REACH+: Extending the Reachability of Encountered-type Haptics Devices through Dynamic Redirection in VR

Eric J. Gonzalez
Stanford University
Stanford, USA
ejgonz@stanford.edu

Parastoo Abtahi
Stanford University
Stanford, USA
parastoo@stanford.edu

Sean Follmer
Stanford University
Stanford, USA
sfollmer@stanford.edu

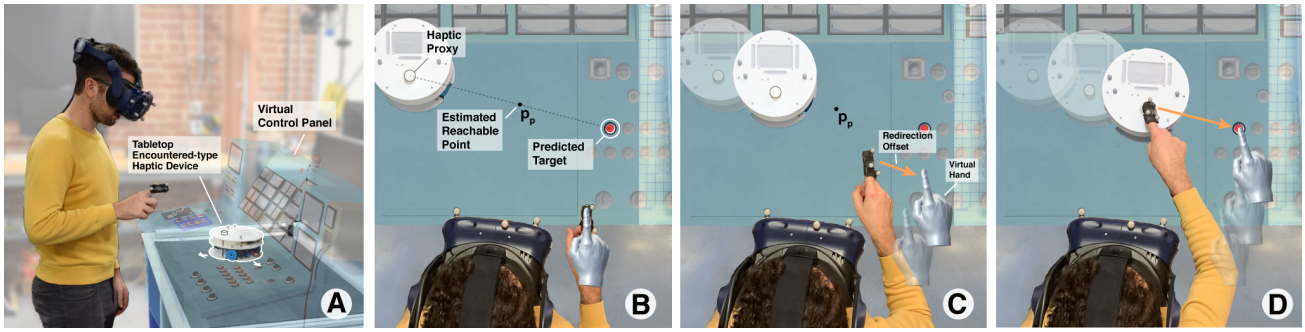

Figure 1. (A) A user in VR interacts with a virtual control panel rendered by an encountered-type haptic device. Our framework: (B) predicts their intended target and arrival time, estimates a point p_p reachable by the device, and then (C) redirects the user's hand to the reachable point to minimize visuo-haptic discrepancy and ensure (D) haptic feedback is rendered.

ABSTRACT

Encountered-type haptic devices (EHDs) face a number of challenges when physically embodying content in a virtual environment, including workspace limits and device latency. To address these issues, we propose REACH+, a framework for dynamic visuo-haptic redirection to improve the perceived performance of EHDs during physical interaction in VR. Using this approach, we estimate the user's arrival time to their intended target and redirect their hand to a point within the EHD's spatio-temporally reachable space. We present an evaluation of this framework implemented with a desktop mobile robot in a 2D target selection task, tested at four robot speeds (20, 25, 30 and 35 cm/s). Results suggest that REACH+ can improve the performance of lower-speed EHDs, increasing their rate of on-time arrival to the point of contact by up to 25% and improving users' self-reported sense of realism.

Author Keywords

Encountered-type Haptics; Virtual Reality; Visuo-haptic illusion; Redirection; Retargeting

CCS Concepts

•**Human-centered computing** → **Human computer interaction (HCI);** *Haptic devices;* User studies;

UIST '20, October 20–23, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7514-6/20/10 ...$15.00.
http://dx.doi.org/10.1145/3379337.3415870

INTRODUCTION

Supporting immersive, general-purpose force feedback for VR is a difficult challenge. Traditional haptic devices such as the PHANToM [35] can render force accurately, but often only at a single point in a small workspace, while wearable devices can be cumbersome and have difficulty rendering externally grounded forces [26]. An ideal haptic system for VR should support a diverse array of physical interactions and manipulations throughout the user's workspace, through physical feedback that closely matches that expected in reality.

One promising approach to addressing this problem is the paradigm of *encountered-type haptics* [36]. Rather than requiring the user to hold or wear a device, encountered-type haptic devices (EHDs) autonomously position haptic proxies (i.e., physical stand-ins for virtual objects) in the real world. This allows users to reach out and "encounter" virtual objects physically just like in the real world. While the type of haptic feedback provided by EHDs can vary greatly (e.g., shape [20], texture [2, 3]), the use of such haptic proxies for virtual content has generally been shown to benefit user's sense of immersion and engagement in VR [23, 43].

In practice, however, EHDs face a number of challenges to supporting real-time physical interaction. We focus on one challenge critical to interaction with an EHD: *reachability*. We define the reachability of an EHD as its ability to appropriately position itself at a target, both in space and time. If the EHD is unable to reach the target, this can lead to spatial discrepancies (i.e., position mismatch between the virtual object and haptic proxy) and latency between what the user sees and feels, both

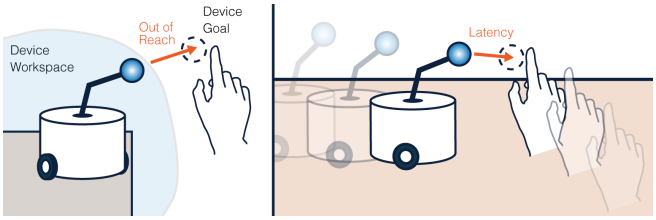

Figure 2. Common sources of spatial discrepancy. Left: The goal position (dashed) for the haptic proxy (blue) is outside the workspace. Right: The device is not fast enough to arrive at the goal in time for contact.

of which negatively impact user experience [34, 12]. Two common device limitations resulting in reachability issues are:

1. **Workspace Constraints**: Kinematic constraints limit the workspace of the EHD, making certain targets unreachable.
2. **Speed Limitations**: Actuator limitations and safety considerations may limit EHD speed, delaying the device's arrival to certain targets.

While these issues can be addressed in part by improving device hardware, factors such as cost, safety, and complexity often lead to design decisions that make device workspace and speed constraints unavoidable.

