

Effects of Small Latency Variations in 2D Target Selection Tasks

Andreas Schmid
University of Regensburg
Regensburg, Germany
andreas.schmid@ur.de

Isabell Röhr
University of Regensburg
Regensburg, Germany
isabell.roehr@stud.uni-
regensburg.de

Martina Emmert
University of Regensburg
Regensburg, Germany
martina.emmert@ur.de

Niels Henze
University of Tübingen
Tübingen, Germany
niels.henze@uni-tuebingen.de

Raphael Wimmer
University of Regensburg
Regensburg, Germany
raphael.wimmer@ur.de

Abstract

Systems' latency – the time between user input and system response – slows down the human-computer interaction loop. Several studies revealed negative objective and subjective effects of high latency, typically treating latency as a constant delay. Because latency varies significantly in practice, recent work also assessed the effects of large and sudden latency changes. In practice, however, latency variations are small but frequent. As the effects of such variations are unclear, we investigate how small latency variations (± 50 ms) affect users' performance and perceived task load for 2D target selection tasks with static and moving targets. For static targets, we found that latency variation causes significantly higher completion times and less efficient trajectories, however with small effect sizes. In contrast, we found no significant effects on any performance measure for moving targets. Our findings indicate that the effect of latency variation is generally very small and quickly disappears for non-trivial tasks.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**;
Graphical user interfaces; Pointing.

Keywords

Latency, Pointing, Fitts' Law

ACM Reference Format:

Andreas Schmid, Isabell Röhr, Martina Emmert, Niels Henze, and Raphael Wimmer. 2026. Effects of Small Latency Variations in 2D Target Selection Tasks. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3772318.3791712>

1 Introduction

The interaction between a human and a computer can be modeled as a feedback cycle [11], comprising user input, processing time, and system response. Even small temporal delays within this loop affect the whole interaction process, as they delay the system's

response and therefore the user's next action. Such delays negatively impact user experience [17, 22] and performance [25, 45, 59]. Therefore, reducing and controlling system latency is relevant in many application areas, such as video games [9, 13, 30, 74], audio [7, 39, 40], virtual reality [72, 82], target selection [25, 49, 55], and teleoperation [64, 65, 80]. Plant et al. [61–63] warn that latency could also influence the outcome of user studies (e.g., reaction time experiments) and therefore contribute to the replication crisis [56]. Consequently, it is important to measure latency in interactive systems, understand its effect on users, and find ways to reduce or compensate it [29, 34, 47, 67].

Depending on the task, humans are able to perceive latency as low as 2 milliseconds for dragging objects on a touch screen [58], ~ 20 ms for clicking [42, 43], and ~ 30 ms for audio signals [36, 68]. Negative effects on quantitative measures, such as task time or error rate, can be measured for latencies as low as 25 – 50 ms [42].

A system's end-to-end latency – the time between user input and system response – comprises several partial delays. They originate from individual processing times of input devices [10, 12, 62, 81], software frameworks [71], network communication [37, 53, 78, 85], and displays [75]. Additionally, task scheduling, update rates, and polling intervals lead to variations in latency across multiple interactions [81]. Therefore, end-to-end latency should not be viewed as a constant time delay, but rather a probability distribution of possible latencies for each individual interaction. These latencies do not necessarily have a normal distribution and often are multi-modal due to the interplay of polling rates [71, 81].

Cunningham et al. [19] have shown that users are able to quickly adapt their behavior in environments with high latency to counteract its negative effects. Consequently, users should be able to compensate constant high latency better than varying latency, as the system's behavior is more predictable. However, studies investigating the effects of varying latency are scarce and their findings are inconclusive. Several studies found that large and sudden changes in latency affect users more than a constant high latency [20, 32, 79]. However, latency variations are much smaller in practice, with standard deviations in the order of 10 to 100 ms [38, 71]. A recent study found no effects of such small latency variations on player performance in a first-person shooter [69]. Aforementioned studies have tried to emulate real-world applications, such as operating a graphical user interface, driving a vehicle, or playing a video game. Even though this increases external validity, confounding factors due to the higher task complexity could have influenced the results.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791712>

To isolate the effects of latency and latency variation on users' performance, we conducted a controlled lab study ($n=32$). Participants completed a 2D Fitts' Law task [24], as well as a 2D target selection task with moving targets – each under one of four latency conditions: an average latency of 50 or 150 ms combined with a latency variation of either ± 0 or ± 50 ms. We found small effects of latency variation for the pointing task with static targets. In line with Schmid et al. [69], however, we found no evidence that small latency variations affect task time, error rate, or pointing efficiency for moving targets.

2 Related Work

Negative effects of latency on user experience and performance have been revealed in numerous studies, covering different interaction scenarios and input methods. In this section, we first describe the main sources of latency in interactive systems and its inherent complexity due to being comprised of partial latencies contributed by different components. We then describe how and to which degree latency affects users during interaction, before focusing on research on varying latency in particular.

2.1 Sources of Latency in Interactive Systems

The time from an input event being triggered by a user until a response by a system being displayed is commonly referred to as the system's end-to-end latency [73]. Each component involved contributes to this latency. Input devices' microcontrollers need to read out sensor data and translate it to USB events [81] and the computer's operating system processes those events and passes them on to user space. Applications process those inputs and prepare a response. Depending on context, this can involve time-consuming steps such as disk or database access, network communication, or rendering a complex scene [71]. Lastly, visual output is composed and passed to a display, which in turn has a non-zero response time [21, 75]. Communication between those components includes polling intervals, such as the typical USB polling rate of 125 – 1000 Hz [81], and vertical display synchronization to avoid screen tearing, which adds up to 16.67 ms latency with a 60 Hz monitor [75]. Furthermore, compositing window managers and game engines oftentimes use double buffering, delaying the output by one additional frame [71]. Interplay between such polling rates, as well as process scheduling [78], leads to slight variations in latency, oftentimes with multi-modal distributions [81].

Besides the local latency caused by a system's hardware and software components, network communication can add a substantial amount of delay time, depending on connection quality, communication protocol, and the physical distance between client and server [5, 27, 44]. Due to packet loss, network latency can significantly contribute to variations in latency, with common jitter magnitudes of above 100 ms [37, 41, 78].

2.2 Effect of Latency on Users

As interaction between human and computer forms a feedback loop [11], every latency within this loop delays the whole interaction process. Therefore, latency has a direct influence on task time. In case of time-critical tasks with a temporal deadline, it can also have a negative impact on effectiveness, as deadlines might be

missed, causing the task to fail [15]. Lastly, it has been repeatedly shown that high latency can affect subjective measures, such as user experience [13, 17, 22] or game experience [29, 31, 50].

2.2.1 Latency in Target Selection Tasks. One of the first studies investigating the effect of latency on users' performance was the seminal work by MacKenzie and Ware [55]. Participants completed 1D pointing tasks under different latency conditions, uncovering a linear effect of latency on task time and thereby the task's index of difficulty. Subsequent studies followed a similar scheme: Participants complete a pre-defined task under different latency conditions, and the impact of latency on performance metrics and questionnaire results is analyzed.

Pavlovych et al. [59] found that latency significantly affects performance in terms of throughput, error rate, and movement time in a 2D Fitts' law pointing task under various latency conditions ranging from 33 ms to 133 ms [59]. A similar effect was shown for steering tasks [2] when investigating the impact of latencies between 50 ms and 250 ms [84]. However, even a latency of only 16 ms can have a measurable impact on performance in both pointing and steering tasks [25]. For moving targets, Kundu et al. [45] confirmed the negative effect of latency on performance. A delay of 100 ms already leads to a 50% increase of task difficulty for moving targets, and increases even more when targets are moving faster [49]. Furthermore, not only performance is affected by latency but also subjective measurements such as quality of experience in target selection tasks [13].

2.2.2 Perception Thresholds of Latency. To derive design guidelines for low-latency systems, it is essential to know how much latency can be perceived, from which point it starts to influence users, and what the maximum tolerable amount is. Perception thresholds can be determined using psychophysical measures, such as the *Just Noticeable Difference* (JND) [23]. Such thresholds strongly depend on the task. For example, the JND for latency when tapping on a touch screen has been determined to be at about 20 ms in two different studies [42, 43]. However, when dragging a virtual object over a touch-sensitive surface, delays as low as 2 ms can be perceived due to the spatial offset between finger and object [58].

Even though humans can perceive such small amounts of latency in a controlled environment via direct comparison, an impact on the aforementioned performance and UX metrics can only be measured at much higher latencies. For example, Annett et al. [4] replicated the dragging study by Ng et al. [58], but with a more complex task: Instead of dragging an object and focusing purely on differences in latency, participants were asked to draw shapes and write text. In this context, a much higher perception threshold of 59 ms was found. Similarly, Kaaresoja et al. [43] found that latency starts to impact the perceived quality of a virtual button starting only at 70 ms, regardless of the perception threshold being much lower.

2.2.3 Latency in Video Games. As latency has a major impact on real-time applications, the field of video games research is particularly interested in studying its effects. Especially networked multiplayer games and cloud games streamed from a server are always affected by latency [9, 29, 46, 51]. When it comes to latency studies in the context of video games, researchers oftentimes try to maximize ecological validity, using a particular video game as

the study apparatus, manipulating only the system's latency as an independent variable, and measuring in-game metrics, as well as specialized questionnaires such as the player experience inventory (PXI) [1]. Similar to lab studies with standardized pointing tasks, player performance is deteriorated by latency, leading to less frags, less accurate aiming and worse kill/death ratio in first person shooters [5, 9, 31, 52], longer lap times in racing games [14], as well as less efficient navigation through a virtual world [50]. Furthermore, high latency deteriorates perceived quality of the game and player experience [13, 29, 31, 50].

