

The Effects of Low Latency on Pointing and Steering Tasks

Sebastian Friston, *Student Member, IEEE*, Per Karlström, and Anthony Steed, *Member, IEEE*

Abstract—Latency is detrimental to interactive systems, especially pseudo-physical systems that emulate real-world behaviour. It prevents users from making quick corrections to their movement, and causes their experience to deviate from their expectations. Latency is a result of the processing and transport delays inherent in current computer systems. As such, while a number of studies have hypothesized that any latency will have a degrading effect, few have been able to test this for latencies less than ~50 ms. In this study we investigate the effects of latency on pointing and steering tasks. We design an apparatus with a latency lower than typical interactive systems, using it to perform interaction tasks based on Fitts’s law and the Steering law. We find evidence that latency begins to affect performance at ~16 ms, and that the effect is non-linear. Further, we find latency does not affect the various components of an aiming motion equally. We propose a three stage characterisation of pointing movements with each stage affected independently by latency. We suggest that understanding how users execute movement is essential for studying latency at low levels, as high level metrics such as total movement time may be misleading.

Index Terms—Latency, Indirect Input, HCI, Fitts’s law, Human Factors.



1 INTRODUCTION

ADVANCED graphical interfaces are commonly used to facilitate intuitive visualisation and manipulation of data as efficiently as possible. Some do this with abstractions such as widgets or manipulators. Others, such as pseudo-physical interfaces, exploit knowledge about natural object behaviour to allow more intuitive interaction techniques. For example, by constraining virtual objects by the laws of physics [1]. When successful, a user will take the same approach to tasks in this system as they would to such a system in the real world. This is achieved through the formation of the sensorimotor loop - the “continued correlation between proprioception and sensory data” [2]. To form and maintain this loop, the responses of the system to user input must meet the expectations the interface creates.

One way in which these responses can deviate from such expectations is in how fast the user receives them - the latency. Latency is defined as the time between a user’s action and the response to this action. Keeping latency low is important to maintain the perception of a correlation between a user’s action and the response to it. As a product of the inherent processing and transport delays within a computer system, latency will never reach zero [3]. Previous studies examined the effects of latency on a number of sensory modalities, from latency detection in immersive virtual environments, to its effects on indirect physical interaction. This latter modality has received considerable attention due to its ubiquity and importance, with many previous studies using motion primitives such as pointing tasks to investigate the effects of latency [4], [5], [6], [7], [8]. Only recently though has it become practical to build apparatus with latencies low enough that the limits of its effects may be found [9].

The human motor system has been modelled as a control loop, with inherent delays that place natural limitations on performance; movement cannot be coordinated on time-scales smaller than the inherent delay [10]. We therefore hypothesize that there may be a non-zero external latency which has no perceptible effect on the sensorimotor loop. Latency cannot be removed given current technology, but we can compensate for it. By understanding how latency affects the different modalities that create an effective user interface, we can distribute resources of computer systems to minimize negative effects and create a better user experience.

In this study we investigate the modality of indirect physical interaction, using familiar desktop based pointing and steering tasks. We create an apparatus similar to previous studies [4], [7], [8] but capable of much lower latencies. Our results, taken with those from studies on other modalities, will help guide the requirements for future interactive systems and better estimations of the effectiveness of existing ones. Further, our results have implications for future studies using physical tasks to investigate latency, as we show considering only total movement time and not its constituent parts may result in inconclusive measurements which hide the effects of latency.

2 PREVIOUS WORKS

An interface affording natural interaction can have a number of advantages. Exploiting the user’s intuition may reduce learning time or improve performance. For example, Smith et al. applied real world physical constraints to the objects in a 3D editor, decreasing the degrees of freedom of the objects in a familiar way. Users showed improved performance when interacting with the editor using 2D interaction techniques [11]. Where interfaces emulate physical systems, users will likely interact with their motor system. Even when abstractions are present, actions are still predominantly basic motion primitives such as reaching and pointing [1]. To encourage natural interaction, an interface must exercise the same functionalities of the user that are exercised day to day by the

- *Sebastian Friston is with University College London.
E-mail: sebastian.friston.12@ucl.ac.uk*
- *Per Karlström is with Maxeler Technologies Ltd.
E-mail: pkarlstrom@maxeler.com*
- *Anthony Steed is with University College London.
E-mail: a.steed@cs.ucl.ac.uk*

Manuscript received January 08, 2015; revised June 03, 2015.

real world. To do this the system must provide stimuli from which users can form percepts and react to them as if they were real [2]. Limitations of current technology prevent perfect emulation of these stimuli however. One of these limitations is latency, which is an unavoidable result of the inherent processing and transport delays in the computer system itself [3].

2.1 Models of the Motor System

A number of authors have constructed theoretical models to explain the operation of the visuomotor system. One such model is that of Botzer & Karniel [12]. The authors derived their model from observations of delay compensation behaviours. Participants performed Fitts's law style tests [13]. They were allowed to adapt to different latency conditions, and then the visual feedback was removed and at the same time the latency changed. By observing how user motion changed under this new condition, the authors tested where in the hypothesized control loop delay compensation was performed. Whether in the feedforward model, which plans the trajectory, or the feedback loop, where correction commands are issued based on visual feedback. Overshoot and undershoot were present in reaching tasks in unexpected delay conditions. This demonstrates dominance of an adapted visual feedback stage over the feedforward planning stage. They also found that while discrete reaching movements returned to baseline conditions (that is, the users no longer overshoot or undershoot), rhythmic ones do not. This suggests there is adaptation in the forward model, but it is dependent on movement type, leading to their model incorporating multiple pathways.

Beamish et al. considered the motor system as a Vector Integration to Endpoint (VITE) circuit. In the VITE circuit a continuous outflow of commands to the muscles are a result of the motor system attempting to reduce the difference vector between the intended target position and the present position. The commands are generated by the neuron population calculating the difference vector, based on the present position estimation from a population which integrates all previous movement commands. They note the VITE circuit as one of the earliest models to suggest how the movement characteristics described by Fitts's law are a result of underlying neurobiological mechanisms. The authors introduce time delay between the two populations into this model. By drawing comparisons with a servomechanism model, they show that for a system to be stable the gain (magnitude) of the movement commands must be below a value which is a function of delay. It should be noted that the model described above does not take into account visual feedback - or indeed any external delays. That is, even considering a system based only on proprioceptive cues the authors demonstrate a hard upper limit on performance [10].

