.
- Kim, J., Kim, H., Tay, B., Muniyandi, M., Jordan, J., Mortensen, J., Oliveria, M. & Slater, M (2004). Transatlantic touch: a study of haptic collaboration over long distance. *Presence*, 13, 3, 328-337.
- MacKenzie, S. & Ware, C. (1993). Latency as a Determinant of Human Performance in Interactive Systems. In *Proceedings of the ACM Conference on Human Factors in Computing Systems - INTERCHI*, 488-493.
- Noma, H. & Miyasato, T. (1997). Cooperative object manipulation in virtual space using virtual physics. *ASME-DSC-Vol.61*, pp.101-106.
- Ottensmeyer, M. P., Hu, J., Thompson, J. M., Ren, J., & Sheridan, T. B. (2000) Investigations into performance of minimally invasive telesurgery with feedback time delays. *Presence* 9, 4, 369-382.
- Park, K. & Kenyon, R. (1999). Effects of Network Characteristics on Human Performance in a Collaborative Virtual Environment, In *Proceedings of IEEE VR*, 104-111.
- Sallnas, E., Rassmuss-Grohn, K. & Sjostrom, C. (2000). Supporting presence in collaborative environments by haptic force feedback. *ACM Transactions on Computer-Human Interaction*, 7, 4, 461-476.
- Sallnas, E. & Zhai, S. (2003). Collaboration meets Fitts' Law: passing virtual objects with and without haptic force feedback. In *Proceedings of INTERACT, IFIP Conference on Human-Computer Interaction*, 97-104.
- Sheridan, T. B. & Ferrell, W. R. (1963). Remote manipulative control with transmission delay. *IEEE Transactions on Human Factors in Electronics*, 4, 1, 25-29.
- Smith, W., McCrary, J. & Smith, K. (1960). Delayed visual feedback and behaviour. *Science*, 132, 1013-1014.

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

Ware, C. & Balakrishnan, R. (1994). Reaching for objects in VR displays: lag and frame rate. In *ACM Transactions on Computer-Human Interaction*, 1, 4, 331-356.

## Appendix A: Results of the Regression Analysis

This appendix contains all the significant regression results obtained with data from the collaborative target acquisition experiment. The best fitting model for each measure is displayed in bold.

Model	R <sup>2</sup>
MT = 3616.33 + 24.1657LAG	.98
<b>MT = 3224.72 + 31.7348LAG - 0.0199 LAG<sup>2</sup></b>	<b>.99</b>
MT = 3505.3 + 20.6393LAG + 0.0541LAG <sup>2</sup> - 0.0001 LAG <sup>3</sup>	.99

Table 9: Regression models for the movement time data (MT).

Model	R <sup>2</sup>
<b>O_PER = 11.3860 + 0.1875LAG - 0.0010LAG<sup>2</sup> + 0.0000014LAG<sup>3</sup></b>	<b>.83</b>
P_PER = 1.9966 + 0.0064LAG	.73
P_PER = 1.6432 + 0.0132LAG - 0.00002LAG <sup>2</sup>	.81
<b>P_PER = 1.3148 + 0.0262LAG - 0.0001LAG<sup>2</sup> + 0.00000014LAG<sup>3</sup></b>	<b>.87</b>
S_PER = 10.0416 + 0.1615LAG - 0.0009LAG <sup>2</sup> + 0.0000012 <sup>3</sup>	.79

Table 10: Regression results for the percentage of time spent in error during the initiation period. O\_PER = overall; P\_PER = penetration; S\_PER = separation