Recently, however, researchers have begun exploring the use of visuo-haptic illusions to augment physical interactions in VR [1, 4]. Reach redirection [29] is one such illusion that augments the mapping between the user's physical hand and their virtual hand. This can be used to "redirect" one's reaching motion to particular point in space. While most existing research has focused on interactions with passive physical props, such as mapping multiple virtual objects to a single real object [4, 10], some recent work has explored the use of illusions to compensate for limitations of EHDs, include altering the perceived resolution of shape-changing displays [1] and compensating for known spatial discrepancies [33].

We propose an approach to compensate for both (1) known spatial discrepancies caused by device workspace limitations and (2) dynamic spatial discrepancies caused by device speed limitations. We refer to these as the reachability limitations of an EHD; namely, they impact whether the EHD can reach the user's intended target in time for contact. We present REACH+: a framework leveraging dynamic reach redirection to compensate for reachability limitations of EHDs on the fly. We first present our general framework, which models basic user interaction, predicts user intent, estimates time until contact, and pre-emptively compensates for expected spatial discrepancies using redirection, dynamically. We then describe and validate our implementation of a desktop-scale EHD using an omnidirectional mobile robot. Finally, we evaluate the effectiveness of REACH+ at compensating for device latency across multiple speeds of the EHD.

RELATED WORK

Encountered-Type Haptics

To provide haptic feedback in VR, McNeely proposed that users can interact with robot-actuated physical proxies that position themselves for interaction on an as-needed basis [36]. He referred to this concept as *robotic graphics*, drawing an analogy between graphical displays that simulate objects visually and robotic displays that can simulate their feel. Such robotic displays have since been termed *encountered-type haptic devices* (EHDs), as users "encounter" the device when they make contact with virtual objects, without the need for wearing or holding any equipment. This has the added benefit of allowing unimpeded motion of the user's hand in free space.

EHDs have most commonly been grounded robotic arms that enable rendering of various physical objects [20, 48, 50]. For example, Tachi et al.'s Shape Approximation Device [48] uses an end-effector with a collection of surfaces and edges to render various geometric shapes. Snake Charmer renders surface characteristics including texture and temperature, using detachable end-effectors and an off-the-shelf robotic arm [3]. More recently, researchers have explored new forms of EHDs, including shape changing displays [44], surrounding platforms [22], mobile robots [19, 47, 52], and drones [2, 21].

To enable "what you see is what you feel" interactions, EHDs need to achieve a high level of spatial-temporal consistency between the visual and haptic sensory inputs [55]. However, EHDs have limitations that lead to visuo-haptic discrepancies, including limited workspace volumes, positional inaccuracy, and low speeds that may not be able to support real-time interactions. While there has been considerable work on the motion planning [42, 54, 56] and design [57] of novel EHDs, there has been less focus on hardware limitations that tend to make their use challenging. Some researchers have explored the use of multiple robots to better enable just-in-time haptics [24, 47] and dynamically assemble props [61], the additional cost and complexity of a multirobot system may still be impractical, and individual device limitations (e.g., accuracy, speed) may still negatively impact system performance.

Visuo-Haptic Illusions

The human sensory system receives and integrates streams of information from multiple senses simultaneously. When visual input is contradicted by another sensory input, vision tends to dominate such that people perceive the visual information, often without even noticing the conflict; a phenomenon called *visual dominance* [15, 41]. Once the magnitude of the discrepancy exceeds a certain threshold however, such illusions no longer remain undetected [17]. Psuedo-haptics is a technique that leverages the visual dominance effect, below the noticeable thresholds, to create perceived haptic sensations that differ from the real haptic feedback provided [31].

Psuedo-haptics can be utilized to enable a single passive object to represent multiple distinct virtual objects, by modifying the users' perception of various physical properties such as the size, shape, texture, and stiffness of the passive prop [27, 32]. One common technique for creating such illusions is *redirection*, in which the position of the user's virtual hand is altered to create a warped mapping between the physical and virtual spaces. To achieve this effect, prior research has used the boundaries of the physical object as constraints and mapped every point on the physical surface to a different virtual point [60, 28, 29]. Ban et al. use redirection to create the illusion of curvature on flat surfaces [6, 7] and to modify the perceived angle between two parallel lines [5].

Redirection has been applied during reach, prior to users contacting the passive haptic prop. Reach redirection, or *haptic retargeting* [4], takes advantage of the adaptation of the users' sensorimotor loop, guiding their hand to a different position in space such that one passive prop can act as a proxy for multiple virtual objects at different locations. Prior work has proposed different approaches for retargeting: modifying the virtual world's coordinate system (world manipulation), modifying the virtual representation of the users' body (body manipulation), and a hybrid technique [4, 10]. While redirection with passive haptics is effective, there are limits to how much the user's hand can be redirected without becoming too noticeable [59], sacrificing user comfort [10] and performance [18].

Beyond passive haptics, visuo-haptic illusions have been utilized to overcome the limitations of EHDs. Dynamic redirection has been applied to overcome the position inaccuracies of quadcopters, when used as non-grounded EHDs in VR [2]. Researchers have also used control-to-display ratio modifications to compensate for the low resolution and speed of shape displays in VR [1]. Perhaps most relevant to our work, Lee et al. recently explored the use of redirection and visual guidance to address known spatial discrepancies during EHD interaction [33]. While their algorithm compensates for discrepancies known apriori (e.g., workspace constraints), our framework is focused on predicting such discrepancies as they arise in dynamic haptic interaction. Specifically, we address the broader challenge of reachability by also accounting for the dynamic spatial discrepancies induced by the limited speed of EHDs.