Beamish et al. pursued this model, using it to estimate the inherent effective feedback delays in the motor system based on the results of previous Fitts's law style experiments. They expressed the performance of the VITE circuit (movement time) in terms of difference vector neuron population time constant, and feedback delay. They could then relate these parameters to the observed Fitts's law constants a & b . Using the measurements available from over 25 previous Fitts's law style studies, they found feedback delays between 0-112 ms, generally below 60 ms. They also found that the nature of the VITE circuit imposes a limit on the performance of unidirectional movement [14]. When this limit is expressed as a Fitts's law Index of Difficulty (ID), it happens to be the typical range employed by previous experimenters.

2.2 Measuring the Effects of Latency

An advantage of assessing an interface that emulates a physical system is that there is a clear baseline to compare it to: the real system. Considering latency and physical interaction, we can measure a user's performance to see how this degrades from the 'real world standard' as latency increases. A number of studies have done this using typical motion primitives, such as pointing and reaching tasks. Performance is defined in terms of completion time and error rate.

For example, Jay et al. used task completion time and error rate to measure the performance of users in collaborative physical manipulation tasks while experiencing delays in haptic feedback. 25-400 ms of latency was added during the experiments. The authors estimated a base latency of 14 ms (7 ms from the projector and 7 ms from the network). They found a strong interaction between latency and both error rate and task completion time. The authors also documented users employing the impact-perceive-adapt model of latency compensation. This states that latency begins to degrade performance, before the user is aware of it. Once the latency becomes large enough to cause a "breakdown of the perception of immediate causality" (the sensorimotor loop), the user adopts a 'wait and see' pattern. They act based entirely on predictions of the result of their motion, wait for a response, and then make corrections using the same technique. At this point real time correction ceases, and the user's response time consists almost entirely of the system delay [15].

2.3 Fitts's law

Most studies on physical interaction, such as that of Jay et al., use Fitts's law style tests. A good review of Fitts's law is by Seow [16]. Fitts's law is an emergent property rather than a description of the motor system operation. This is discussed by Bootsma et al. [17] and Huys et al [18]. Both sets of authors demonstrate that by observing the patterns of motion directly under different conditions, Fitts's law is a good summary of complex motor processes. However there is increased asymmetry in the amount of time spent in the acceleration stage vs. the deceleration stage as latency increases. The pattern of movement is significantly different between rhythmic and non-rhythmic movement, and as ID increases rhythmic pattern becomes more like the discrete pattern. This suggests multiple functionalities acting in parallel, such as in the model proposed by Botzer & Karniel [12]. Botzer & Karniel referred to Rhythmic/Non-Rhythmic as Slicing and Reaching respectively.

As a characterisation of the motor performance, Fitts's law has been observed a number of times under a range of conditions. Its repeatability and invariance make it valuable for testing the effects of various factors on user interaction. For example, Adam et al. measured the difference between egocentric guided movement and allocentric guided movement [19]. This was expanded on by Blinch et al., who found that the most significant effects occurred between the presence of allocentric markers and the preparation stage of movement [20]. Perrault et al. tested the scale effect using Fitts's law [21]. Jax et al. tested the effects of obstacles in the movement path [22].

2.3.1 Fitts's law and Latency

For the same reasons described above, Fitts's law has been used extensively to investigate the effects of latency. MacKenzie & Ware did one of the first studies in this area, reformulating Fitts's law to

account for additional movement time delay [4]. They estimated the base latency at 8.3 ms, and between 16 and 225 ms of latency was added. Pavlovych & Stuerzlinger suggest that the base latency could actually have been ~ 60 ms though [8]. Performance began to decrease significantly at the 75 ms condition. Ware & Balakrishnan used 3D reaching tasks in order to compare the effects of hand tracking delay with head tracking delay in an immersive Virtual Environment. They tested latencies between 87 and 337 ms. Teather et al. measured the effect of latency and jitter on performance in Fitts's style 2D tasks, and 3D object movement tasks, while looking for an effect of the type of tracker used. They measured the latency of their system at 73 ms and found that the performance degradation was equal for the tracker devices [7]. Pavlovych & Stuerzlinger performed a Fitts's law style test to determine the effects of jitter and latency. They found a strong interaction with latency and jitter. Further, with low jitter the effects of latency were dominant, but the jitter degraded performance at a higher rate than latency. The authors measured the base latency of their system at 33 ms, and added up to 100 ms [8]. Chung & So considered that latency may affect the stages of movement differently. They studied the effects of latency in Fitts's law style tests but on target width and distance separately. There was strong interaction between latency and target width, but not target distance [6].

2.4 The Steering law

There is evidence ([12], [10]) that the motor control system consists of multiple complex elements, some acting in parallel, and that the effects of latency on these is not equivalent. Thus in our experiment, aside from a Fitts's law-style task, we introduce a second task based on the Steering law. It is designed to exercise the real-time correction functionalities predominantly and force the user to continually change goals as they move.

The Steering law was introduced by Accot & Zhai. It was originally derived from Fitts's law, considering a path as a sequence of goal crossing tasks. The completion time was the measure of performance, and was estimated to be the sum of the time to complete the individual goal crossing tasks, that make up a path [23]. It was extended by Kulikov et al., who used the concept of effective width to demonstrate that the Steering law was even more accurate than originally shown [24].

Like Fitts's law the Steering law has been used to investigate the effect of specific factors on user performance. Liu et al. investigated which path properties affected user performance. The path properties considered were curvature and width [25]. Liu & Liere continued to investigate the effect of these properties changing within a path. In their test the path was presented as a tube. Participants were encouraged to remain within it by pushing a ball through it with the cursor. We model our implementation of the Steering law task on theirs. On examining the user movements, they assert that the behaviour does not resemble a goal crossing task, as much as a set of small ballistic movements [26].

Pavlovych & Stuerzlinger investigated the impact of latency on tracking tasks. While this task is analogous to the Steering law task, the authors point out that the Steering law itself does not apply. This is because there are no boundaries to movement outside of the target area, and the user is required to correct velocity as well as direction. The experimental setup had a base latency of 20 ms, and an additional latency of 30-150 ms. The authors observed a significant effect of latency on tracking accuracy, and that it was not symmetric: users had a smaller error perpendicular to the

target, than tangential. The latencies that could be tolerated before a significant interaction was visible were higher than in previous studies (50 ms for latency and 40 ms for jitter). Another interesting observation was that performance decreased for the condition with the lowest additional latency (20 ms), improved between 20-50 ms, then for latencies above 50 ms degraded again but at a slow rate [27].