Model	R <sup>2</sup>
O_FREQ = 1.5359 + 0.0021LAG	.46
O_FREQ = 1.2178 + 0.0082LAG - 0.00002LAG <sup>2</sup>	.82
<b>O_FREQ = 1.0289 + 0.0157LAG - 0.00007LAG<sup>2</sup> + 0.000000083LAG<sup>3</sup></b>	<b>.93</b>
P_FREQ = 0.3733 - 0.0015LAG	.75
P_FREQ = 0.2513 - 0.0038LAG - 0.000006LAG <sup>2</sup>	.93
<b>P_FREQ = 0.1990 + 0.0059LAG - 0.00002LAG<sup>2</sup> + 0.000000023LAG<sup>3</sup></b>	<b>.96</b>
S_FREQ = 0.9550 + 0.0045LAG - 0.00001LAG <sup>2</sup>	.66
<b>S_FREQ = 0.8172 + 0.0099LAG - 0.00005LAG<sup>2</sup> + 0.00000006LAG<sup>3</sup></b>	<b>.87</b>

Table 11: Regression results for error frequency/sec during the initiation period. O\_FREQ = overall; P\_FREQ = penetration; S\_FREQ = separation.

Model	R <sup>2</sup>
M MAG = 0.0905 + 0.0001LAG	.48
<b>M MAG = 0.0813 + 0.0003LAG - 0.0000005LAG<sup>2</sup></b>	<b>.61</b>
P MAG = 0.1242 + 0.000092LAG	.80
P MAG = 0.1224 + 0.0001LAG - 0.00000009LAG <sup>2</sup>	.81
<b>P MAG = 0.127 - 0.00006LAG - 0.0000011LAG<sup>2</sup> - 0.000000002LAG<sup>3</sup></b>	<b>.87</b>
S MAG = 0.1097 + 0.0002LAG	.75
S MAG = 0.0936 + 0.0005LAG - 0.0000008LAG <sup>2</sup>	.94
<b>S MAG = 0.0869 + 0.0008LAG - 0.000003LAG<sup>2</sup> + 0.0000000029LAG<sup>3</sup></b>	<b>.97</b>

Table 12: Regression results for error magnitude during the initiation period. M\_MAG = mean; P\_MAG = penetration; S\_MAG = separation.

Model	R <sup>2</sup>
O PER = 27.5725 + 0.0479LAG	.65
O PER = 22.6374 + 0.1433LAG - 0.0003LAG <sup>2</sup>	.88
<b>O PER = 19.0438 + 0.2854LAG - 0.0012LAG<sup>2</sup> + 0.0000016<sup>3</sup></b>	<b>.99</b>
A PER = 7.0992 + 0.0161LAG	.72
A PER = 5.8506 + 0.0403LAG - 0.00006LAG <sup>2</sup>	.86
<b>A PER = 5.6387 + 0.0487LAG - 0.0001LAG<sup>2</sup> + 0.000000093LAG<sup>3</sup></b>	<b>.87</b>
P PER = 2.5412 + 0.0131LAG	.78
P PER = 1.8854 + 0.0258LAG - 0.00003LAG <sup>2</sup>	.85
<b>P PER = 1.5750 + 0.0381LAG - 0.0001LAG<sup>2</sup> + 0.00000014LAG<sup>3</sup></b>	<b>.86</b>
<b>S PER = 11.8301 + 0.1987LAG - 0.001LAG<sup>2</sup> + 0.0000013LAG<sup>3</sup></b>	<b>.76</b>

Table 13: Regression results for the percentage of time spent in error during the movement period. O\_PER = overall; A\_PER = aiming; P\_PER = penetration; S\_PER = separation.

Model	R <sup>2</sup>
<b>O_FREQ = 0.0087LAG - 0.00005LAG<sup>2</sup> + 0.00000007LAG<sup>3</sup></b>	<b>.86</b>
A_FREQ = 0.1062 + 0.0001LAG	.53
A_FREQ = 0.0899 + 0.0004LAG - 0.0000008LAG <sup>2</sup>	.90
<b>A_FREQ = 0.0827 + 0.0007LAG - 0.000003LAG<sup>2</sup> + 0.0000000032LAG<sup>3</sup></b>	<b>.96</b>
P_FREQ = 0.1215 + 0.001 LAG - 0.000002LAG <sup>2</sup>	.66
<b>P_FREQ = 0.0890 + 0.0023LAG - 0.00001LAG<sup>2</sup> + 0.000000014LAG<sup>3</sup></b>	<b>.86</b>
<b>S_FREQ = 0.5567 + 0.0057LAG - 0.00004LAG<sup>2</sup> + 0.000000053LAG<sup>3</sup></b>	<b>.80</b>

Table 14: Regression results for error frequency/sec during the movement period. O\_FREQ = overall; A\_FREQ = aiming; P\_FREQ = penetration; S\_FREQ = separation.