THE REACH+ FRAMEWORK

Overview

We propose REACH+: a framework for addressing reachability limitations of EHDs, namely limited workspace and speed, via dynamic reach redirection. This approach synthesizes and builds upon work in encountered-type haptics and visuo-haptic illusions to (1) *predict* which element of a virtual scene the user is reaching to touch and when they will make contact, (2) *anticipate* the inability of an EHD to arrive at that element, due to either workspace or speed constraints, (3) *remap* the virtual element to a physical point within the EHD's reachable space, and (4) *redirect* the user's real hand to the reachable point as they reach for the virtual element. The goal of REACH+ is to expand the perceived capabilities of EHDs, allowing devices ordinarily constrained by their workspace or speed to render elements in a virtual scene with greater success on-the-fly.

Assumptions

In this work, we assume a known, discrete set of N virtual objects which the user can interact with using a single hand, shown in Figure 3. We assume each object can be represented by a single point (e.g., a button) within a virtual workspace volume (V) and that the EHD has a known workspace (P). Furthermore, we consider only translational offsets between objects in V and P in this work, as such offsets are the most likely to arise from device limitations (as shown in Figure 3).

Interaction Modeling & Prediction

REACH+ leverages the prediction of user interaction. An estimate of the user's intended reach target and arrival time

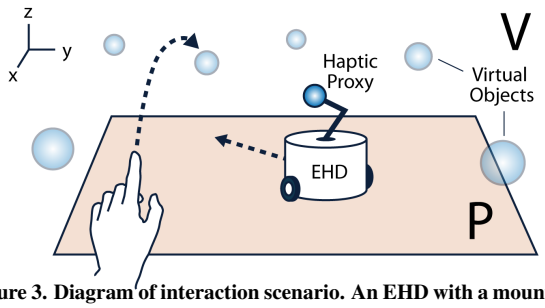

Figure 3. Diagram of interaction scenario. An EHD with a mounted haptic proxy moves within its workspace to render a set of virtual objects.

(ETA) to determine the reachable space of the EHD and apply subsequent redirection. We consider the scenario of a user freely reaching to touch one of N objects in a virtual scene with one hand. In order to make these predictions and apply the results in real-time, a suitable interaction model is needed.

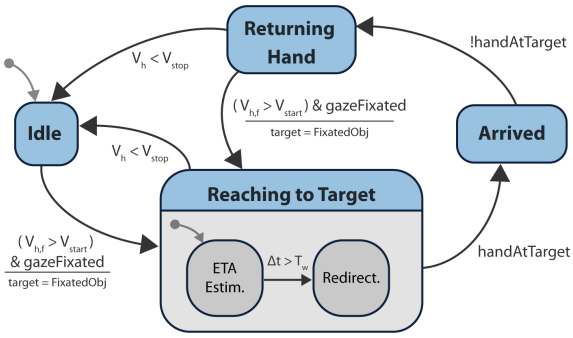

Figure 4. State machine used to model goal-directed reach interactions with REACH+. V_h is hand speed and $V_{h,f}$ is hand speed in the direction of the object. Blue is interaction states; gray is sub-states used in REACH+.

User Reach State Machine

We model this interaction using the state machine in Figure 4, where each state represents a different stage of interaction: (1) *Idle*, during which the user is visually scanning the scene, (2) *Reaching to Target*, which occurs after the user fixates on an object and begins moving their hand toward it, (3) *Arrived at Object*, and (4) *Returning Hand*. As humans tend to fixate on a target prior to initiating a goal-directed reach [39, 40], gaze plays an important role both in the early detection of reaching movements and prediction of reach targets. Prior research has also shown hand velocity to be an important indicator of reach intent [10, 45]. We use both gaze fixation and hand velocity thresholds to model transitions between the above states.

Intended Target Prediction

In our model, target prediction is encoded in the state transition from *Idle* to *Reaching to Target*. Cheng et al. [10] demonstrated high target prediction accuracy (97%) for a similar task using gaze fixation and a hand velocity threshold as a heuristic predictor; we use this strategy to predict intended targets. We use Binsted et al.'s approach for detecting gaze fixation [8]: if the standard deviation of gaze angle over a moving window of 0.6 s falls below 1°, gaze is considered fixated. If the gaze ray collides with a potential target during fixation, then that object is considered to be fixated on. In our implementation, an HTC Vive Pro Eye headset was used with the Tobii XR SDK to measure and process gaze data (detailed in the EHD

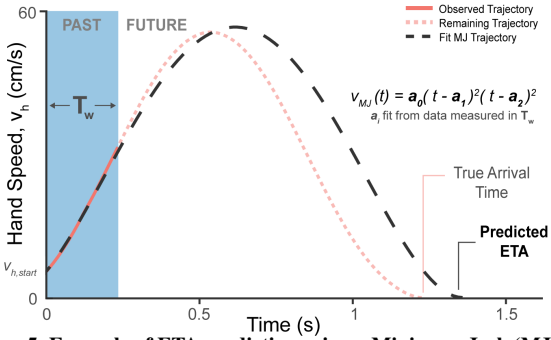

Figure 5. Example of ETA prediction using a Minimum Jerk (MJ) fit on measured hand speed, collected from the start of reach ($v_h > v_{h,start}$) over window T_w. An MJ model is then fit and predicted ETA determined.

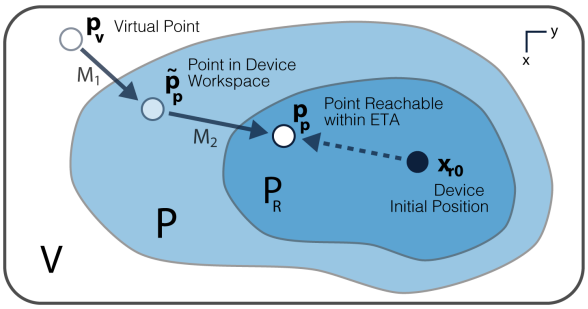

Figure 6. Predicted target is first mapped from virtual space (V) into the EHD's workspace (P), then into the space of points reachable within the predicted ETA (P_R). Boundaries shown in 2D for clarity.