2.5 Investigation of very low latencies

The closest study to ours is that of Jota et al [9]. They studied the effects of latency on direct interaction surfaces, with their HPT (High Performance Touch prototype) - a touchscreen with a latency of less than 1 ms. A number of previous studies have investigated the effects of latency on direct touch interaction, but none at such low levels. Participants performed Fitts's law style tests. Of particular interest in this study, is that the user received visual feedback from both their non-latent hand and the latent cursor simultaneously. How the potentially conflicting stimuli affect performance is not clear. Participants showed a range of behaviours in response to the latent cursor, from ignoring it completely, to leading it, to slowing their movement so that it remained under their finger at all times. The additional latencies were between 1-50 ms. The authors reported no observable difference in performance between latencies of 1 ms and 10 ms. A linear regression fit suggested the performance floor may not exist. By segmenting the movement into stages, the authors demonstrate the effects of increasing latency on these are not symmetric, as Chung & So and Bootsma et al. showed for increasing ID [6], [17].

3 EXPERIMENT

A number of studies have used performance in motor tasks to detect the effects of latency. Few though have investigated latencies at very low levels. Jota et al. found a potential floor for direct interaction tasks [9]. Indirect interaction techniques however remain important for both 2D and 3D interfaces. They can exceed direct interaction in both efficiency and precision [1]. We therefore continue the investigation into indirect interaction.

3.1 Apparatus

To conduct the investigation an interface with very low controllable latency was required. The indirect input Fitts's law and Steering law tests require a 2D interface. The participants interacted through a cursor, which had to respond to the user within the shortest amount of time possible. As described by Mine, latency consists of tracker delay, processing delay, rendering & display scan-out delay, and transport delays between those stages [28]. By probing and optimising the latency between different parts of our system we constructed a system with a latency of ~ 6 ms using mostly off-the-shelf components (Figure 1).

3.1.1 Tracker

The tracker was a Kingston Mouse-in-a-Box optical mouse, with the Control-Display gain set to 1. Many newer mice, such as this one, can be sampled at 1kHz. The mouse device, connected via USB, was polled directly by our application, avoiding the event system of the operating system.

3.1.2 Rendering

To drive the display we implemented our own display controller on a Maxeler Dataflow Engine (DFE) [29]. DFEs are processing cards which execute dataflow computations. Algorithms are described as dataflow graphs, which are implemented as pipelines of single-purpose cores executing in parallel in space, rather than sets of operations executed by a small number of multipurpose cores such as on CPUs. This spatial parallelism provides high performance, and a deterministic latency at levels lower than that achievable by conventional GPUs. We designed an algorithm to render 2D sprites, driven by an application running on the CPU, and described it as a dataflow graph using MaxCompiler, Maxeler’s toolchain. Our algorithm is deterministic. Knowing exactly how long it takes to compute one pixel, we can begin computation of a pixel using the latest tracking data, that much time before it is required for transmission. At no point in the system is a frame buffered, on each clock tick a new pixel is completed and transmitted to the display.

We used the parallel port of the host computer and an output from the DFE to probe the latency of the rendering stage of our system. The DFE illuminated an LED on receipt of a specific input. High speed video monitored the input device, and the LED. This arrangement was chosen as it allowed us to monitor both the input device and the scan-out of the display, with no further instrumentation. The latency between the input and the LED was below the temporal resolution of the video (1 ms). In the best case scenario the user begins just prior to the cursor is drawn. In this case the latency is between 1-2 ms - predominantly the mouse sampling time. In the worst case the user moves immediately after. In this case the latency is 8-9 ms. This is the mouse sampling time (1-2 ms), the rendering time (<1 ms) and the period of one frame on our display (6.9 ms). We expect the latency to be ~5 ms on average. We measured the total end-to-end latency of our apparatus using the cross-correlation variant of Steed’s Method. Correct operation of the apparatus was confirmed by measuring the latency throughout the investigation, between each participant. The baseline latency was measured at 6 ms, with the tolerances described for the measurement method [30].

In our renderer, a number of sprites and a background map are composited to make a frame. The content and transformations of these sprites make up the renderer state. The renderer maintains the state, which is updated asynchronously by the CPU. With Maxeler’s assistance, a modification was made to MaxCompiler, which allowed DVI compliant display data to be output directly from the DFE. We also constructed an electrical interface that would allow the DFE to drive any DVI receiver.

At ~1 ms, the latency of our system from input to video signal output is much lower than previous apparatus. We are limited by display technology however. The display scan-out time increases our end-to-end latency to 6 ms. Beyond this, persistence of the image on the monitor can cause the perception that latency is greater than the average frame period. This is because the stimuli at any time is a blur between the current stimuli and the previous one. We selected a highly responsive, high frame rate monitor (an ASUS VG248QE), minimizing perceived latency due to both scan-out delay and persistence. The limitations in available display technology are shared by previous authors. Out of the aforementioned studies only Jota et al. secured a better performing display than the VG248QE. They did this by building a custom display based on a Digital Micro-mirror Device driven in a very low chromatic range [9].

3.1.3 Processing and Transport

Our system was based around an Intel Core i7 PC running CentOS 6. The tests were implemented in a thread running with real-time priority, controlled by a non-realtime manager application. The renderer was accessed using Maxeler’s low-latency API for communicating with the DFE via the PCIe bus. Like the mouse access, this makes use of polling, rather than events. The real-time thread communicated with the managing application via flags in memory. We profiled the thread to ensure that we only used calls which would not cause it to yield unintentionally. The thread was given the highest priority. The result was that the thread was never pre-empted, and latency due to time-slicing of the CPU was not introduced.



Fig. 1. Experimental apparatus that the participants interacted with.

3.2 Participants

30 participants (19 M/11 F) with an average age of 27 (Standard Deviation: 4 years) from within University College London were recruited for the study. Participants were paid £5 for taking part.

3.3 Procedure

Participants were seated ~0.6m in front of the display, their right hand being obscured by a black cloth. They were invited to move the chair, display and mouse. Once comfortable, they were shown the two tasks and allowed to practice each for as long as they wished. All participants were instructed to move as fast as possible. The participants spent 20-30 minutes completing the actual tests. The time to complete the whole experiment was 30-50 minutes. Our experimental design is very similar to the one-directional tapping task described in ISO9241-9 [31]. We deviated by asking users to make discrete movements, rather than repeated rhythmic movements. This is because the motor system behaves differently during these two types of motion [18], [12]. Further, the seminal works using Fitts’s law to investigate latency, such as that of MacKenzie & Ware [4], use discrete tasks.

3.4 Tasks

3.4.1 Fitts’s law

For the Fitts’s law style tests, participants saw a box on the screen ~2cm x 2cm, which remained throughout all the tests (the staging area). Clicking on this box would start the test, and a target would become visible to the right. Participants were instructed to click on the target as fast as possible, then in their own time move back to the staging area. Clicking the staging area a second time would begin the second test, and they were to repeat this until all tests were complete.