2.3 Varying Latency

Users can adapt to high-latency environments up to a certain degree [19, 66, 76] by intentionally or subconsciously performing inputs slightly earlier. Accordingly, random variations in latency, also referred to as *jitter* in the context of network delays, should have a greater impact on performance than a constant latency with the same mean. This is because the exact temporal deadline is unknown and therefore predictions of the system's behavior become less precise [79]. However, only few studies investigated the effects of varying latency and findings from those studies are inconclusive.

Weber et al. [79] investigated the effect of large variability of system response time (300 to 3000 ms) in a GUI dialog system. Participants could complete tasks faster with low latency variability, even though average latency was higher in this condition. The authors argue that this effect is caused by temporal expectancy, with users being able to predict the system's behavior more precisely when variability is lower. Similarly, Davis et al. [20] compared the effects of latency and latency variation on mental load and performance in a driving simulator. Average latency was 700 ms in both latency conditions, but was varied between 400 and 1100 ms following a sinusoidal function in the condition with varying latency. While latency variation led to significantly higher lane offset, no other differences compared to the condition with constant latency were found. Halbhuber et al. [32] found that in a 2D shooting game, sudden changes in otherwise constant latency affect performance and user experience more than constant latency with an even higher mean.

All aforementioned studies used sudden latency changes and/or very large ranges of latency variation. Although this choice of conditions is suitable for researching usage scenarios involving breakdowns of network connection or extensive disk access, findings are not applicable to small latency variations caused by local system components or a stable network connection. Only three studies investigated this type of latency variation. Pavlovyh and Stuerzlinger [60] conducted a user study with a path following task, manipulating latency (20 – 170 ms), latency variation (± 0 – ± 60 ms), as well as dropout rate (0 – 20%) and dropout duration (0 – 160 ms). They found that higher latency variation leads to significantly worse performance in following the path for high base latencies. Additionally, they found that the negative effect of latency variation is amplified at higher movement velocities. While significant effects of base latency on participants' response delay could be found, latency variation had no measurable effect here. Schmid et al. [69] systematically investigated the effect of small latency variations on player performance and experience in a

first-person shooter. They deliberately chose latency conditions to represent realistic values that can be caused by local components of a desktop computer. Participants played a first-person shooter with two levels of constant base latency (50 and 150 ms) and two levels of latency variation (± 0 or ± 50 ms). While high base latency had negative effects on hit rate, total amount of kills, kill/death ratio, and player experience, latency variation did not influence those factors. Lastly, Beech et al. [8] compared participants' performance in a 2D pointing task with static targets for a high average latency of 167 ms, and a latency varying by ± 100 ms around the same mean. They found that this latency variation caused increased task times, slower cursor movement, and more overshooting.

2.4 Summary

As latency is an inherent property of human-computer systems, several studies investigated its effect on user experience and performance. Direct consequences of high latency include increased task time and error rate [17, 25, 55], as well as negative effects on user experience [22].

Interplay between different system components [71, 81], as well as network communication [35, 37] leads to variations in latency during interaction. As humans can perceive very low amounts of latency under certain conditions [6, 43, 58], the effects of such latency variations on users have to be studied. However, previous research on such effects is sparse and the results are inconclusive. While the negative effects of large and sudden switches in latency have been confirmed in multiple studies [20, 32, 79], effects for small latency variations in first-person shooters have been ruled out by Schmid et al. [69]. In contrast, there is evidence for effects of small latency variation on simple pointing and steering tasks [8, 59]. Consequently, there is a research gap concerning the complexity of tasks for which small latency variations still play a role.

3 Method

To investigate the impact of varying latency on 2D pointing tasks, we conducted a within-subject study. 32 Participants completed a 2D Fitts' Law pointing task [54] with static targets, as well as a simple 2D pointing task with moving targets. Each task was repeated four times, with two levels of base latency and two levels of latency variation. Compared to the study by Schmid et al. [69], who found no effects of latency variation on any performance measure in a first-person shooter, our tasks were deliberately chosen to be as simple as possible. Thereby, confounding effects are reduced to a minimum and potential effects of latency variation should be found in case they exist. In terms of base latency, we expect to replicate the negative effects on performance found in numerous previous studies [17, 22, 26, 55, 77].

3.1 Study Design

Each participant completed two target selection tasks for each of the four latency conditions. Those conditions are comprised of our two independent variables, BASE latency and latency VARIATION, with two levels each. This results in a within-subjects 2×2 study design. To circumvent sequence effects, the order of latency conditions was varied between participants based on a balanced Latin square.

BASE latency represents the average amount of latency added to the study apparatus, regardless of its variation. *Low* BASE latency refers to an average latency of 50 ms^1 , whereas *high* BASE latency refers to 150 ms. The second IV, latency VARIATION, also comes in two levels: *low* VARIATION represents a constant latency added to the system ($\pm 0 \text{ ms}$). *High* VARIATION means that latency varies by $\pm 50 \text{ ms}$ around the BASE latency. Consequently, during the condition with *high* BASE latency and *high* VARIATION, the system's latency varied between 100 and 200 ms. All aforementioned latency values are added on top of our system's latency, as described in subsection 3.3.4.

Latency was applied to all mouse interactions: clicks and movement events. In case of *high* VARIATION, a new latency value was randomly drawn from a uniform distribution for each mouse event. We selected the same latency conditions as Schmid et al. [69] for easier comparison. Furthermore, the values chosen for BASE latency are in a realistic range [12, 38]. While latency VARIATION of $\pm 50 \text{ ms}$ is higher than in most local desktop contexts [71], such variations are not uncommon for network communication [3, 35, 37, 57].

We measured the dependent variables *task time*, *error rate*, and *distance ratio* – the ratio between the ideal path to a target and the actual pointing trajectory – to quantify participants' performance. Additionally, we applied the raw NASA-TLX questionnaire [33] to measure perceived task load on the six scales *mental demand*, *physical demand*, *temporal demand*, *performance*, *effort*, and *frustration*.

3.2 Tasks

To investigate the impact of varying latency on 2D target selection, we designed two pointing tasks: one with *static* and one with *moving* targets.

3.2.1 Static Targets. For the *static* task, we opted for a 2D Fitts' Law task [54] with nine circular targets arranged in a circle (Fig. 1, left). Participants' objective was to click the currently active target (highlighted by a different color) as quickly as possible. Missing the target resulted in an error being recorded and the next target becoming active.

Targets were either 80 or 160 pixels in diameter (W) and the distance between targets was either 400, 600, or 800 pixels (d). Due to an overlap in index of difficulty (ID) ($W = 80, d = 400$ and $W = 160, d = 800$), we ended up with six different task rounds and five different IDs. Each permutation of target width and diameter was repeated four times, resulting in a total of 24 rounds ($2 \text{ widths} \times 3 \text{ distances} \times 4 \text{ repetitions}$) with 9 targets each. The order of rounds was randomized. The first target of each round served as a set-up target and was therefore excluded from the analysis. Thus, for each participant and latency condition, 192 individual pointing trajectories and clicks were recorded.

3.2.2 Moving Targets. The standardized and well-established 2D Fitts' Law task serves as a reliable proxy for operating static graphical user interfaces. To represent real-time applications, such as video games or teleoperation, we added a second, more complex

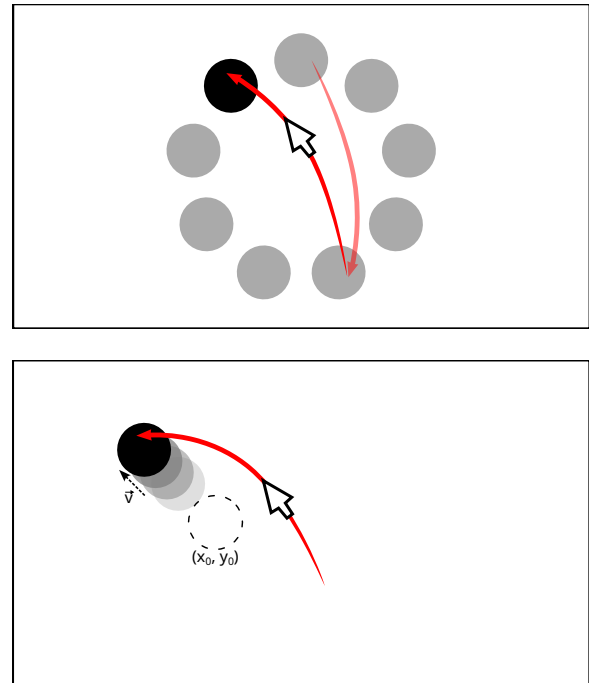


Figure 1: Illustration of tasks with static (top) and moving (bottom) targets.

target selection task with moving targets. Targets were presented one at a time, with a fixed width, starting distance from the cursor, movement direction, and velocity (Fig. 1, right). Participants' objective was to click the target as quickly as possible. If the target was missed or the target moved out of the screen, the trial was recorded as an error.

Similar tasks have been used in previous studies. For example, Hajri et al. [28] and Claypool et al. [13, 16] used simple 2D games with participants chasing circular targets with their mouse cursor. We opted for targets following a linear path with a constant velocity to reduce task complexity. Furthermore, we introduced a temporal deadline [15] of targets leaving the screen to increase potential effects of latency.