Implementation section). Similar to [10], if hand speed in the direction of the fixated object ($v_{h,f}$) is above a threshold (v_{start}), that object is predicted to be the user's reach target and the state transitions from *Idle* to *Reaching to Target*. Based on empirical testing, in this work we use a threshold of 3 cm/s.

Arrival Time Estimation via Minimum Jerk

Once the user begins *Reaching to Target*, their intended target has been predicted, but an estimate of their arrival time (ETA) is needed to determine if the EHD can position itself in time.

Similar to [9], we use a brief segment of the user's reach trajectory to inform ETA prediction by fitting this partial trajectory to a Minimum Jerk (MJ) model. The MJ model [14] is a well-studied kinematic model for goal-directed human limb motion, used to accurately explain experimental data [25] and predict features of human trajectories [9, 16, 46]. The MJ model suggests humans minimize the derivative of their hand's acceleration (i.e., jerk) during point-to-point motion. When starting and stopping at rest, this results in a symmetric, bell-shaped hand speed profile defined by a 4$^{\text{th}}$ order polynomial:

$$v(t) = a_0(t-a_1)^2(t-a_2)^2 \tag{1}$$

Here t is time, $v(t)$ is tangential hand speed, and a_i are free parameters to be fit; a_0 is proportional to predicted peak speed, and a_1 and a_2 encode predicted reach start and stop times.

This approach is illustrated in Figure 5, where T_w represents the window of time after entering *Reaching to Target* in which data is collected to make a prediction. Larger windows yield more accurate ETAs at the expense of a greater delay in prediction. Larger delays are less desirable because they leave less time and distance over which to apply redirection. Based on pilot testing and preliminary study [16], we selected a window of 100 ms, or about 10% of a 1 second reach trajectory.

Reachability Remapping

With the user's intended target predicted and ETA estimated, the next step is to map the target's position in virtual space ($\mathbf{p_v}$) to a point within the EHD's temporally reachable space ($\mathbf{p}_p$). We refer to this process as *reachability remapping*. We first generate a point within the device's workspace, and then adjust based on the device's ability to arrive to that point.

Workspace Constraints

We first address the issue of $\mathbf{p_v}$ lying outside of the EHD's workspace volume due to kinematic limitations. Our task is to find a suitable mapping from point $\mathbf{p_v}$ in the virtual workspace V to a point $\tilde{\mathbf{p}}_\mathbf{p}$ in the EHD's workspace P (Figure 6).

A number of spatial remapping techniques have been explored to map points from virtual space to physical space, including ergonomic optimization [38], spatial consistency [38], line-of-sight [10], and nearest-point [10]. An in-depth discussion of each is beyond scope here, but it is important to note that mapping functions may be used if desired.

We use the nearest-point method since it is simple and has been shown to be well-preferred among users [10]. With this method, objects in V are mapped to the nearest point in P, resulting in minimal spatial discrepancy between $\mathbf{p_v}$ and $\tilde{\mathbf{p}}_\mathbf{p}$:

$$\tilde{\mathbf{p}}_\mathbf{p} := \arg\min_{\mathbf{x}\in P} ||\mathbf{x} - \mathbf{p_v}||^2, \tag{2}$$

resulting in a point $\tilde{\mathbf{p}}_\mathbf{p}$ corresponding to $\mathbf{p_v}$ which can *feasibly* be reached by the EHD, but does not gaurantee on-time arrival.

Temporal Constraints

To ensure the EHD can arrive in time for contact, the final proxy point for $\mathbf{p_v}$ should be reachable by the device within the predicted ETA. That is, our task is to determine a second mapping from point $\tilde{\mathbf{p}}_\mathbf{p}$ in P to a point $\mathbf{p_p}$ in P_R, where P_R is the subset of P reachable within the predicted ETA (Figure 6).

Let t_0 be the time at which ETA estimate $\hat{t}_f$ was obtained, and let $\mathbf{x_r}(t)$ be the position of the EHD as a function of time. In general, P_R can be computed for an EHD given its initial position $\mathbf{x_{r0}}$ and a kinematic model of its motion. To simplify this calculation, here we assume an omnidirectional device moving with constant speed S along the line between $\mathbf{x_r}(t_0)$ and $\tilde{\mathbf{p}}_\mathbf{p}$. Let $\hat{\mathbf{d}}$ be the unit vector from $\mathbf{x_{r0}}$ to $\tilde{\mathbf{p}}_\mathbf{p}$. The furthest point reachable by the EHD within $\hat{t}_f$ is estimated as:

$$\hat{\mathbf{x}}_\mathbf{r}(\hat{t}_f) = \mathbf{x_{r0}} + S(\hat{t}_f - t_0)\hat{\mathbf{d}} \tag{3}$$

This point can be considered the boundary of P_R in the $\hat{\mathbf{d}}$ direction. If $\tilde{p}_p$ lies within that boundary, no additional mapping is necessary. Else, we set p_p to the reachable point $\hat{x}_r(\hat{t}_f)$:

$$\mathbf{p_p} = \begin{cases} \hat{\mathbf{x}}_\mathbf{r}(\hat{t}_f) & ||\hat{\mathbf{x}}_\mathbf{r}(\hat{t}_f) - \mathbf{x_{r0}}|| < D, \\ \tilde{\mathbf{p}}_\mathbf{p} & ||\hat{\mathbf{x}}_\mathbf{r}(\hat{t}_f) - \mathbf{x_{r0}}|| \geq D \end{cases} \tag{4}$$

where D be the distance between $\tilde{\mathbf{p}}_\mathbf{p}$ and $\mathbf{x_{r0}}$. The point $\mathbf{p_p}$ serves as the physical proxy location corresponding to the predicted virtual object at $\mathbf{p_v}$. This then becomes the goal

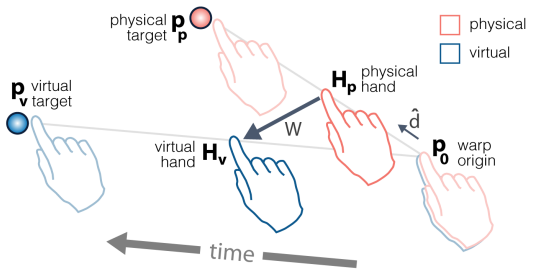

Figure 7. Diagram of reach redirection [4]. At each timestep, virtual hand position is computed as a vector offset W from the physical hand position, based on the hand's displacement in the target direction $\hat{d}$.

position of the EHD, as well as the physical target point for reach redirection, described in the following section.