3.4.2 Steering law

For the Steering law tests, users were presented immediately with a 2D path, and at the start of the path, a green ball. They were instructed to push the ball through the path, by placing the cursor behind the ball and moving it forward through the path. Users were again told to maximise speed, and were told that keeping the cursor within the path would be the fastest way to complete the tests.

Examples of the stimuli seen by the users are in Figure 2.



Fig. 2. Images of the stimuli the participants were exposed to.

3.5 Design

The experiments had three independent variables: Latency, Width and Distance (Fitts's)/Curvature (Steering). For both Fitts's law and the Steering law there were four conditions of spatial difficulty, summarised in Table 1. For each condition there were six additional latencies (0, 10, 20, 30, 50, 80) for a total of $2 \times 2 \times 6$ (24) unique conditions. Unique Fitts's law conditions were repeated 8 times, and Steering law conditions 5 times. The repetitions were averaged for each participant, resulting in 720 data points for the Fitts's law tests and 720 for the Steering law tests. The tasks had low entertainment value, and fatigue was a concern. Since we expected the effect to be small, we optimised for a high number of latencies and repeats at the expense of spatial difficulty range. The widths and distances were informed by pre-trials. The range of IDs found by these matched those of MacKenzie & Ware, and those estimated by Beamish et al. [4], [10]. The IDs were calculated using MacKenzie's method [32].

The Steering law paths were manually created, with one designed to emphasize sharper higher rate turns (predominantly exercising the wrist) and other sweeping turns to exercise the upper arm (classed as curvatures 2 & 1 respectively). The IDs were calculated with Accot & Zhai's method [23]. The curves are not produced from any predictable function. This was deliberate, to prevent any unanticipated motor process (such as that used for reciprocal movement) from hiding the effect of latency on the on-line correction processes. Conditions were distributed to maximize the difference between sequential latencies. Within this constraint the widths and distances/curvatures were distributed randomly. All participants received the same conditions in the same order.

4 RESULTS & DISCUSSION

4.1 Pointing Tasks

We measure Movement Time (MT) to be from the time the user clicks the staging area, to the time they click the target. Figure 3 shows MT for each of the latencies. As expected MT increases with ID, with a jagged appearance due to the small number of spatial (Width & Distance) conditions that do not have overlapping IDs [8]. We clearly see an increase in MT with high latencies, but not for low latencies. This is better illustrated in Figure 4 which shows how MT changes with latency for each condition.

TABLE 1
Spatial Difficulty Conditions for both types of task.

Condition	Fitts's law		Index of Difficulty
	Width (cm)	Distance (cm)	
1	0.25	4	4.09
2	0.25	11	5.49
3	0.9	4	2.44
4	0.9	11	3.72
Condition	Steering law		Index of Difficulty
	Width (cm)	Curvature	
1	0.4	1	45.37
2	0.4	2	50.63
3	0.7	1	25.92
4	0.7	2	28.92

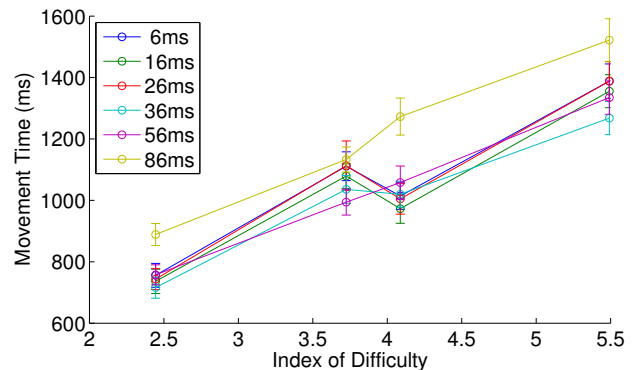


Fig. 3. Movement Time against Index of Difficulty, for all latency conditions. Error bars indicate confidence intervals at 95%.




Fig. 4. Movement times for each condition against latency. Error bars indicate confidence intervals at 95%.

4.1.1 Comparison with Previous Works

Studies conducting experiments most comparable with ours include [4], [7], [8]. All studies included Fitts's law style tests using mice, with latency as the independent variable.

- MacKenzie & Ware [4] investigated latencies estimated to be between 68 - 315 ms [8].
- Teather, et al. [7] investigated latencies measured at 35 - 255 ms.
- Pavlovych & Stuerzlinger [8] investigated latencies measured between 33 - 133 ms.

For [7] and [8] the latencies measured are the total end-to-end system delay, the same as measured by us. We first consider only the higher latency conditions (36, 56, 86 ms) which are directly comparable with the previous studies above.

We fit a model using multiple linear regression and show a significant interaction with width ($\beta = -499.81, t(356) =$

$-22.22, P < 0.001$), distance ($\beta = 37.48, t(356) = 17.95, P < 0.001$) and latency ($\beta = 4.02, t(356) = 11.30, P < 0.001$). We then fit MaxKenzie & Ware’s model to our data and show an almost identical r^2 value (0.995 (ours) vs. 0.967 (theirs)). Finally we perform a one-way ANOVA as done by Teather, et al. and Pavlovych & Stuerzlinger showing a similarly significant interaction $F_{2,357} = 21.32, P < 0.001$. All studies showed a significant almost identical multiplicative effect of latency with ID. We show the same effect for the overlapping latency conditions in our study. This is illustrated in Figure 5, which compares our results with those of previous studies. At lower difficulties our results appear slightly higher than previous works. The relationship though is identical, and our results are well within the inter-study variance. Thus for the higher latency conditions our experiments reproduce previous results.




Fig. 5. Movement time for the 36, 56 & 86 ms conditions of the current study, compared with the overlapping conditions from Pavlovych & Stuerzlinger (PS) [8], MacKenzie & Ware (MW) [4] and Teather, et al. (T) [7].

4.1.2 Deviation at Low Latency Conditions

When we consider the low latency conditions, our results begin to diverge from the expectations of ourselves and other authors. Considering only the lower levels of latency (6, 16, 26 ms), multiple linear regression demonstrates no significant interaction between movement time and latency ($t(356) = -0.27, P = 0.78$) and neither does ANOVA ($F_{2,357} = 0.48, P = 0.62$). We hypothesised a non-linear response as latency decreases. As shown in Figure 4 though, the linear relationship dissolves at higher latencies than we would expect. A clear correlation between MT and latency does not form until the latencies reach 50-80 ms, while studies have found an effect at far lower levels [9]. Further, in some cases user performance appears better in high latency conditions than in low latency conditions.