Similar to the Fitts' Law task with static targets, target width (80 or 160 px) and initial distance to the current cursor position (400, 600, or 800 px) were varied between trials. Additionally, as movement direction relative to the initial cursor position, as well as target velocity affect task difficulty [28], movement direction was selected from one of *towards*, *diagonally towards*, *perpendicular*, *diagonally away*, or *away*. Movement velocity was set to 200 or $400 \frac{\text{px}}{\text{s}}$. Each permutation of distance, width, velocity, and direction was repeated four times, which results in a total of 240 targets per latency condition ($2 \text{ widths} \times 3 \text{ distances} \times 5 \text{ directions} \times 2 \text{ velocities} \times 4 \text{ repetitions}$). The order of permutations was, again, randomized. Targets always spawned at the vertical center of the screen. If the mouse pointer was too far from this center line, a static setup target was inserted to re-position the cursor at the center.

¹The lowest possible amount of BASE latency has to be at least half the range of the highest VARIATION as otherwise, combinations of *low* BASE and *high* VARIATION would require negative latency.

3.3 Apparatus

In this section, we describe our apparatus, including the study software, technical specification of the system, our method for adding latency to the system, as well as results of latency measurements.

3.3.1 Study Software. We implemented both, the static and moving target selection tasks, in C using the SDL2 graphics framework with an OpenGL backend. We opted for this framework as it was shown to be suitable for low-latency applications [71]. In both conditions, the application was running in full-screen mode, displaying a white background. Targets were rendered in gray, with the active target in the Fitts' Law task being highlighted in black. On click, participants were provided feedback with the active target changing its color to blue in case of a hit, or orange for a miss. To enable left-handed and right-handed mouse usage, targets could be selected by either clicking with the left or the right button of a symmetric mouse.

Cursor position, clicks, information about the current condition, as well as timestamps were logged continuously by storing them in a large array. Log data was only saved to disk at the end of a task to avoid unintended latency spikes.

3.3.2 Adding Latency to the System. We used a software tool similar to Liu's *EvLag* [48] to introduce latency to our system. This program runs independently of the study software and delays all mouse events by a specified amount and variation. Input events are detected and captured using the *evdev* library. Instead of propagating those events to the remaining system, our tool starts a thread that waits for a specified amount of time before invoking an input event using the *uinput* library. It is worth noting that when latency variation is applied, it is possible for input events to be invoked out of order. For movement events, which are triggered in quick succession, this leads to slight spatial jitter of the movement path.

3.3.3 Hardware. We conducted the study on an *HP Pavillon Gaming 790* desktop PC². It ran *Debian Buster 5.10* with proprietary *Nvidia* graphics drivers (version 470.103.01). The external monitor *ASUS ROG Strix XG248Q* had a resolution of 1920 × 1080 pixels and a refresh rate of 240 Hz. We chose the *Logitech G5* gaming mouse, since it features both low latency and low variety in latency ($m = 2.17\text{ ms}$, $SD = 0.3\text{ ms}$, reported by Wimmer et al. [81]).

3.3.4 Latency Measurements. To measure the end-to-end latency of our system, we used Schmid and Wimmer's *Yet Another Latency Measuring Device* (YALMD) [70]. It triggers a click by automatically closing electrical contacts, while a photo sensor attached to the monitor detects a change in brightness as a result of the click. The time elapsed between these events is recorded as the system's end-to-end latency. To use YALMD, we adapted our SDL2 program to display a black rectangle that turns white on click. We displayed this rectangle and attached the photo sensor in the bottom right corner of the screen for measurements.

We conducted a series of 50 measurements each for all four latency conditions, as well as without any added latency. Without added latency, our system has a mean latency of 8.4 ms ($min = 6.7\text{ ms}$, $max = 10.7\text{ ms}$, $SD = 1.2\text{ ms}$). Latency distributions are depicted in Figure 2.

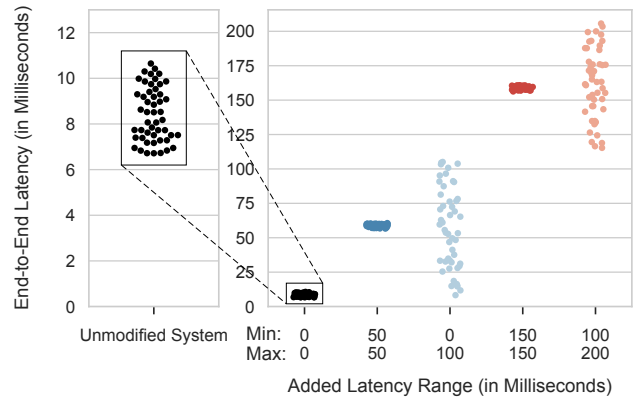


Figure 2: Latency measurements without added latency (6.7 – 10.7 ms, left), as well as all four latency conditions used in the study (right). Latency of 50 or 150 ms was added to user input for conditions with constant latency. Uniformly distributed values between 0 and 100 ms or 100 and 200 ms were added for the conditions with varying latency.

3.4 Procedure

The study was conducted in a controlled laboratory environment. Upon arrival in our lab, we informed participants about our study's procedure, which data we collect, and that they could withdraw from the study at any time. After providing informed consent and demographic data, participants completed a short warm-up phase. Then, the first latency condition was selected from a balanced Latin square and participants completed the Fitts' Law task. Afterward, the latency added to the system was reset to zero and participants answered the NASA-TLX questionnaire. If needed, they were allowed to take a break. Then, they completed the moving target task with the same latency condition, again followed by a NASA-TLX questionnaire and an optional break. This procedure was repeated for all four latency conditions. We concluded the study with a debriefing, informing participants about the study's exact objectives. We estimated a duration of 45–60 minutes to complete the study.

3.5 Participants

Previous work has shown that effects of latency in 2D pointing tasks can be detected reliably even with a small number of participants (e.g., $n=12$ [59, 60]) with most studies using within-subjects designs with about 30 participants [16, 25, 45, 84]. Therefore, we recruited 32 participants (16 men, 15 women, 1 non-binary) for our study via a university mailing list. They were between 19 and 34 years old (mean: 25.8). Among participants, 29 were right-handed, 2 were left-handed and one was ambidextrous. All of them used their right hand to operate the computer mouse. While all participants were experienced with using a mouse, their gaming habits were heterogeneous (action games: 4 playing *often*, 11 *sometimes*, 7 *rarely*, 10 *never*; strategy games: 4 *often*, 5 *sometimes*, 8 *rarely*, 15 *never*).

Participants were rewarded with sweets, as well as credit points for study participation. The study was designed and conducted following the ethical guidelines for research, as provided by our institution.

²Intel i7-8700 (3.2 GHz), Nvidia GTX 1080, 16 GB DDR4 RAM

4 Results

We analyzed the collected data using Python and the *pingouin*³ package. We used Shapiro-Wilk tests to test for normal distribution of residuals for each dependent variable. If data was normally distributed, we used *pingouin*'s repeated measures ANOVA for inferential analysis. Otherwise, we applied an ART ANOVA [83] using the R package *ARTool*⁴. Violation of normality and consequent use of ART ANOVA are stated explicitly in the remainder of the section. In case of violated sphericity, we applied Greenhouse-Geisser correction. We used an alpha level of 0.05 for all statistical tests. Effect sizes (η_p^2) are reported according to Cohen [18]. We structure the results by dependent variable to improve comparability between the static and moving targets tasks.

4.1 Task Time

Task Time for the static target task for different latency conditions are depicted in Fig. 3a. As normal distribution was violated for the static target task ($p < 0.001$), we used an ART ANOVA for inferential analysis. We found a significant effect of BASE latency on *Task Time* ($F(1, 31) = 459.001, p < 0.001, \eta_p^2 = 0.83$). Participants were significantly faster with *low* BASE latency ($M = 0.66\text{ s}, SD = 0.09\text{ s}$) compared to *high* BASE latency ($M = 0.93\text{ s}, SD = 0.13\text{ s}$). We also found a significant effect of latency VARIATION on *Task Time* ($F(1, 31) = 9.277, p = 0.003, \eta_p^2 = 0.09$). For *low* VARIATION, the *Task Time* was significantly lower ($M = 0.77\text{ s}, SD = 0.17\text{ s}$) than for *high* VARIATION ($M = 0.81\text{ s}, SD = 0.18\text{ s}$). We, however, found no significant BASE \times VARIATION interaction effect on *Task Time* ($F(1, 31) = 0.797, p = 0.374, \eta_p^2 < 0.01$).

Task Time for the moving target task for different latency conditions are depicted in Fig. 3b. We found a significant effect of BASE LATENCY on *Task Time* ($F(1, 31) = 753.654, p < 0.001, \eta_p^2 = 0.96$). Again, participants were faster with *low* BASE latency ($M = 0.69\text{ s}, SD = 0.08\text{ s}$) compared to *high* BASE latency ($M = 0.98\text{ s}, SD = 0.12\text{ s}$). We found no significant effect of LATENCY VARIATION on *Task Time* ($F(1, 31) = 2.566, p = 0.119, \eta_p^2 = 0.08$). We also found no significant BASE \times VARIATION interaction effect on *Task Time* ($F(1, 31) = 1.117, p = 0.299, \eta_p^2 = 0.03$).

4.2 Error Rate

We define *Error Rate* as the ratio between clicks missing the target and the total number of clicks. In the static target task, mean error rates were consistent across all conditions, ranging from 2.85% ($SD = 2.09\%$) for *low* BASE to 3.45% ($SD = 2.83\%$) for *high* BASE. Consequently, we found no significant main or interaction effects on *Error Rates* for the static target task (all $p > 0.1$) (Fig. 4a).