Limiting Redirection Magnitude

While the remapping described thus far emphasizes the reachability of point $\mathbf{p_p}$ by the EHD, there is also an implicit tradeoff in that larger offsets between $\mathbf{p_p}$ and $\mathbf{p_v}$ require more redirection. Prior work has made it clear that large hand redirection is less tolerable to users [10, 59] and negatively impacts realism, embodiment, and task performance [18, 29]. Thus, limiting the maximum distance of $\mathbf{p_p}$ from $\mathbf{p_v}$ may be desirable to ensure redirection remains undetectable [59]. Critically, limiting redirection leads to a greater risk of delay, however. Such a limit can be applied by updating $\mathbf{p_p}$ according to:

$$\mathbf{p_p} := \mathbf{p_v} + \min(||\mathbf{T}||, T_{max})(\mathbf{T}/||\mathbf{T}||) \quad (5)$$

where $\mathbf{T} = \mathbf{p_p} - \mathbf{p_v}$ and T_{max} is the maximum offset.

Reach Redirection

With the EHD goal position $\mathbf{p_p}$ computed, our final step is to redirect the user's hand to $\mathbf{p_p}$ over the remaining reach. For this, we use the haptic retargeting technique presented by Azmandian et al. [4]. Illustrated in Figure 7, this technique gradually adds a vector offset $\mathbf{W}$ between the user's physical hand $\mathbf{H_p}$ and their virtual hand $\mathbf{H_v}$ as they reach to the virtual target $\mathbf{p_v}$. At each timestep during redirection, we compute:

$$\mathbf{H_v} = \mathbf{H_p} + \mathbf{W} \quad (6)$$

$$\mathbf{W} = \left(\frac{D_s}{D_s + D_p}\right)(\mathbf{p_v} - \mathbf{p_p}) \quad (7)$$

where $D_s = ||proj_{\hat{\mathbf{d}}}(\mathbf{H_p} - \mathbf{p_0})||$ and $D_p = ||proj_{\hat{\mathbf{d}}}(\mathbf{p_p} - \mathbf{H_p})||$. Here $\mathbf{p_0}$ is the warp origin and $\hat{\mathbf{d}}$ is the unit vector point from $\mathbf{p_0}$ to $\mathbf{p_p}$. These can be thought of as the distance travelled by the hand towards the target, and the distance remaining to the target, respectively. Due to visual dominance, the user corrects for this gradual offset, redirecting their reach such that their physical hand arrives at $\mathbf{p_p}$ as their virtual hand arrives at $\mathbf{p_v}$.

EHD IMPLEMENTATION

We built an omnidirectional mobile robot to serve as a desktop EHD to demonstrate and evaluate the REACH+ framework.

The device, shown in Figure 8, consists of three 12V high-power gearmotors with integrated quadrature encoders (Pololu #4845) mounted to an acrylic chassis. Each motor powers

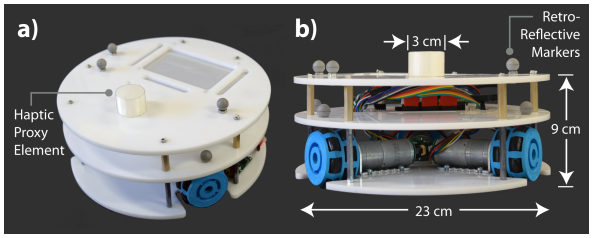

Figure 8. Custom desktop-scale omnidirectional mobile robot with a 3D printed cylinder as a haptic proxy, used for encountered-type haptics.

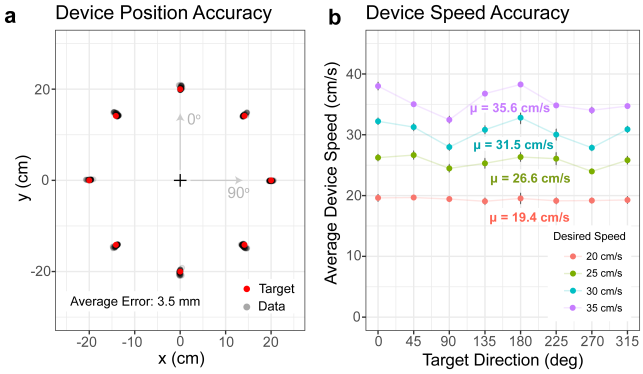

Figure 9. Results of device validation. (a) Final device positions (black) when commanded to each target (red). (b) Average recorded speed ± SD in each direction for each commanded speed.

an omnidirectional wheel and is driven by a VNH5019 high-power motor driver. The entire system is controlled by an on-board microcontroller (Teensy 3.6). Mounted on top of the chassis are three retro-reflective markers for optical tracking, and a passive 3D printed cylinder (30 mm dia., 20 mm height), which serves as the haptic end-effector in the present study; we selected 30 mm as it is a standard pushbutton diameter.

It is important to note that this EHD is *planar* – meaning it can freely traverse in the x & y dimensions, but cannot adjust the end effector in the z dimension. We chose to use a planar, omnidirectional device for the present evaluation because control of the end-effector position is simple and accurate, without requiring a more costly or complex device.