There have been hints of this effect in previous studies. In Fitts’s law tests by MacKenzie & Ware [4] and Teather et al. [7] there were IDs at which users had near identical performance at two different values of latency. The differences were slight though and the equivalence could be argued to be measurement error. In pointing tests by Pavlovych & Stuerzlinger [8] the effect is more pronounced, with the 83 & 33 ms conditions appearing to alternately outperform each other depending on the ID. In a target tracking test [27] users had a reduced tracking error at 50 ms compared to 25 ms. Although this interaction was proved not statistically significant, the authors suggest it could be caused by the users overcompensating and moving in front of the target at lower latencies. Another explanation is that users are more familiar with computer mice having latencies around 50 ms.

Until now these anomalous results have not warranted further investigation. Our tests however show a pronounced & repeatable effect. Observing the behaviour of the user during the task more closely reveals a possible cause. We suggest it is a result of the independent affects of latency on different stages of movement, happening at levels well below those at which performance supposedly improves.

While movement time is a useful metric, it does not allow for appreciation of the underlying processes. A number of works have hypothesised the motor system as a feedback loop, with an initial impulse followed by some form of continuous control. One way to characterise this has been to examine the symmetry of movement velocity profiles around the point of peak-velocity. As task difficulty increases so does the proportion of time spent correcting movement in the second - deceleration - stage [33]. An example of this is given by Bootsma et al [17]. We show that this deceleration stage may be further subdivided, into what we term the acquisition, and correction stages (defined below). Further, the impact of latency on these is not symmetrical.

Various parameters of kinematic profiles have been examined to gain insight into motor system functionality. There are different schemes to partition kinematic profiles. Partitioning based on peak-velocity is one example [33]. Another is that used by Meyer et al to partition motion into a primary movement and optional correction sub-movements for the two-component motor system model [34].

Bootsma et al previously used the peak-velocity scheme to quantify the effects of Fitts’s law test parameters on the kinematic profile [17], and it has been used to investigate multiple theories of motor system operation [33]. Examining the effects of latency on the kinematic profile with respect to specific motor system models such as Meyer et al’s however may provide new insights and is a subject for future work. For this investigation we partition the aiming motion into three stages:

- Acceleration** The time between the user beginning to move, and reaching their peak velocity.
- Acquisition** The time between the peak velocity and the user first reaching the target.
- Correction** The time it takes the user to settle and complete the task once the target has been reached.

Under very low latency conditions the majority of the time is spent in the acceleration and acquisition stages, so the correction stage is typically the time it takes the user to click the mouse button. Under high latency conditions the user overshoots and so the time in this stage is extended. The breakdown of the total MT into stages can be done by defining kinematic markers (e.g. the sample with peak-velocity) and using the position and timing data in the log files. The breakdown is shown in Figure 6.

As would be expected of a pre-planned impulse, multiple linear regression shows a strong interaction between the acceleration stage and distance ($\beta = 8.5, t(687) = 30.02, P < 0.001$), but not width ($t(687) = 0.97, P = 0.33$) or latency ($t(687) = 1.27, P = 0.2$). Performing multiple regression on the acquisition and correction periods independently, show the effects of latency are strong but asymmetric. The results are shown in Table 2. The r^2 values for the acquisition stage and correction stage are 0.832 and 0.682, respectively. The error degrees of freedom for both is 687.

As latency increases, the time in the acquisition stage decreases. This is accompanied by an increase in average velocity for the stage. That is, the user covers the same distance during this stage as before, but makes the motion in a shorter amount of time. Conversely, time




Fig. 6. Mean time in each task stage, for each latency. Error bars indicate confidence intervals at 95%.

TABLE 2

Multiple linear regression results for separate stage movement times.

Variable	Stage			
	Acquisition		Correction	
	Coefficient	P-Value	Coefficient	P-Value
Width	-181.82	<0.001	-243.09	<0.001
Distance	41.28	<0.001	-6.06	<0.001
Latency	-2.25	<0.001	3.19	<0.001

in the correction stage increases. Total movement time is the sum of all three stages. Since the correction time typically increases faster than acquisition time decreases, higher latencies generally result in higher movement times. The effect on correction time is non-linear however, with large increases occurring only at high latencies. Therefore there is a subset of latencies, within which acquisition time decreases faster than correction time increases, resulting in a lower movement time overall. This is shown in Figure 7.

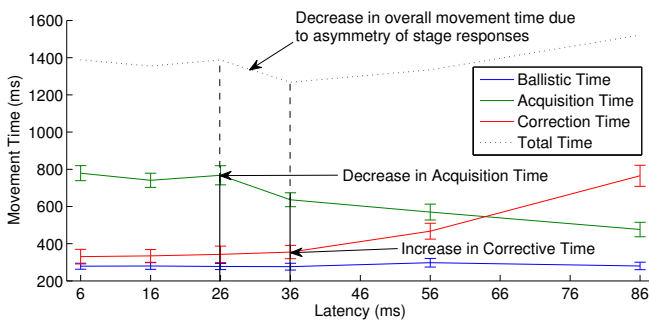


Fig. 7. Mean time for the different stages of movement for condition 2 with latency. The graph has been annotated to illustrate how movement time changes with latency due to the differences in the response of the stages. Error bars indicate confidence intervals at 95%.

At latencies between 26-36 ms, the user does not need to make significant corrections once the target is reached, but neither does their deceleration profile match the conditions between 0-26 ms. They continue their quick movements causing them to move farther and faster than they likely intended. However, the latency is still low enough that the overshoot, if present at all, is marginal, and the correction stage is not significantly confounded. We cannot say with certainty the cause of this change in profile. One possibility is the transition to a third and unanticipated compensation process. Another is that the beginning of the deceleration is delayed due to interference with the motor processes. The result though, is that total movement time decreases with the decrease in acquisition time, until the point at which the correction stage is significantly

affected, negating and then eclipsing the acquisition time gains. If this is the case, it is likely not optimal functioning of the control loop however. Interfering with the deceleration process may prevent normal trajectory modifications as the target is approached, benefiting only the subset of tasks which can be completed without these. In most cases without the controlled deceleration stage, overshoots are likely to occur and take considerable time to correct.

4.1.3 Latency Thresholds

With the movement stages split up we are in a better position to observe when latency begins to impact the function of the motor system. We perform ANOVA with the various pairs of latency for both movement stages. (Recall that the time in the acceleration stage is independent of latency.) Pairs between which the time in the stage differs significantly are shown in Table 3. For all tests *between group degree of freedom* = 1 and *within group degree of freedom* = 58. From this it is clear that for these tests latency begins to have an effect at ~16 ms, ~6 ms larger than that found by Jota et al. for direct interaction.