We found a highly significant effect of BASE latency on *Error Rate* for moving targets ($F(1, 31) = 159.457, p < 0.001, \eta_p^2 = 0.84$). Significantly fewer errors were made with *low* BASE latency ($M = 13.07\%, SD = 5.23\%$) than with *high* BASE latency ($M = 23.41\%, SD = 7.92\%$). We found no significant effect of VARIATION on *Error Rate* ($F(1, 31) = 2.672, p = 0.112, \eta_p^2 = 0.08$). We also found no significant BASE \times VARIATION interaction effect on *Error Rate* ($F(1, 31) = 0.153, p = 0.698, \eta_p^2 < 0.01$). Results are depicted in Fig. 4b.

³<https://pingouin-stats.org/>

⁴<https://depts.washington.edu/acelab/proj/art/>

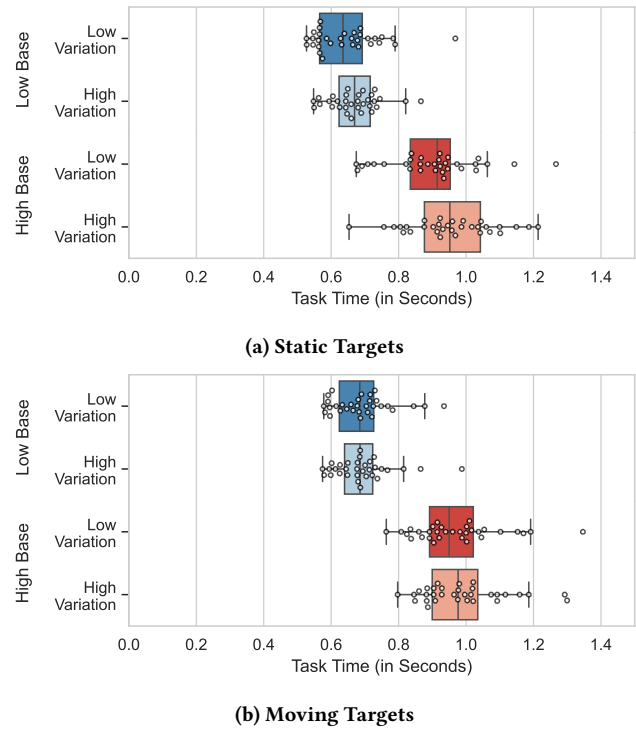


Figure 3: Task time of both tasks for the four latency conditions. Base latency is a major contributor to higher task time. For static targets, we found a small effect of latency variation on task time.

4.3 Mouse Trajectories

We analyzed pointing trajectories by comparing the total length of each trajectory to the length of the ideal path – the Euclidean distance between the initial cursor position and the closest point within the target. By dividing the actual path length by the length of the ideal path, we calculated a *Distance Ratio*. The shortest possible trajectory has a *Distance Ratio* of 1. Higher ratios are caused by excessive mouse movement, for example, due curved trajectories, overshooting, or correcting behavior.

Normal distribution was violated for both tasks (both $p < 0.001$). Therefore, we used an ART ANOVA for inferential analysis.

Distance Ratio for the static target task for all latency conditions are depicted in Fig. 5a. Using an ART ANOVA, we found a significant effect of BASE latency on *Distance Ratio* ($F(1, 31) = 17.003, p < 0.001, \eta_p^2 = 0.15$). Furthermore, we found a significant effect of latency VARIATION on *Distance Ratio* ($F(1, 31) = 42.873, p < 0.001, \eta_p^2 = 0.32$). Participants could follow the ideal path significantly better with *low* VARIATION ($M = 1.23, SD = 0.05$) compared to *high* VARIATION ($M = 1.26, SD = 0.12$). No significant interaction effect of BASE \times VARIATION were found for *Distance Ratio* ($F(1, 31) = 0.115, p = 0.735, \eta_p^2 < 0.01$).

Similar to the task with static targets, we analyzed mouse trajectories for the task with moving targets. To operationalize participants' pointing efficiency, we again investigated the ratio between the total distance covered and a straight line between the initial

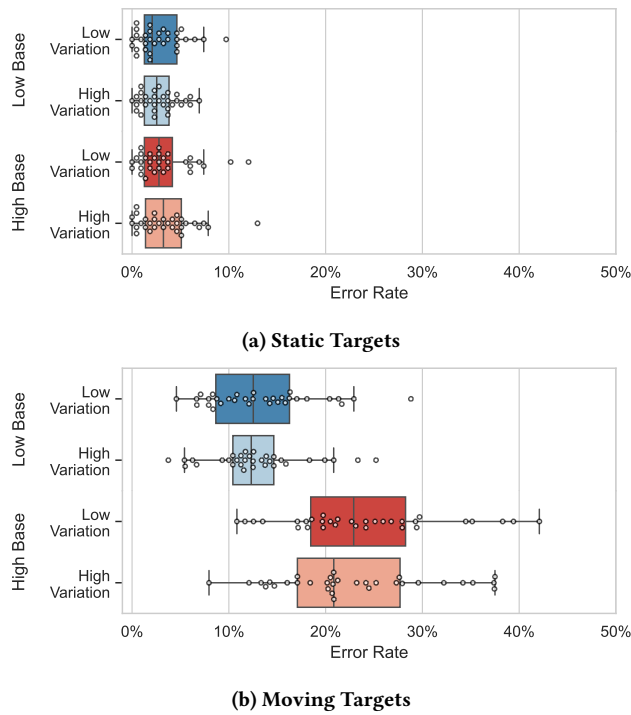


Figure 4: Error rates for both tasks. During the task with static targets, each click that missed the current target was counted as an error. For moving targets, it was additionally possible for the target to leave the screen before it was clicked, which was also counted as an error. Due to this temporal deadline, base latency only influences error rate for moving targets.

position and the target. However, as targets were moving, the ideal path constantly changed. To reduce complexity and keep comparability with the analysis of the task with static targets, we simplified the trajectory analysis using a straight line between the initial and final cursor position as the "ideal trajectory". Consequently, the *Distance Ratio* for moving targets is defined as the ratio between the actual covered distance and the Euclidean distance between the initial and final cursor position.

We found a significant effect of *BASE LATENCY* on *Distance Ratio* ($F(1, 31) = 94.978, p < 0.001, \eta_p^2 = 0.51$). Participants follow the ideal path more closely with *low BASE* latency ($M = 1.23, SD = 0.09$) compared to *high BASE* latency ($M = 1.43, SD = 0.26$). We found no significant effect of *LATENCY VARIATION* on *Distance Ratio* ($F(1, 31) = 0.478, p = 0.491, \eta_p^2 < 0.01$). We also found no significant *BASE* \times *VARIATION* interaction effect on the *Distance Ratio* ($F(1, 31) = 0.358, p = 0.551, \eta_p^2 < 0.01$). Results are shown in Fig. 5b.

4.4 Task Load

For the static target task, we found a significant main effect of *BASE* latency on *Task Load* – the average of the NASA-TLX sub-scales ($F(1, 31) = 39.231, p < 0.001, \eta_p^2 = 0.56$). Participants experienced significantly lower *Task Load* with *low BASE* latency ($M = 30.29, SD =$

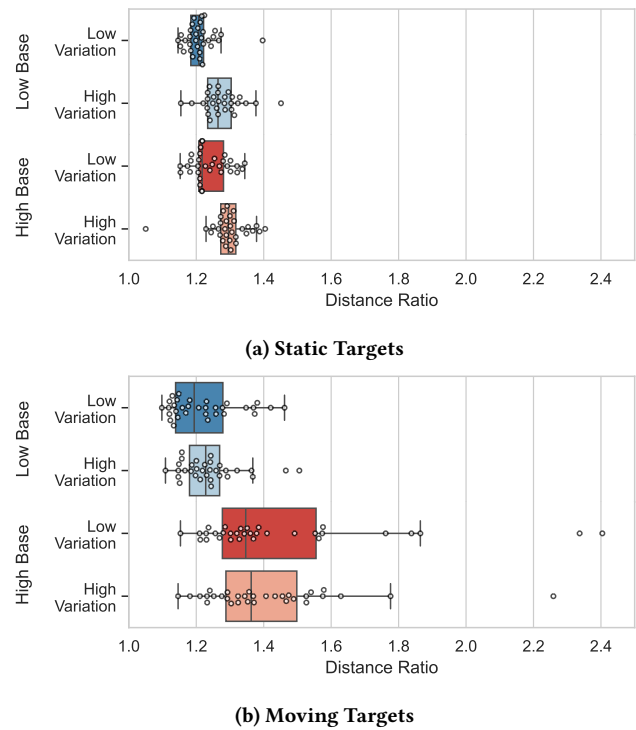


Figure 5: *Distance Ratio* is the length ratio between the the actual pointing trajectory, and the ideal path from the initial cursor position to the target. For static targets, we found significant effects of both, base latency and latency variation, on distance ratio. For moving targets, only high base latency significantly increased distance ratio.

15.40) compared to *high BASE* latency ($M = 40.38, SD = 16.08$). *LATENCY VARIATION* had no significant effect on *Task Load* ($F(1, 31) = 1.002, p = 0.325, \eta_p^2 = 0.03$). Furthermore, we found a significant *BASE* \times *VARIATION* interaction effect on *Task Load* ($F(1, 31) = 4.285, p = 0.047, \eta_p^2 = 0.12$). Bonferroni-Holm corrected pairwise t-tests show that a significant effect of *LATENCY VARIATION* on *Task Load* only exists with *low BASE* latency ($p = 0.03, d_{cohens} = 0.337$). With *low BASE* / *low VARIATION*, *Task Load* ($M = 27.71, SD = 14.22$) was significantly lower compared to *low BASE* / *high VARIATION* ($M = 32.87, SD = 16.31$).