Device Validation

Since we aim to evaluate the REACH+ framework using this custom device, it is important to characterize its performance and ensure commanded positions and speeds were sufficiently accurate. We evaluated the positioning and speed accuracy of the device across four speeds (20, 25, 30, and 35 cm/s) and 8 target locations. The 8 targets were spaced evenly about a circle (d = 40 cm). The EHD was commanded to move from the center of this circle to each target 20 times for each desired speed, for a total of 640 trials. Results are shown in Figure 9.

Positioning Accuracy

A PID position controller was used to govern EHD motion, with a saturation limit used to limit device speed to a specified amount. The average positioning error across all trials and speeds was 3.5 mm ± 2.1 mm (SD). Given the average human index finger tip width of 20 mm [11] and proxy diameter of 30 mm, we determined this to be an acceptable amount of error.

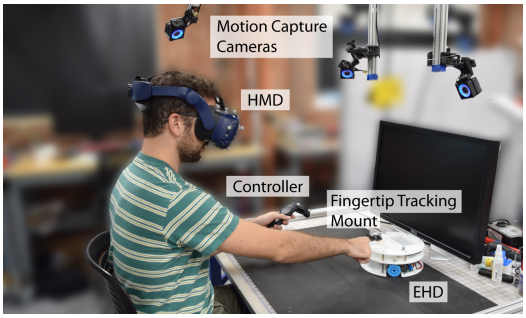

Figure 10. The experiment setup used in the evaluation.

Speed Accuracy

Average speed errors of -0.6 $\pm$ 0.6 cm/s, 1.6 $\pm$ 1.1 cm/s, 1.5 $\pm$ 1.8 cm/s, and 0.6 $\pm$ 1.9 cm/s were found for the commanded speeds of 20, 25, 30, and 35 cm/s, respectively. Illustrated by Figure 9b, this variance can be explained in part by moderate directionality of the device, particularly at higher speeds; note that there was little variance within each direction. While a more sophisticated controller could easily account for such errors and directionality, we chose to keep the simple PID control scheme to mirror the kinds of modeling and uncertainty expected in low-cost practical devices. Furthermore, since errors compared to desired speed were reasonably small ($<$10%), we found this to be an acceptable level of accuracy.

EVALUATION

As prior work has explored the use of redirection to compensate for known workspace limits of EHDs [33], we focused on the limited device speed, which is more involved due to its temporal nature: the redirection required depends on the expected time the device has to reach the user's target, estimated from their movement. We designed a target selection task to evaluate the effectiveness of REACH+ across different EHD speeds. We tested three redirection conditions: a control ("No Redirection"), our framework with no limit on redirection ("REACH+"), and our framework with a 3 cm redirection limit ("Limited REACH+"). We evaluated the impact of each on EHD performance, as measured by device arrival latency and subjective user experience.

Assumptions

As intent prediction is not the focus of this evaluation, we assume the participant's intended target is known. We assume a constant velocity model of EHD motion, and compute reachability based on the ETA obtained from using the minimum jerk method described previously. We also note that all virtual targets lie within the planar physical workspace of the EHD (i.e., $\mathbf{p_p} = \mathbf{p_v}$); offsets between physical and virtual targets therefore arise from speed limitations of the EHD.

Participants

We recruited 11 right-handed participants (8 male, 3 female), ages 22-44 ($\mu = 28.7$, $\sigma = 6.6$), for a $15 compensation.

Experimental Setup

Apparatus

The setup (Figure 10) consisted of the EHD described in the previous section, an HTC Vive Pro Eye head-mounted display (HMD), an HTC Vive controller for user input, noise-cancelling headphones, an OptiTrack motion capture system, and retro-reflective markers. To enable real-time tracking of the participant, markers were fit to the user's right index finger using a 3D printed mount. Markers were positioned on the EHD to enable position and orientation tracking of the device. Motion data was streamed using Motive software to Unity.

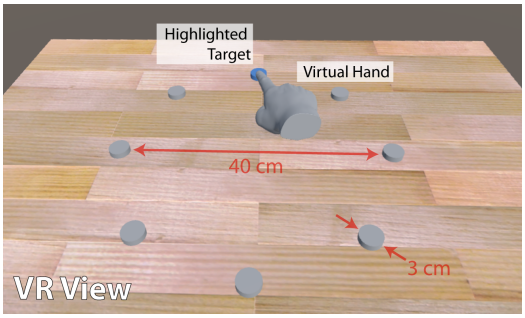

Figure 11. Participant's VR viewpoint in the study. Labels and red annotations not visible to users. Trials began with the EHD at the center of the ring of targets; the user reached to the blue highlighted target.

Target Layout

Participants observed 8 targets arranged in a 40 cm diameter circle (Figure 11). For each trial, the EHD proxy was repositioned to be at the center of the ring of targets; the orientation of the EHD was fixed in the study. This selected configuration controlled for directionality and distance to each target, while maximizing workspace coverage. The reach starting point was 20 cm above and 35 cm back from the center of the ring.

Conditions

Four EHD speeds were evaluated: 20, 25, 30, and 35 cm/s. These were selected based on pilot testing to achieve a spread of EHD performance; additionally, 35 cm/s is the maximum speed achievable by the tested EHD. Three redirection conditions were also investigated: None, Limited REACH+, and REACH+. In the None condition, no redirection was applied during the course of the reach. In the REACH+ condition, arrival time to the highlighted target was estimated and the physical location of the virtual target was remapped for device reachability according to the previous section. In the Limited REACH+ condition, we took the same approach but imposed a limit of 3 cm on the amount of offset between the virtual target and its remapped physical position. This limit was chosen by averaging the estimated detection threshold for each target, computed based on [59]. Thus a total of 12 experimental conditions (4 speeds $\times$ 3 redirection methods) were evaluated.

Procedure

Participants were first seated at the study table and donned the HMD and index-fingertracking markers. Following an introduction to the study, participants completed a reach speed training session. The goal was to roughly standardize the reach speed of each participant to ensure results could be compared; however, they were instructed to reach as naturally as possible. A target reach speed of 45 cm/s was selected from pilot trials. Participants completed a minimum of 30 practice reaches, until the target speed was reliably reached within $\pm$5 cm/s.