TABLE 3

P-Values for ANOVA between pairs of latencies within each condition.

Condition	Latency Condition Pairs				
	6-16	16-26	26-36	36-56	56-86
	Acquisition Stage				
1		<0.05		<0.05	
2			<0.05	<0.05	<0.05
3			<0.05	<0.05	
4			<0.05	<0.05	
	Correction Stage				
1		<0.05	<0.05	<0.05	<0.05
2			<0.05	<0.05	<0.05
3			<0.05	<0.05	<0.05
4				<0.05	<0.05

4.1.4 Effects of Target Width & Distance

From Figure 4 we see that the unexpected decrease in movement time is largest for conditions 2 & 4, which have the largest target distance. This is intuitive. The gain is a result of confounding the acquisition stage: the longer this stage lasts, the larger the effect on total movement time. From Table 3 we see that correction time begins to be affected at higher latencies than acquisition time. The velocity for the correction stage is lower than for the acceleration or acquisition stages. This may make the processes of this stage more tolerant to delay. Any benefits are short lived though. When latency does begin to affect this stage, the performance degradation is severe and increases rapidly. As shown in Table 2, while distance has a larger effect on acquisition time than correction time as would be expected, neither stage is a product only of width or distance.

4.1.5 Modelling the effects of latency

MacKenzie & Ware modified Fitts's law to account for the multiplicative affects of latency. In addition we created a linear model, with predictors of width, distance and latency. The response variable is the total MT. This model is the identical to one consisting of the sum of the linear models for each of the movement stages. We consider how well both models fit our data. We selected three latencies, and for each in turn, removed the conditions with those latencies from our results. The models were fitted to the remaining data, and then used to estimate the results of omitted conditions. The error of these estimations were averaged. The results are shown in Table 4.

TABLE 4

Mean estimation errors of our linear model and MacKenzie & Ware's variant of Fitts's law, after estimating the movement time of a specific condition, the results of which had been removed.

Latency Condition Predicted	Estimation Error (ms)	
	Our Linear Model	MacKenzie & Ware's Fitts's law variant
6	81.5934	78.4651
36	68.1741	68.3698
86	199.0222	201.4407

As the latency of the condition being predicted increases, so does the prediction error. In both cases this error is caused by an underestimation of MT. We hypothesise the response to latency is non-linear, and if it is, this is what would be expected. Since the models are created from a region of the non-linear response that has a weaker correlation with latency, they underestimate MT as latency increases. We do not attempt to model this non-linear relationship, as we do not have data from a wide enough range of conditions to do so. Acquisition time is shown to decrease with latency, though clearly this cannot continue indefinitely. We would expect it to continue down until it reaches a floor at which it remains a constant multiple of width and distance. Since the acceleration stage is constant, and we expect the acquisition time to degenerate to constant, we must conclude that correction time will come to resemble the relationship described by MacKenzie & Ware's model. Our latency conditions are not extensive enough to test these hypotheses however so any model we created would be incomplete. This is a subject for future work. The models created if all observations are considered are shown in Equation 1 (MacKenzie & Ware) and 2 (Ours). These models have r^2 values of 0.384 and 0.685, respectively. The error degrees of freedom of our model is 716.

$$MT = 296 + (184 + 0.38LAG)ID \quad (1)$$

$$MT = 941 - 453Width + 45Distance + 1.6Latency \quad (2)$$

Two additional Fitts's law metrics that are commonly used are throughput and error. In our study, a trial was not complete until the target had been acquired. Therefore error is approximated by Correction Time. Throughput (or Bandwidth) is given as the ration between ID and MT [8]. We calculated both error and throughput and found that they had similar profiles to the MT for each condition. It is not clear how we could unambiguously separate these high level metrics into their constituent parts though, so we unable to determine any more from them than we are total MT.

4.2 Steering Tasks

The MT for a Steering law test is considered to be time between the cursor first touching ball, and the ball reaching the last point on the path. Like the Fitts's law tests, MT increases with ID, and there is generally a multiplicative effect of latency with ID (Figure 8). Although this degenerates at lower latencies (Figure 9). We are not aware of any previous studies that have investigated the effect of latency on the Steering law itself. The experiment closest to ours is that of Pavlovych & Stuerzlinger, in which the authors investigated the effects of latency and jitter on performance in tracking tasks [27]. As velocity was fixed in their tracking task, the performance measure was the error rate, defined as the distance from the target. Their participants performed best with 50 ms

of latency (the lowest latency was 25 ms). We observe similar profiles for our conditions, in terms of MT (Figure 9) and error rate (Figure 10) though our participants performed best at slightly lower latencies.




Fig. 8. Movement Time against Index of Difficulty for all latencies. Error bars indicate confidence intervals at 95%.




Fig. 9. Movement Time against Latency for each Steering law Condition. Error bars indicate confidence intervals at 95%.

Path tracing is likely to require constant acceleration. Intuitively then steering tasks could be thought of as a sequence of correction movements. However, after comparing Fitts's law and the Steering law, Liu et al. postulate that behaviour is more like a series of ballistic tasks [25]. The similarity of the responses in the Steering law & Fitts's law tests in our experiments suggest both tasks use similar processes. We performed the Steering law tests as it was hypothesized they would exercise different motor processes than Fitts's law. This could help disambiguate the Fitts's law results. There is no evidence to suggest this is the case however, and segmenting movement stages in steering tasks is not as straightforward as in discrete pointing tasks. We are unable to offer an explanation for the apparent non-linear effects of latency, other than it is possibly the same interaction between motor processes seen in the discrete tests.

5 CONCLUSION & FUTURE WORK

Latency is known to impact performance in motor tasks. User performance has been characterized by models such as Fitts's law, which have been extended to include the effects of latency. Authors have commented previously that Fitts's law is likely to degrade at extreme values. One example is the negative intercept of the linear model, which would result in a zero or negative MT for low enough IDs. Clearly the real response must deviate from this model. We suspected a similar deviation would occur from the latency model for very low values, and designed an experiment to test this.

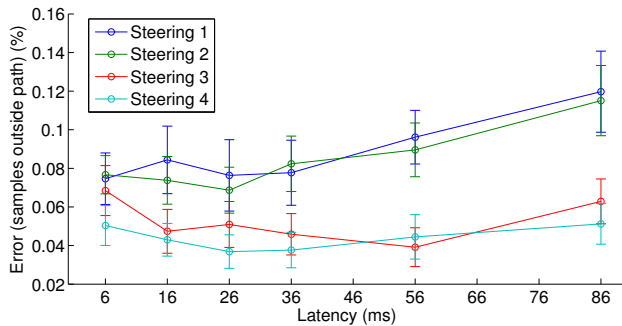


Fig. 10. Error Rate in percentage of samples that occurred outside the path, against latency for all Steering law conditions. Error bars indicate confidence intervals at 95%.