Regarding the individual sub-scales of the NASA-TLX questionnaire, we found a significant main effect of *BASE* latency on every scale except *Temporal Demand*. Additionally, we found significant *BASE* \times *VARIATION* interaction effects for *Performance* and *Frustration*. Using Bonferroni-Holm corrected pairwise t-tests, we could not confirm the interaction effect regarding *Performance* (both $p > 0.1$). Post-hoc tests regarding the interaction effect on *Frustration* show that *LATENCY VARIATION* only has an influence when *BASE* latency is *low* ($p = 0.006, d_{cohens} = 0.65$). With *low BASE* / *low VARIATION*, *Frustration* ($M = 18.91, SD = 14.47$) was significantly lower compared to *low BASE* / *high VARIATION* ($M = 30.78, SD = 21.37$).

Likewise to the static task, we found a significant effect of *BASE* latency on *Task Load* for the moving target task ($F(1, 31) = 31.134,$

Table 1: Results of the inferential analysis for all dependent variables using repeated measures ANOVAs. If the normal distribution of residuals was violated, we applied an ART ANOVA [83]. In this case, the line is marked with a dagger (†).

Performance-Measures	Task	BASE			VARIATION			BASE × VARIATION		
		$F(1, 31)$	p	η_p^2	$F(1, 31)$	p	η_p^2	$F(1, 31)$	p	η_p^2
†Task Time	static	459.001	< 0.001	0.83	9.277	0.003	0.09	0.797	0.374	< 0.01
	moving	753.653	< 0.001	0.96	2.566	0.119	0.08	1.117	0.299	0.03
Error Rate	static	2.628	0.115	0.08	0.154	0.697	< 0.01	0.19	0.666	< 0.01
	moving	159.457	< 0.001	0.84	2.672	0.112	0.08	0.153	0.698	< 0.01
†Distance Ratio	static	17.003	< 0.001	0.15	42.873	< 0.001	0.32	0.115	0.735	< 0.01
†	moving	94.978	< 0.001	0.51	0.478	0.491	< 0.01	0.358	0.551	< 0.01
NASA-TLX Scale										
Mental Demand	static	27.568	< 0.001	0.47	0.519	0.477	0.02	2.13	0.154	0.06
	moving	25.261	< 0.001	0.45	4.568	0.041	0.13	2.464	0.127	0.07
Physical Demand	static	7.539	0.01	0.2	0.169	0.684	< 0.01	0.159	0.693	< 0.01
	moving	15.506	< 0.001	0.33	1.841	0.185	0.06	1.773	0.193	0.05
†Temporal Demand	static	0.066	0.798	< 0.01	2.040	0.157	0.02	0.139	0.710	< 0.01
	moving	1.357	0.253	0.04	0.001	0.971	< 0.01	0.295	0.591	< 0.01
Performance	static	12.338	0.001	0.28	0.047	0.83	< 0.01	4.641	0.039	0.13
	moving	16.089	< 0.001	0.34	0.531	0.472	0.02	1.107	0.301	0.03
Effort	static	19.13	< 0.001	0.38	0.002	0.969	< 0.01	0.854	0.362	0.03
	moving	24.124	< 0.001	0.44	0.099	0.755	< 0.01	1.563	0.221	0.05
Frustration	static	29.703	< 0.001	0.49	1.336	0.257	0.04	8.039	0.008	0.21
	moving	23.691	< 0.001	0.43	0.663	0.422	0.02	2.549	0.121	0.08
Total	static	39.231	< 0.001	0.56	1.002	0.325	0.03	4.285	0.047	0.12
	moving	31.134	< 0.001	0.50	1.021	0.320	0.03	3.710	0.063	0.11

$p < 0.001$, $\eta_p^2 = 0.50$). Participants, again, experienced significantly lower *Task Load* with low BASE latency ($M = 45.48$, $SD = 17.84$) compared to high BASE latency ($M = 57.04$, $SD = 17.55$). We found no significant effect of VARIATION on *Task Load* ($F(1, 31) = 1.021$, $p = 0.320$, $\eta_p^2 = 0.03$). In contrast to the static task, we also found no BASE × VARIATION interaction effect on *Task Load* ($F(1, 31) = 3.710$, $p = 0.063$, $\eta_p^2 = 0.11$).

As for the static target task, we found significant main effects of BASE latency on all NASA-TLX sub-scales except for *Temporal Demand*. Additionally, we found a significant effect of latency VARIATION on *Mental Demand* ($F(1, 31) = 4.568$, $p = 0.041$, $\eta_p^2 = 0.13$). We found no significant BASE × VARIATION interaction effect for any of the sub-scales. Detailed results of the NASA-TLX questionnaire, including all sub-scales, are listed in Table 1 and Figure 6.

5 Discussion

We investigated the effects of latency and latency variation on performance and perceived task load in two different 2D pointing tasks: one with static targets, and one with moving targets. In this section, we discuss our findings on an abstract level, contextualize them with existing literature, and interpret our results' implications for future research. Lastly, we discuss limitations of our study and provide an outlook on potential future work.

5.1 Base Latency

For both tasks, we found significant effects of high base latency on almost all performance and task load measures. The only exceptions

are *Error Rate* for the task with static targets, as well as the NASA-TLX sub-scale *Temporal Demand*. As there was no temporal deadline during the task with static targets, participants could take as long as they wanted to accurately click at the target, regardless of latency condition. Therefore, latency did not affect error rate during this task. Related studies also only found effects of latency on error rate for target selection tasks, if time was limited [59]. The missing effect of latency on temporal demand is in line with this explanation. Participants did not feel more time pressure when latency was added to the system. Overall, our results regarding base latency replicate findings of previous studies, with users being slower [45, 55, 59, 84] and navigating less efficiently [8, 59] when operating a system with high latency.

5.2 Latency Variation

Interestingly, we found diverging results for the two tasks regarding the effect of latency variation. While task time and distance ratio were significantly higher with high latency variation for static targets, we found no such effects for moving targets. However, the effect of latency variation on task time ($\eta_p^2 = 0.09$) is small when compared to the effect of base latency ($\eta_p^2 = 0.83$). We therefore hypothesize that even though small latency variations do have an effect on performance in target selection tasks, their effects are so small that they only surface in very controlled conditions. Similar phenomena have been observed regarding the perception threshold for latency for dragging tasks: Ng et al. [58] determined a just noticeable difference of 2 ms when isolated. However, if dragging interaction is combined with a high-level task such as drawing or

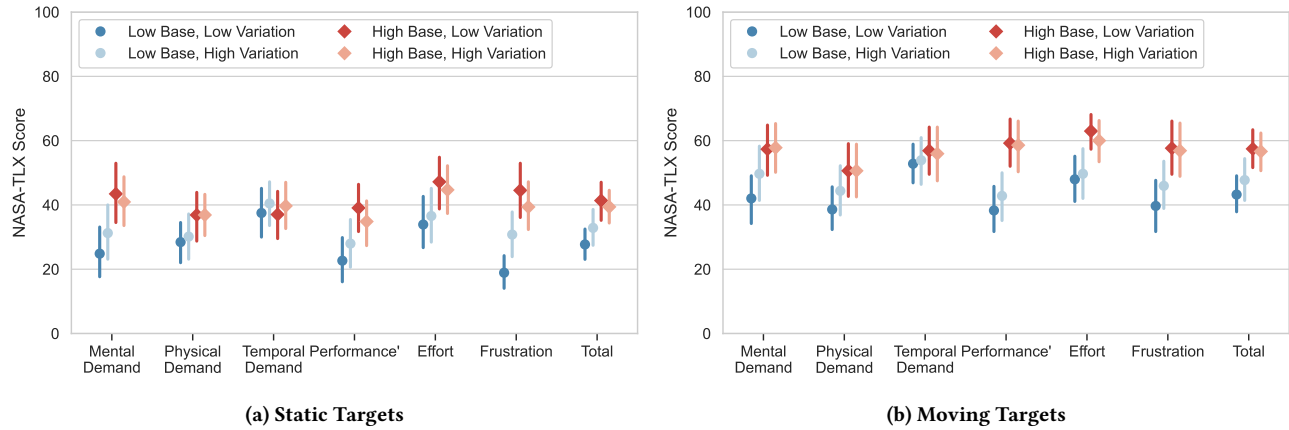


Figure 6: Results of all sub-scales of the NASA-TLX questionnaire for both tasks. We found that high base latency increases perceived task load for all sub-scales except temporal demand, for both tasks. Additionally, we found a significant effect of latency variation on mental demand for the moving target task, as well as interaction effects for performance and frustration for the static target task. Note that the performance score was inverted for better comparability.

writing, latency differences of about 50 ms are required for users to be able to perceive the difference [4].

Besides task time, latency variation had a significant effect on the distance ratio for static targets. Notably, the effect size for latency variation is higher than for the (also significant) influence of base latency. While it is possible that this effect was partially caused by slight spatial jittering due to movement events arriving out of order, this is most certainly not the only factor, as we could not observe a similar effect for the moving target task. A comparable effect of latency variation on longitudinal tracking error was found by Pavlovych and Stuerzlinger [60] for a target following task.

5.3 Interaction between Base Latency and Latency Variation

Furthermore, we found interaction effects on task load between base latency and latency variation for the task with static targets. High latency variation caused perceived frustration and total task load to be significantly higher, but only when base latency was low. This observation is another indicator for the effects of small latency variation being overshadowed by the larger effects of base latency.

5.4 Implications of our Results

We classify our study mainly as foundational research towards a deeper understanding of the influence of latency variation on performance in pointing tasks. However, we can also draw conclusions regarding concrete implications for future HCI research.