For each of the 12 conditions, participants completed 2 reaches to each of the 8 targets for a total of 16 trials per condition, and $12 \times 16 = 192$ total trials per participant. The experiment was blocked by condition, which were randomized. Target order within each condition was also randomized.

Each trial began with the participant's right index finger at the indicated starting point. One of the targets was then highlighted, cueing the user to begin their reach. Users were instructed to smoothly reach and touch the center of the target circle with their index finger. Upon the user's finger leaving the starting point, the EHD began moving toward the highlighted target. In the REACH+ and Limited REACH+ conditions, after 100 ms a temporally-reachable point ($\mathbf{p_p}$) was computed using an estimate of the user's arrival time to the target (ETA) and the constant-velocity model of EHD motion, as detailed in the previous sections. Reach redirection was then initiated to map the virtual target to the temporally-reachable point.

Similar to other work in haptics [21, 23, 53] and redirection [4, 10, 60], we used a self-report tool to gauge subjective user experience after each trial. After arriving at the virtual target, participants indicated their agreement on a 5 point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree) with the following two statements:

Q1. The movement of my virtual hand matched my physical hand.
Q2. The physical contact matched my expectations of touching a real object.

Q1 captures how noticeable the redirection was to the participant. Note that participants were told at the start that there may be times where the movement of the physical and virtual hands may not be one-to-one. Q2 captures how realistic the physical touch interaction felt. For example, if the EHD was significantly delayed in arriving to the target, a participant would rate their agreement as low since that violates their expectations of normal touch interaction. Participants were instructed not to consider offsets between their physical and virtual hands when answering Q2. They logged their answers using an HTC Vive controller held in their left hand. After logging their response to Q1 and Q2, they returned their fingertip to the starting point and the EHD returned to the center of the ring of targets. After both arrived, the next trial began following a 2 second delay. Participants had three 2-minute breaks during the study.

Results

Out of the 2112 trials, 25 were discarded due to recording errors. Furthermore, reaches lasting more than 5 seconds (8 trials) were discarded as outliers, leaving 2079 trials in our analysis. We first present a summary of participants' reach characteristics to provide context for later results. We then report the accuracy of our arrival time prediction and describe the range of redirection magnitudes computed and applied during the study. Finally, we present the study results on device latency, on-time arrival to target, and subjective ratings of redirection noticeability (Q1) and physical realism (Q2).

Reach Characteristics

Figures 12a and 12b illustrate the reach times and speeds for all participants, respectively, grouped by redirection method.

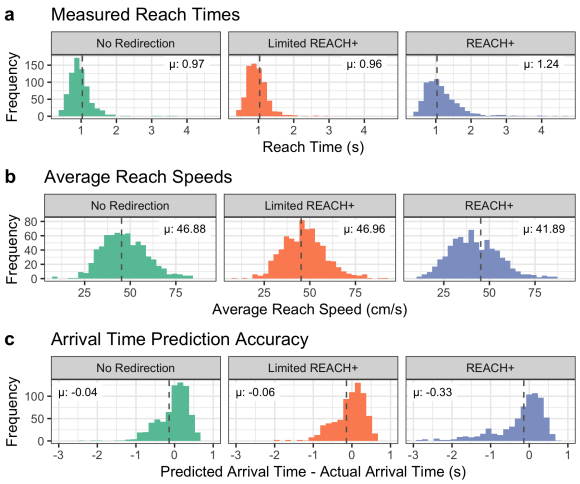

Figure 12. (a) Reach times recorded for each redirection method. (b) Average reach speed recorded for each redirection method. (c) Error in arrival time prediction for each redirection method. Vertical lines indicate averages. Data shown for all participants.

Note that reach time represents between reach onset and arrival of the virtual finger to the virtual target, and reach speed represents the average recorded hand speed during reaching in a trial. A pair of one-way repeated measures ANOVAs revealed a significant effect of redirection method on reach time ($F(2,2068) = 72.13$, $p < 0.001$) and reach speed ($F(2,2068) = 40.305$, $p < 0.001$). Post-hoc comparisons using Tukey's HSD test indicated this resulted from significant differences between REACH+ and the other two methods for both reach time ($p < 0.001$) and reach speed ($p < 0.001$). No difference was found between No Redirection and Limited REACH+ either for reach time ($p = 0.96$) or speed ($p = 0.99$).

Arrival Time Prediction

For each trial, we estimated participants' arrival time to the highlighted target. Following the method detailed in the previous section, we fit a MJ model to the first 100 ms of hand speed data and use the zeros of the fitted polynomial (1) to estimate when the hand will arrive at the target. We use this to determine if the EHD can arrive at the target in time and, if not, how much redirection is needed. This estimation was performed for every reach; note that no redirection occurred in the first 100 ms of any reach. Figure 12c shows the arrival time prediction errors found for all trials. Here, a negative value indicates a predicted arrival that was earlier than the true arrival time. On average, the predicted arrival time was early by 40 ms in No Redirection, 60 ms in Limited REACH+, and 330 ms in REACH+. These correspond to a 4.1%, 6.3%, and 26.6% error, respectively, with respect to the average reach time of each condition. A one-way repeated measures ANOVA confirmed a significant effect of redirection method on arrival time prediction accuracy ($F(2,2068) = 38.19$, $p < 0.001$), post-hoc comparisons using Tukey's HSD test revealed no significant difference between arrival time accuracies of the No Redirection and Limited REACH+ methods ($p = 0.86$). Significant differences were found between REACH+ and No Redirection ($p < 0.001$), as well as Limited REACH+ ($p < 0.001$).