We constructed a system with a system latency of ~ 1 ms and display latency of ~ 5 ms. Informed by pre-trials, we selected conditions which matched those from previous works, and which happen to be those at which the motor system can theoretically operate optimally. Where conditions overlap we compare our results and find no significant differences. For our lower latency conditions however, we find a significant and unexpected effect.

Our results show that for some conditions, higher latencies can result in lower movement times. This effect has been hinted at in previous studies, though not significantly enough to pursue. We caution that movement time is just one metric and does not necessarily mean performance is improved. On closer inspection we suggest that the effect may be explained by the independent, degrading effects of latency on the processes of the motor system.

5.1 The non-linear response of latency

Even the simplest models of the motor system consider a set of separate processes, cooperating to execute smooth motion. So far though, the impact of latency has been modelled as linear. It is theorized that latency interferes with the ability to make quick corrections to motion, slowing the user down. Our results show that while this is true, it may be obscured by the interaction of latency with a stage of movement which we term the acquisition stage. When the user enters this stage they are moving at their highest velocity, and during the stage begin to decelerate. Interfering with this stage then could result in a decrease in movement time, as the user moves further and faster than they would if they had full control of their motor system. Usually this results in overshoot. With certain task difficulties and low enough latencies however, the corrections required are minimal, resulting in a lower overall movement time. It appears as if latency is improving user performance. The reality though is just another trade-off analogous to that between speed and accuracy - but one the user has no control over. This explanation also implies that the motor system is naturally conservative.

5.2 Thresholds of latency

Our conditions are not extensive enough to derive better models of the effects of latency on MT than the existing linear ones. By performing ANOVA for latency condition pairs though we can determine an initial range where latency does not appear to have a significant effect. The value of this threshold (16-26 ms) is not as important as the fact that such a threshold can exist. If a threshold exists for our apparatus, one may exist for more complex installations.

We hypothesize that above specific latencies, the impact of latency on the acquisition stage will degenerate. It will become constant, like the acceleration stage, and the correction stage can be modelled by the original Fitts's law with latency introduced by MacKenzie & Ware. Fitts's law describes an observation rather than the operation of the underlying system, and so it transcends the revisions of motor system models. Conversely though, it only applies to conditions within a certain range. If predicting user performance across the full range of conditions is important, models which describe the contribution of all movement processes will have to be derived.

5.3 Investigating Latencies at Low Levels

The Steering law tests were included in order to disambiguate the results of the Fitts's law tests, where the effects of latency on different motor processes may not be clear. In fact though, it was the Steering law test behaviour we could not explain, due to the inability to quantify participant behaviour beyond MT and error rate. The Fitts's law test is valuable in investigating latency, so long as metrics beyond that of MT are considered. It is not clear though, whether the floors we have supposedly found are functions of the motor system, or the motor system & task difficulty & apparatus. Ideally a task would be designed, which both approximated real interaction primitives, and exercised the motor processes in such a way that the response to latency remained linear, or at least linear with a clearly defined floor. That the effect of latency on the movement stages is independent is significant. If only the sum of these stages is considered, the effect of latency may be obscured. Recall that no interaction with latency was demonstrated when considering the total movement time for the lowest latency conditions. What is not clear is the extent to which each movement stage is dependent on pre-planning. In the previous tests we have considered, the requisite movements are predictable for the user. Even in Pavlovych & Stuerzlinger's tracking task, Lissajous Curves were used which made the motion of the target predictable for most of the experiment. With improved models of the motor system it may be possible to isolate and test the processes involved with each movement stage separately. Latency is detrimental to interfaces facilitating continuous or pseudo-physical interaction with complex datasets or systems. This is especially true for those affording natural interaction such as virtual environments. If we are to continue to investigate low latency, it may be worthwhile to pursue a new interaction benchmark.

REFERENCES

- [1] J. Jankowski and M. Hachet, "A Survey of Interaction Techniques for Interactive 3D Environments," *Eurographics 2013 - State of the Art Reports*, pp. 65–93, 2013.
- [2] M. Slater, B. Lotto, M. M. Arnold, and M. V. Sanchez-Vives, "How we experience immersive virtual environments: the concept of presence and its measurement," *Anuario de Psicología*, vol. 40, no. 2, pp. 193–210, 2009.
- [3] G. Papadakis, K. Mania, and E. Koutroulis, "A system to measure, control and minimize end-to-end head tracking latency in immersive simulations," *VRCAI '11 Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry*, pp. 581–584, 2011.
- [4] I. S. MacKenzie and C. Ware, "Lag as a determinant of human performance in interactive systems," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. New York: ACM Press, 1993, pp. 488–493. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=169059.169431>
- [5] C. Ware and R. Balakrishnan, "Reaching for Objects in VR Displays : Lag and Frame Rate," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 1, no. 4, pp. 331–356, 1994.