As previous studies on the effects of latency variation led to inconclusive results, our findings are an important step towards a deeper understanding of its practical impact. We found a small effect of latency variation on task time and movement trajectories for static targets. This goes in line with findings by Beech et al. [8], who found similar effects, but with twice as much variation. Additionally, similar to Pavlovych and Stuerzlinger [60], we found an effect of latency variation on users' ability to efficiently navigate along a path. We also show that latency variation does not affect

performance in a very simple 2D pointing task with moving targets. We regard this as a generalization of previous findings by Schmid et al. [69], who could show that small variations in latency do not influence player performance in a first-person shooter.

Our findings have a direct impact on future user studies that include simple pointing tasks. As latency is omnipresent in human-computer systems, including study apparatus, and it is never constant due to the interaction of different polling rates throughout the system, it can potentially influence the outcome of experiments [61–63]. On the one hand, just like Schmid et al. [69], our study shows that it is unlikely for small latency variations to have affected the outcome of past experiments. On the other hand, even though we found only few significant effects of latency variation on our dependent variables, and effect sizes were comparatively small, we do not advocate for ignoring it as a factor in the design of interactive systems. Previous studies revealed that the influence of latency is highly dependent on context and can even influence performance if it is not directly perceived [6]. Furthermore, we found a significant effect of latency variation on the NASA-TLX sub-scale *Mental Demand* for the moving target task. This indicates that the task was perceived as subjectively more demanding regardless of there being no measurable effect on performance.

5.5 Limitations and Future Work

In our study, we defined latency variation as a uniform distribution with a comparatively small range of ± 50 milliseconds. Similar designs have been used by Schmid et al. [69] (same latency conditions), Beech et al. [8] (167 ± 100 ms), and Pavlovych and Stuerzlinger [60] ($\pm 0 - \pm 60$ ms). While this replicates typical latency variations caused by a system's software and hardware components, other studies have used different definitions and thus implementations for latency variations. For example, Halbhuber et al. [32] investigated switches between different constant latency levels, and Weber et al. [79] as well as Davis et al. [20] implemented latency variation as

large leaps in settings with high base latency. While both phenomena are relevant in practice, they are inherently different and one should not draw conclusions from one to the other. Consequently, our study did not provide novel insights regarding large and sudden variations in latency, as they can be caused by network dropouts or high system load.

Furthermore, while a latency range of ± 50 milliseconds is realistic when considering a usage context with network communication, it is large when only local system components are involved [71, 81]. As the effects of latency variation are very small, larger ranges could uncover those effects more effectively. For example, Beech et al. [8] found larger effects using twice the latency range. As such large latency variations rarely occur in real-world systems, practical applications of those findings are limited.

We deliberately decided to compromise on the total task duration to cover both tasks and all four latency conditions within a reasonable total study duration. However, our tasks might have been too short for adaptation effects to fully manifest. In comparison, Rohde et al. [66] used a priming phase of 17 minutes before measuring temporal adaption which is longer than a whole latency condition in our study. Thus, future studies should focus on fewer but longer tasks to ensure potential adaptation effects have time to manifest. Additionally, it would be interesting to investigate how users adapt to latency and latency variation over the duration of exposure.

We used a fixed task order (static targets first, moving targets afterwards) for each condition. Therefore, we can not fully exclude potential sequence effects: In theory, participants could have adapted to the new latency condition during the static target task, leading to more homogeneous performance in the moving target task across latency conditions. However, as participants were switching between the static and moving targets tasks four times over the duration of the study, as well as the short duration of individual trials, we regard the pragmatic influence of such an effect on our results as negligible.

Lastly, our participant sample was quite homogeneous, consisting mainly of university students. While it is possible for different results to emerge from other demographics, this does not invalidate the results for our particular sample.

6 Conclusion

We investigated the effect of small latency variations on users' performance and task load in two 2D target selection tasks. We could replicate the effects of base latency on performance measures as found in numerous previous studies. Similar to Beech et al. [8], we found that latency variation can affect performance in a target selection task with static targets. In contrast, we found no effects of latency variation in a slightly more complex task with moving targets. This finding is in line with Schmid et al. [69], who found no effects of latency variation on in-game metrics in a first-person shooter. Therefore, our findings are an important step towards understanding which tasks are influenced by small latency variations.

Acknowledgments

Partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project "Pervasive Touch", grant number 466608952.

References

- [1] Vero Vanden Abeele, Katta Spiel, Lennart Nacke, Daniel Johnson, and Kathrin Gerling. 2020. Development and validation of the player experience inventory: A scale to measure player experiences at the level of functional and psychosocial consequences. *International Journal of Human-Computer Studies* 135 (March 2020), 102370. doi:10.1016/j.ijhcs.2019.102370
- [2] Johnny Accot and Shumin Zhai. 1997. Beyond Fitts' Law: Models for Trajectory-based HCI Tasks. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems (CHI EA '97)*. ACM, Atlanta, Georgia, 250–250. doi:10.1145/1120212.1120376
- [3] Jay Aikat, Jasleen Kaur, F. Donelson Smith, and Kevin Jeffay. 2003. Variability in TCP round-trip times. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement (IMC '03)*. Association for Computing Machinery, New York, NY, USA, 279–284. doi:10.1145/948205.948241
- [4] Michelle Annett, Albert Ng, Paul Dietz, Walter F. Bischof, and Anoop Gupta. 2020. How Low Should We Go? Understanding the Perception of Latency While Inking. In *Graphics Interface 2014* (1 ed.), Paul G. Kry and Andrea Bunt (Eds.). A K Peters/CRC Press, 167–174. doi:10.1201/9781003059325-22
- [5] G. Armitage. 2003. An experimental estimation of latency sensitivity in multiplayer Quake 3. In *The 11th IEEE International Conference on Networks, 2003. ICON2003*. 137–141. doi:10.1109/ICON.2003.1266180
- [6] Christiane Attig, Nadine Rauh, Thomas Franke, and Josef F. Krems. 2017. System Latency Guidelines Then and Now – Is Zero Latency Really Considered Necessary?. In *Engineering Psychology and Cognitive Ergonomics: Cognition and Design (Lecture Notes in Computer Science)*, Don Harris (Ed.). Springer International Publishing, Cham, 3–14. doi:10.1007/978-3-319-58475-1_1
- [7] Christopher Bartlett, Dave Headlam, Mark Bocko, and Gordana Velickic. 2006. Effect of Network Latency on Interactive Musical Performance. *Music Perception* 24, 1 (Sept. 2006), 49–62. doi:10.1525/mp.2006.24.1.49
- [8] Sam Beech, Danaë Stanton Fraser, and Iain D. Gilchrist. 2025. Visuomotor adaptation to constant and varying delays in a target acquisition task. *Journal of Vision* 25, 6 (May 2025), 8. doi:10.1167/jov.25.6.8
- [9] Tom Beigbeder, Rory Coughlan, Corey Lusher, John Plunkett, Emmanuel Agu, and Mark Claypool. 2004. The effects of loss and latency on user performance in unreal tournament 2003®. In *Proceedings of 3rd ACM SIGCOMM workshop on Network and system support for games (NetGames '04)*. Association for Computing Machinery, New York, NY, USA, 144–151. doi:10.1145/1016540.1016556
- [10] Florian Bockes, Raphael Wimmer, and Andreas Schmid. 2018. LagBox – Measuring the Latency of USB-Connected Input Devices. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–6. doi:10.1145/3170427.3188632
- [11] Stuart K. Card, Thomas P. Moran, and Allen Newell. 1986. *The model human processor: An engineering model of human performance*. John Wiley & Sons, Oxford, England.
- [12] Géry Casiez, Thomas Pietrzak, Damien Marchal, Sébastien Poulmane, Matthieu Falce, and Nicolas Roussel. 2017. Characterizing Latency in Touch and Button-Equipped Interactive Systems. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*. ACM, New York, NY, USA, 29–39. doi:10.1145/3126594.3126606
- [13] Mark Claypool. 2018. Game Input with Delay—Moving Target Selection with a Game Controller Thumbstick. *ACM Transactions on Multimedia Computing, Communications, and Applications* 14, 3s (June 2018), 1–22. doi:10.1145/3187288
- [14] Mark Claypool and Kajal Claypool. 2006. Latency and player actions in online games. *Commun. ACM* 49, 11 (Nov. 2006), 40–45. doi:10.1145/1167838.1167860
- [15] Mark Claypool and Kajal Claypool. 2010. Latency can kill: precision and deadline in online games. In *Proceedings of the first annual ACM SIGMM conference on Multimedia systems*. ACM, Phoenix Arizona USA, 215–222. doi:10.1145/1730836.1730863
- [16] Mark Claypool, Ragnhild Eg, and Kjetil Raaen. 2016. The Effects of Delay on Game Actions: Moving Target Selection with a Mouse. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts - CHI PLAY Companion '16*. ACM Press, Austin, Texas, USA, 117–123. doi:10.1145/2968120.2987743
- [17] Mark Claypool, Ragnhild Eg, and Kjetil Raaen. 2017. Modeling User Performance for Moving Target Selection with a Delayed Mouse. In *MultiMedia Modeling, Laurent Amsaleg, Gylfi Þór Guðmundsson, Cathal Gurrin, Björn Þór Jónsson, and Shin'ichi Satoh* (Eds.). Springer International Publishing, Cham, 226–237. doi:10.1007/978-3-319-51811-4_19
- [18] Jacob Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences* (2 ed.). Routledge, New York. doi:10.4324/9780203771587
- [19] Douglas W. Cunningham, Vincent A. Billock, and Brian H. Tsou. 2001. Sensorimotor Adaptation to Violations of Temporal Contiguity. *Psychological Science* 12, 6 (Nov. 2001), 532–535. doi:10.1111/1467-9280.d01-17
- [20] James Davis, Christopher Smyth, and Kaleb McDowell. 2010. The Effects of Time Lag on Driving Performance and a Possible Mitigation. *IEEE Transactions on Robotics* 26, 3 (June 2010), 590–593. doi:10.1109/TRO.2010.2046695