Model	R <sup>2</sup>
M MAG = 0.3344 + 0.0011LAG	.87
<b>M MAG = 0.2846 + 0.0021LAG - 0.000003LAG<sup>2</sup></b>	<b>.93</b>
M MAG = 0.2700 + 0.0026LAG - 0.000006LAG <sup>2</sup> + 0.0000000064LAG <sup>3</sup>	.93
P MAG = 0.2514 + 0.0020	.86
P MAG = 0.1609 + 0.0037LAG - 0.000005LAG <sup>2</sup>	.92
<b>P MAG = 0.1325 + 0.0049LAG - 0.00001LAG<sup>2</sup> - 0.000000012LAG<sup>3</sup></b>	<b>.93</b>
S MAG = 0.1052 + 0.0010LAG	.99
S MAG = 0.0936 + 0.0012LAG - 0.0000006LAG <sup>2</sup>	.99
S MAG = 0.1025 + 0.0009LAG + 0.0000018LAG <sup>2</sup> - 0.000000004LAG <sup>3</sup>	.99

Table 15: Regression results for error magnitude during the movement period. *M\_MAG* = mean; *P\_MAG* = penetration; *S\_MAG* = separation.

Model	R <sup>2</sup>
ID = 1.8575 + 0.0081LAG	.98
<b>ID = 0.7046 + 0.0111LAG - 0.000008LAG<sup>2</sup></b>	<b>.99</b>
ID = 1.7901 + 0.0077 + 0.000015LAG <sup>2</sup> - 0.00000004LAG <sup>3</sup>	.99
IV = 1.9238 + 0.0073LAG	.97
<b>IV = 1.7889 + 0.0099LAG - 0.000007LAG<sup>2</sup></b>	<b>.98</b>
IV = 1.8040 + 0.0093LAG - 0.000003LAG <sup>2</sup> - 0.000000007LAG <sup>3</sup>	.98
IH = 1.9779 + 0.0074LAG	.96
<b>IH = 1.7522 + 0.0118LAG - 0.00001LAG<sup>2</sup></b>	<b>.99</b>
IH = 1.7800 + 0.0107LAG - 0.000004LAG <sup>2</sup> - 0.00000001LAG <sup>3</sup>	.99
MD = 1.9598 + 0.0084LAG	.98
<b>MD = 1.7930 + 0.0116LAG - 0.000008LAG<sup>2</sup></b>	<b>.99</b>
MD = 1.7694 + 0.0125LAG - 0.00001LAG <sup>2</sup> + 0.00000001LAG <sup>3</sup>	.99
MV = 2.0006 + 0.0079LAG	.98
<b>MV = 1.8459 + 0.0109LAG - 0.000008LAG<sup>2</sup></b>	<b>.99</b>
MV = 1.8675 + 0.0100LAG - 0.000002LAG <sup>2</sup> - 0.000000009LAG <sup>3</sup>	.99
MH = 2.1692 + 0.0073LAG	.94
MH = 1.9469 + 0.0116LAG - 0.00001LAG <sup>2</sup>	.97
<b>MH = 1.7633 + 0.0189LAG - 0.00006LAG<sup>2</sup> + 0.000000081LAG<sup>3</sup></b>	<b>.99</b>

Table 16: Regression results for difficulty and disruption ratings. *I* = initiating contact; *M* = maintaining contact; *D* = overall difficulty; *V* = visual disruption; *H* = haptic disruption.