- [6] K. M. Chung and R. H. So, "Effects of Hand Movement Lag on Discrete Manual Control Tasks in Virtual Environments," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 43, no. 22, pp. 1210–1213, Sep. 1999. [Online]. Available: <http://pro.sagepub.com/lookup/doi/10.1177/154193129904302209>
- [7] R. J. Teather, A. Pavlovych, W. Stuerzlinger, and I. S. MacKenzie, "Effects of tracking technology, latency, and spatial jitter on object movement," *2009 IEEE Symposium on 3D User Interfaces*, pp. 43–50, 2009. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4811204>
- [8] A. Pavlovych and W. Stuerzlinger, "The tradeoff between spatial jitter and latency in pointing tasks," in *Proceedings of the 1st ACM SIGCHI symposium on Engineering interactive computing systems - EICS '09*. New York, USA: ACM Press, 2009, p. 187. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1570433.1570469>
- [9] R. Jota, A. Ng, P. Dietz, and D. Wigdor, "How fast is fast enough?" in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, no. 1, 2013, p. 2291. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2470654.2481317>
- [10] D. Beamish, S. Bhatti, J. Wu, and Z. Jing, "Performance limitations from delay in human and mechanical motor control." *Biological Cybernetics*, vol. 99, no. 1, pp. 43–61, Jul. 2008. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18481080>
- [11] G. Smith, T. Salzman, and W. Stuerzlinger, "3D Scene Manipulation with 2D Devices and Constraints," *Proceedings of Graphics Interface*, pp. 135–142, 2001.
- [12] L. Botzer and A. Karniel, "Feedback and feedforward adaptation to visuomotor delay during reaching and slicing movements." *The European Journal of Neuroscience*, vol. 38, no. 1, pp. 2108–23, Jul. 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23701418>
- [13] P. M. Fitts, "The information capacity of the human motor system in controlling the amplitude of movement." *Journal of Experimental Psychology*, vol. 47, no. 6, pp. 381–391, 1954. [Online]. Available: <http://content.apa.org/journals/xge/47/6/381>
- [14] D. Beamish, S. Bhatti, C. S. Chubbs, I. S. MacKenzie, J. Wu, and Z. Jing, "Estimation of psychomotor delay from the Fitts' law coefficients." *Biological Cybernetics*, vol. 101, no. 4, pp. 279–96, Oct. 2009. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/19862551>
- [15] C. Jay, M. Glencross, and R. Hubbard, "Modeling the effects of delayed haptic and visual feedback in a collaborative virtual environment," *ACM Transactions on Computer-Human Interaction*, vol. 14, no. 2, Aug. 2007. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1275511.1275514>
- [16] S. Seow, "Information Theoretic Models of HCI: A Comparison of the Hick-Hyman Law and Fitts' Law," *Human-Computer Interaction*, vol. 20, no. 3, pp. 315–352, Sep. 2005. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1207/s15327051hci2003_3
- [17] R. Bootsma, L. Fernandez, and D. Mottet, "Behind Fitts' law: kinematic patterns in goal-directed movements," *International Journal of Human-Computer Studies*, vol. 61, no. 6, pp. 811–821, Dec. 2004. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1071581904001065>
- [18] R. Huys, L. Fernandez, R. J. Bootsma, and V. K. Jirsa, "Fitts' law is not continuous in reciprocal aiming." *Proceedings. Biological sciences / The Royal Society*, vol. 277, no. 1685, pp. 1179–84, Apr. 2010. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2842815>
- [19] J. J. Adam, R. Mol, J. Pratt, and M. H. Fischer, "Moving farther but faster: an exception to Fitts's law." *Psychological science*, vol. 17, no. 9, pp. 794–8, Sep. 2006. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/16984297>
- [20] J. Blinch, B. D. Cameron, N. J. Hodges, and R. Chua, "Do preparation or control processes result in the modulation to Fitts' law for movements to targets with placeholders?" *Experimental brain research*, vol. 223, no. 4, pp. 505–15, Dec. 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/2311428>
- [21] H. B. Olafsdottir, S. T. Perrault, and Y. Guiard, "Fitts' with a Twist : An Exploration of Scale Effects using a New Experimental Paradigm," *14ème congrès international de l'ACAPS*, 2011.
- [22] S. a. Jax, D. a. Rosenbaum, and J. Vaughan, "Extending Fitts' Law to manual obstacle avoidance." *Experimental brain research*, vol. 180, no. 4, pp. 775–9, Jul. 2007. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17562027>
- [23] J. Accot and S. Zhai, "Beyond Fitts' Law: Models for Trajectory-Based HCI Tasks," in *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '97*. New York, New York, USA: ACM Press, 1997, pp. 295–302. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=258549.258760>
- [24] S. Kulikov, I. S. MacKenzie, and W. Stuerzlinger, "Measuring the effective parameters of steering motions," *CHI '05 extended abstracts on Human factors in computing systems - CHI '05*, p. 1569, 2005. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1056808.1056968>
- [25] L. Liu, J.-B. Martens, and R. van Liere, "Revisiting path steering for 3D manipulation tasks," *2010 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 39–46, Mar. 2010. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5444724>
- [26] L. Liu and R. V. Liere, "The Effect of Varying Path Properties in Path Steering Tasks," *Proceedings of the 16th Eurographics conference on Virtual Environments & Second Joint Virtual Reality (EGVE - JVRC'10)*, pp. 9–16, 2010.
- [27] A. Pavlovych and W. Stuerzlinger, "Target following performance in the presence of latency, jitter, and signal dropouts," *Proceedings of Graphics Interface 2011 (GI '11)*, pp. 33–40, 2011.
- [28] M. R. Mine, "Characterization of end-to-end delays in head-mounted display systems," University of North Carolina at Chapel Hill, Tech. Rep., 1993. [Online]. Available: <http://www0.cs.ucl.ac.uk/teaching/VE/Papers/93-001.pdf>
- [29] O. Pell and V. Averbukh, "Maximum Performance Computing with Dataflow Engines," *Computing in Science & Engineering*, vol. 14, no. 4, pp. 98–103, Jul. 2012. [Online]. Available: http://link.springer.com/chapter/10.1007/978-1-4614-1791-0_25
- [30] S. Friston and A. Steed, "Measuring Latency in Virtual Environments," *IEEE Transactions on Visualization and Computer Graphics (Proceedings Virtual Reality 2014)*, vol. 20, no. 4, 2014.
- [31] International Organization for Standardization, "ISO9241-9: Ergonomic requirements for office work with visual display terminals (VDTs) — Part 9 Requirements for non-keyboard input devices," Geneva, Switzerland, 2000.
- [32] I. S. MacKenzie, "A note on the information-theoretic basis for Fitts' Law," *Journal of Motor Behavior*, vol. 3, no. 21, pp. 323–330, 1989.
- [33] D. Elliott, W. F. Helsen, and R. Chua, "A century later: Woodworth's (1899) two-component model of goal-directed aiming," *Psychological Bulletin*, vol. 127, no. 3, pp. 342–357, 2001. [Online]. Available: <http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-2909.127.3.342>
- [34] D. E. Meyer, R. A. Abrams, S. Kornblum, C. E. Wright, and J. E. Smith, "Optimality in human motor performance: ideal control of rapid aimed movements." *Psychological review*, vol. 95, no. 3, pp. 340–370, 1988.

Sebastian Friston Sebastian Friston, Student Member, IEEE, graduated from the University of Plymouth in 2010 with a BEng in Computer Engineering. He is currently an Engineering Doctoral student at University College London, sponsored by Maxeler Technologies Ltd.



Per Karlström P. Karlstrom has been active in the software & hardware engineering and research field for more than a decade. He completed his Ph.D. and held an Assistant Professor position at Linköping University where he developed an EDA tool for processor design. After this he joined Maxeler to work on high performance FPGA architectures.



Anthony Steed Professor Anthony Steed, Member of IEEE, graduated from the University of Oxford in 1992 with a B.A. in Mathematics & Computation. He was awarded his Ph.D. in 1996 from Queen Mary College in immersive virtual reality systems. Professor Steed is an Associate Editor of IEEE Transactions on Visualisation and Computer Graphics and a member of the IEEE Visualisation & Graphics Technical Committee.