- [21] Federico Dossena and Andrea Trentini. 2022. OpenLDAT—A system for the measurement of display latency metrics. *Journal of the Society for Information Display* n/a, n/a (Feb. 2022). doi:10.1002/jSID.1104
- [22] Ragnhild Eg, Kjetil Raaen, and Mark Claypool. 2018. Playing with Delay: With Poor Timing Comes Poor Performance, and Experience Follows Suit. In *2018 Tenth International Conference on Quality of Multimedia Experience (QoMEX)*. 1–6. doi:10.1109/QoMEX.2018.8463382
- [23] Gustav Theodor Fechner. 1860. *Elemente der Psychophysik*. Vol. 2. Breitkopf und Härtel.
- [24] Paul M. Fitts. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* 47, 6 (1954), 381–391. doi:10.1037/h0055392
- [25] Sebastian Friston, Per Karlström, and Anthony Steed. 2016. The Effects of Low Latency on Pointing and Steering Tasks. *IEEE Transactions on Visualization and Computer Graphics* 22, 5 (May 2016), 1605–1615. doi:10.1109/TVCG.2015.2446467
- [26] Sebastian Friston and Anthony Steed. 2014. Measuring Latency in Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics* 20, 4 (April 2014), 616–625. doi:10.1109/TVCG.2014.30
- [27] Rohitha Goonatilake and Rafic Bachnak. 2012. Modeling Latency in a Network Distribution. *Network and Communication Technologies* 1, 2 (Aug. 2012), p1. doi:10.5539/nct.v1n2p1
- [28] Abir Al Hajri, Sidney Fels, Gregor Miller, and Michael Ilich. 2011. Moving Target Selection in 2D Graphical User Interfaces. In *Human-Computer Interaction – INTERACT 2011*, Pedro Campos, Nicholas Graham, Joaquim Jorge, Nuno Nunes, Philippe Palanque, and Marco Winckler (Eds.). Springer, Berlin, Heidelberg, 141–161. doi:10.1007/978-3-642-23771-3_12
- [29] David Halbhüser, Niels Henze, and Valentin Schwind. 2021. Increasing Player Performance and Game Experience in High Latency Systems. *Proceedings of the ACM on Human-Computer Interaction* 5, CHI PLAY (Oct. 2021), 1–20. doi:10.1145/3474710
- [30] David Halbhüser, Annika Köhler, Markus Schmidbauer, Jannik Wiese, and Niels Henze. 2022. The Effects of Auditory Latency on Experienced First-Person Shooter Players. In *Proceedings of Mensch und Computer 2022 (MuC '22)*, Association for Computing Machinery, New York, NY, USA, 286–296. doi:10.1145/3543758.3543760
- [31] David Halbhüser, Philipp Schaubhuber, Valentin Schwind, and Niels Henze. 2023. The Effects of Latency and In-Game Perspective on Player Performance and Game Experience. *Proceedings of the ACM on Human-Computer Interaction* 7, CHI PLAY (Oct. 2023), 424:1308–424:1329. doi:10.1145/3611070
- [32] David Halbhüser, Valentin Schwind, and Niels Henze. 2022. Don't Break my Flow: Effects of Switching Latency in Shooting Video Games. *Proceedings of the ACM on Human-Computer Interaction* 6, CHI PLAY (2022), 229:1–229:20. doi:10.1145/3549492
- [33] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*. Vol. 52. Elsevier, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [34] Niels Henze, Markus Funk, and Alireza Sahami Shirazi. 2016. Software-reduced touchscreen latency. ACM Press, 434–441. doi:10.1145/2935334.2935381
- [35] F. Hernandez-Campos and M. Papadopouli. 2005. Assessing the real impact of 802.11 WLANs: a large-scale comparison of wired and wireless traffic. In *2005 14th IEEE Workshop on Local & Metropolitan Area Networks*. 6 pp.–6. doi:10.1109/LANMAN.2005.1541511
- [36] Perttu Hämäläinen and Teemu Mäki-Patola. 2004. Latency Tolerance for Gesture Controlled Continuous Sound Instrument without Tactile Feedback. In *Proceedings of the International Computer Music Conference (ICMC2004)*. The International Computer Music Association, Coral Gables, Florida, USA, 8.
- [37] Toke Høiland-Jørgensen, Bengt Ahlgren, Per Hurtig, and Anna Brunstrom. 2016. Measuring Latency Variation in the Internet. In *Proceedings of the 12th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '16)*, Association for Computing Machinery, New York, NY, USA, 473–480. doi:10.1145/2999572.2999603
- [38] Zenja Ivkovic, Ian Stavness, Carl Gutwin, and Steven Sutcliffe. 2015. Quantifying and Mitigating the Negative Effects of Local Latencies on Aiming in 3D Shooter Games. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, Association for Computing Machinery, New York, NY, USA, 135–144. doi:10.1145/2702123.2702432
- [39] Robert H. Jack, Adib Mehrabi, Tony Stockman, and Andrew McPherson. 2018. Action-sound Latency and the Perceived Quality of Digital Musical Instruments - Comparing Professional Percussionists and Amateur Musicians. *Music Perception* 36, 1 (Sept. 2018), 109–128. doi:10.1525/mp.2018.36.1.109
- [40] Robert H. Jack, Tony Stockman, and Andrew McPherson. 2016. Effect of latency on performer interaction and subjective quality assessment of a digital musical instrument. In *Proceedings of the Audio Mostly 2016*. ACM, Norrköping Sweden, 116–123. doi:10.1145/2986416.2986428
- [41] S. Jaiswal, G. Iannaccone, C. Diot, J. Kurose, and D. Towsley. 2004. Inferring TCP connection characteristics through passive measurements. In *IEEE INFOCOM 2004*, Vol. 3. 1582–1592 vol.3. doi:10.1109/INFCOM.2004.1354571
- [42] Ricardo Jota, Albert Ng, Paul Dietz, and Daniel Wigdor. 2013. How fast is fast enough? a study of the effects of latency in direct-touch pointing tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, Association for Computing Machinery, New York, NY, USA, 2291–2300. doi:10.1145/2470654.2481317
- [43] Topi Kaaresoja, Stephen Brewster, and Vuokko Lantz. 2014. Towards the Temporally Perfect Virtual Button: Touch-Feedback Simultaneity and Perceived Quality in Mobile Touchscreen Press Interactions. *ACM Transactions on Applied Perception* 11, 2 (July 2014), 1–25. doi:10.1145/2611387
- [44] Sebastian Kaune, Konstantin Pusset, Christof Leng, Aleksandra Kovacevic, Gareth Tyson, and Ralf Steinmetz. 2009. Modelling the Internet Delay Space Based on Geographical Locations. In *2009 17th Euromicro International Conference on Parallel, Distributed and Network-based Processing*. 301–310. doi:10.1109/PDP.2009.44
- [45] Aritra Kundu, Ragnhild Eg, and Mark Claypool. 2024. The Interplay Between Delay and Device in a Moving Target Selection Game. In *2024 IEEE Conference on Games (CoG)*. 1–8. doi:10.1109/CoG60054.2024.10645564 ISSN: 2325-4289.
- [46] Ulrich Lampe, Qiong Wu, Hans Ronny, André Miede, and Ralf Steinmetz. 2013. To Frag or to Be Fraggled - An Empirical Assessment of Latency in Cloud Gaming. SciTePress - Science and Technology Publications, 5–12. doi:10.5220/0004345900050012
- [47] Huy Viet Le, Valentin Schwind, Philipp Göttlich, and Niels Henze. 2017. PredictTouch: A System to Reduce Touchscreen Latency using Neural Networks and Inertial Measurement Units. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces (ISS '17)*, Association for Computing Machinery, New York, NY, USA, 230–239. doi:10.1145/3132272.3134138
- [48] Shengmei Liu and Mark Claypool. 2021. EvLag - A Tool for Monitoring and Lagging Linux Input Devices. (2021), 6.
- [49] Shengmei Liu and Mark Claypool. 2021. Game Input with Delay - A Model of the Time Distribution for Selecting a Moving Target with a Mouse. In *MultiMedia Modeling (Lecture Notes in Computer Science)*, Jakub Lokoč, Tomáš Skopal, Klaus Schoeffmann, Vasileios Mezaris, Xirong Li, Stefanos Vrochidis, and Ioannis Patras (Eds.). Springer International Publishing, Cham, 506–518. doi:10.1007/978-3-030-67832-6_41
- [50] Shengmei Liu and Mark Claypool. 2022. The Impact of Latency on Navigation in a First-Person Perspective Game. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*, Association for Computing Machinery, New York, NY, USA, 1–11. doi:10.1145/3491102.3517660
- [51] Shengmei Liu, Mark Claypool, Atsuo Kuwahara, James Scovell, and Jamie Sherman. 2021. L33t or N00b? How Player Skill Alters the Effects of Network Latency on First Person Shooter Game Players. In *Proceedings of the Workshop on Game Systems (GameSys '21) (GameSys '21)*, Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3458335.3460811
- [52] Shengmei Liu, Mark Claypool, Atsuo Kuwahara, Jamie Sherman, and James J Scovell. 2021. Lower is Better? The Effects of Local Latencies on Competitive First Person Shooter Game Players. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–12. doi:10.1145/3411764.3445245
- [53] Shengmei Liu, Atsuo Kuwahara, James Scovell, Jamie Sherman, and Mark Claypool. 2021. Comparing the Effects of Network Latency versus Local Latency on Competitive First Person Shooter Game Players. doi:10.31219/osf.io/bhxcy
- [54] I. Scott MacKenzie and William Buxton. 1992. Extending Fitts' law to two-dimensional tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*, Association for Computing Machinery, New York, NY, USA, 219–226. doi:10.1145/142750.142794
- [55] I. Scott MacKenzie and Colin Ware. 1993. Lag as a determinant of human performance in interactive systems. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems (CHI '93)*, Association for Computing Machinery, New York, NY, USA, 488–493. doi:10.1145/169059.169431
- [56] Scott E. Maxwell, Michael Y. Lau, and George S. Howard. 2015. Is psychology suffering from a replication crisis? What does "failure to replicate" really mean? *The American Psychologist* 70, 6 (Sept. 2015), 487–498. doi:10.1037/a0039400
- [57] Nabil Mesbahi and Hamza Dahmouni. 2016. Delay and jitter analysis in LTE networks. In *2016 International Conference on Wireless Networks and Mobile Communications (WINCOM)*. 122–126. doi:10.1109/WINCOM.2016.7777202
- [58] Albert Ng, Michelle Annett, Paul Dietz, Anoop Gupta, and Walter F. Bischof. 2014. In the Blink of an Eye: Investigating Latency Perception During Stylus Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, ACM, Toronto, Ontario, Canada, 1103–1112. doi:10.1145/2556288.2557037
- [59] Andriy Pavlovych and Wolfgang Stuerzlinger. 2009. The tradeoff between spatial jitter and latency in pointing tasks. ACM Press, 187. doi:10.1145/1570433.1570469
- [60] Andriy Pavlovych and Wolfgang Stuerzlinger. 2011. Target Following Performance in the Presence of Latency, Jitter, and Signal Dropouts. In *Proceedings of Graphics Interface 2011 (GI '11)*, Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 33–40. doi:10.5555/1992917.1992924
- [61] Richard R. Plant. 2016. A reminder on millisecond timing accuracy and potential replication failure in computer-based psychology experiments: An open letter.

- Behavior Research Methods* 48, 1 (March 2016), 408–411. doi:10.3758/s13428-015-0577-0
- [62] Richard R. Plant, Nick Hammond, and Tom Whitehouse. 2003. How choice of mouse may affect response timing in psychological studies. *Behavior Research Methods, Instruments, & Computers* 35, 2 (May 2003), 276–284. doi:10.3758/BF03202553
- [63] Richard R. Plant and Philip T. Quinlan. 2013. Could millisecond timing errors in commonly used equipment be a cause of replication failure in some neuroscience studies? *Cognitive, Affective, & Behavioral Neuroscience* 13, 3 (Sept. 2013), 598–614. doi:10.3758/s13415-013-0166-6
- [64] Markus Rank, Zhuanghua Shi, and Sandra Hirche. 2010. Perception of Delay in Haptic Telepresence Systems. *Presence: Teleoperators and Virtual Environments* 19, 5 (Oct. 2010), 389–399. doi:10.1162/pres_a_00021
- [65] Reiza Rayman, Serguei Primak, Rajni Patel, Merhadd Moallem, Roya Morady, Mahdi Tavakoli, Vanja Subotic, Natalie Galbraith, Aimee van Wynsberghe, and Kris Croome. 2005. Effects of Latency on Telesurgery: An Experimental Study. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y. Vardi, Gerhard Weikum, James S. Duncan, and Guido Gerig (Eds.). Vol. 3750. Springer Berlin Heidelberg, Berlin, Heidelberg, 57–64. http://link.springer.com/10.1007/11566489_8
- [66] Marieke Rohde, Loes C. J. van Dam, and Marc O. Ernst. 2014. Predictability is necessary for closed-loop visual feedback delay adaptation. *Journal of Vision* 14, 3 (March 2014), 4. doi:10.1167/14.3.4
- [67] Saeed Shafiee Sabet, Steven Schmidt, Saman Zadtootaghaj, Babak Naderi, Carsten Griwodz, and Sebastian Möller. 2020. A latency compensation technique based on game characteristics to mitigate the influence of delay on cloud gaming quality of experience. In *Proceedings of the 11th ACM Multimedia Systems Conference (MMSys '20)*. Association for Computing Machinery, New York, NY, USA, 15–25. doi:10.1145/3339825.3391855
- [68] Andreas Schmid, Maria Ambros, Johanna Bogon, and Raphael Wimmer. 2024. Measuring the Just Noticeable Difference for Audio Latency. In *Proceedings of the 19th International Audio Mostly Conference: Explorations in Sonic Cultures (AM '24)*. Association for Computing Machinery, New York, NY, USA, 325–331. doi:10.1145/3678299.3678331
- [69] Andreas Schmid, David Halbhuber, Thomas Fischer, Raphael Wimmer, and Niels Henze. 2023. Small Latency Variations Do Not Affect Player Performance in First-Person Shooters. *Proceedings of the ACM on Human-Computer Interaction* 7, CHI PLAY (Oct. 2023), 381:197–381:216. doi:10.1145/3611027
- [70] Andreas Schmid and Raphael Wimmer. 2021. *Yet Another Latency Measuring Device*. Technical Report. OSF Preprints. <https://osf.io/tkgjh/>
- [71] Andreas Schmid and Raphael Wimmer. 2023. Measuring the Latency of Graphics Frameworks on X11-Based Systems. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3544549.3585779
- [72] Valentin Schwind, David Halbhuber, Jakob Fehle, Jonathan Sasse, Andreas Pfaffelhuber, Christoph Tögel, Julian Dietz, and Niels Henze. 2020. The Effects of Full-Body Avatar Movement Predictions in Virtual Reality using Neural Networks. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology (VRST '20)*. Association for Computing Machinery, New York, NY, USA, 1–11. doi:10.1145/3385956.3418941
- [73] Ben Shneiderman. 1984. Response time and display rate in human performance with computers. *ACM Comput. Surv.* 16, 3 (Sept. 1984), 265–285. doi:10.1145/2514.2517
- [74] Josef Spjut, Ben Boudaoud, Kamran Binaee, Jonghyun Kim, Alexander Majercik, Morgan McGuire, David Luebke, and Joohwan Kim. 2019. Latency of 30 ms Benefits First Person Targeting Tasks More Than Refresh Rate Above 60 Hz. In *SIGGRAPH Asia 2019 Technical Briefs on - SA '19 (SA '19)*. ACM Press, New York, NY, USA, 110–113. doi:10.1145/3355088.3365170
- [75] Patrick Stadler, Andreas Schmid, and Raphael Wimmer. 2020. DispLagBox: simple and replicable high-precision measurements of display latency. In *Proceedings of the Conference on Mensch und Computer*. ACM, Magdeburg Germany, 105–108. doi:10.1145/3404983.3410015
- [76] Chess Stetson, Xu Cui, P. Read Montague, and David M. Eagleman. 2006. Motor-sensory recalibration leads to an illusory reversal of action and sensation. *Neuron* 51, 5 (Sept. 2006), 651–659. doi:10.1016/j.neuron.2006.08.006
- [77] Robert J. Teather, Andriy Pavlovych, Wolfgang Stuerzlinger, and I. Scott MacKenzie. 2009. Effects of tracking technology, latency, and spatial jitter on object movement. *IEEE*, 43–50. doi:10.1109/3DUI.2009.4811204
- [78] Robert Underwood, Jason Anderson, and Amy Apon. 2018. Measuring Network Latency Variation Impacts to High Performance Computing Application Performance. In *Proceedings of the 2018 ACM/SPEC International Conference on Performance Engineering (ICPE '18)*. Association for Computing Machinery, New York, NY, USA, 68–79. doi:10.1145/3184407.3184427
- [79] Florian Weber, Carola Haering, and Roland Thomaschke. 2013. Improving the human-computer dialogue with increased temporal predictability. *Human Factors* 55, 5 (Oct. 2013), 881–892. doi:10.1177/0018720813475812
- [80] Daniel E. Whitney. 1987. Historical Perspective and State of the Art in Robot Force Control. *The International Journal of Robotics Research* 6, 1 (March 1987), 3–14. doi:10.1177/027836498700600101 Publisher: SAGE Publications Ltd STM.
- [81] Raphael Wimmer, Andreas Schmid, and Florian Bockes. 2019. On the Latency of USB-Connected Input Devices. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300650
- [82] Paul Winkler, Philipp Stiens, Nadine Rauh, Thomas Franke, and Josef Krems. 2020. How latency, action modality and display modality influence the sense of agency: a virtual reality study. *Virtual Reality* 24, 3 (Sept. 2020), 411–422. doi:10.1007/s10055-019-00403-y
- [83] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 143–146. doi:10.1145/1978942.1978963
- [84] Shota Yamanaka and Wolfgang Stuerzlinger. 2024. The Effect of Latency on Movement Time in Path-steering. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–19. doi:10.1145/3613904.3642316
- [85] Li Zheng, Liren Zhang, and Dong Xu. 2001. Characteristics of network delay and delay jitter and its effect on voice over IP (VoIP). In *ICC 2001. IEEE International Conference on Communications. Conference Record (Cat. No.01CH37240)*, Vol. 1. 122–126 vol.1. doi:10.1109/ICC.2001.936286