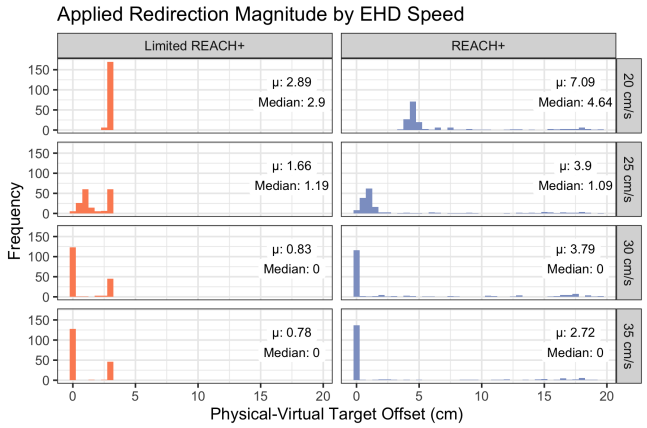

Figure 13. Magnitude of all applied redirection grouped by EHD speed (rows) and redirection method (columns). Bin width is 1 cm.

Applied Redirection

Figure 13 shows the distribution of applied redirection magnitudes in both REACH+ redirection conditions for all EHD device speeds, as well as the mean and median values for each. Note that in the Limited REACH+ condition, 3 cm was set as a saturation limit for the offset.

Looking at the REACH+ trials, we observe the computed redirection magnitude without a saturation. We see that the 20 cm/s EHD on average was predicted to require 7.09 cm of compensation via redirection, while the 25 cm/s EHD was predicted to require 3.9 cm on average. Both the 30 cm/s EHD and 35 cm/s EHD were predicted to require no redirection (0 cm) in over 80% of all trials. In the Limited REACH+ trials, these patterns translate to frequent saturation when using the 20 cm/s EHD ($\mu = 2.89$ cm); that is, more redirection was estimated to be required, but was limited to 3 cm. Redirection with the 25 cm/s EHD yielded the largest spread between 0 cm and 3 cm ($\mu = 1.06$ cm). Similar to the REACH+ trials, the 30 cm/s EHD and 35 cm/s EHD required little redirection ($\mu = 0.83$ cm and $\mu = 0.78$ cm, respectively).

Note that due to our prediction method underestimating the participant's arrival time, larger-than-necessary amounts of required redirection were at times computed. In the REACH+ conditions, these are illustrated by the relatively infrequent but wide spread of offsets ranging between 5 cm and 20 cm. For Limited REACH+, such offsets were saturated to 3 cm.

Device Latency

Latency between what the user sees and feels is an important objective measure of performance in any haptic device [12, 13, 51], but particularly for EHDs which need to physically travel to the contact point. We define visuo-haptic device latency as the time difference between the virtual fingertip colliding with the virtual target, and the participant's physical fingertip contacting the physical proxy mounted on the EHD. Contact was recorded using OptiTrack motion capture data (< 0.5 mm accuracy). If significant latency is present, it reduces visuo-haptic coherence and detracts from user experience.

Figure 14a shows the mean latencies recorded during each of the evaluated conditions. A linear mixed effects model was fit to the data, with device speed and redirection condition as

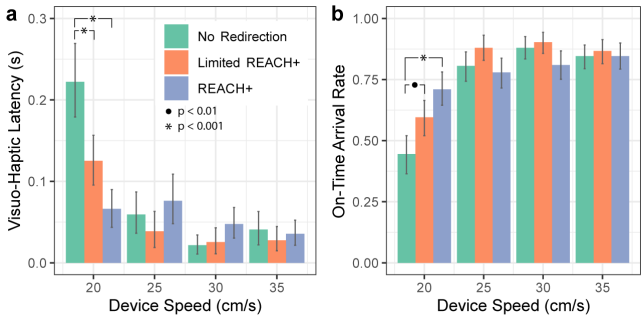

Figure 14. (a) Mean visuo-haptic latency in each condition. (b) On-time arrival rate of EHD in each condition. 95% bootstrap CIs shown.

fixed effects and participant as a random effect. A type II analysis of deviance (ANODE) revealed significant main effects of both device speed ($\chi^2(3, N = 2079) = 115.63, p < 0.001$) and redirection condition ($\chi^2(2, N = 2079) = 9.75, p < 0.01$). A significant interaction effect was also found ($\chi^2(6, N = 2079) = 51.73, p < 0.001$).

The results of post-hoc comparisons using Tukey's HSD test are also shown in Figure 14a. Within the slowest device speed (20 cm/s), latency during No Redirection was significantly greater than during Limited REACH+ or REACH+ ($p < 0.05$). No significant differences between redirection conditions were found within any of the other device speeds.

On-time Arrival

Vogels et al. recommended that latency in visuo-haptic interfaces should be kept below 45 ms, as larger delays are perceptible to users and detract from their sense of realism [49]. To provide further context to the reported latency results, we computed a binary "on-time arrival" variable for each trial, logging whether device latency was below 45 ms (considered to arrive on-time) or above (considered to arrive late).

Figure 14b shows the overall On-time Arrival Rates (OAR) for each condition. A mixed effects logistic regression model was fit to the data with device speed and redirection condition as fixed effects, and participant as a random effect. A type II ANODE revealed significant main effects of both device speed ($\chi^2(3, N = 2079) = 156.53, p < 0.001$) and redirection ($\chi^2(2, N = 2079) = 10.33, p < 0.01$). A significant interaction effect was also found ($\chi^2(6, N = 2079) = 36.59, p < 0.001$).

The results of post-hoc comparisons using Tukey's HSD test are also shown in Figure 14b. Within the slowest device speed (20 cm/s), both REACH+ and Limited REACH+ demonstrated significant improvements in OAR compared to No Redirection. No significant differences between redirection conditions were found within any of the other device speeds.

Subjective Ratings

Figure 15 shows the average ratings participants gave for the two subjective questions asked after each trial. We fit a Cumulative Link Mixed Model (CLMM) to the ordinal data from each question. Device speed and redirection condition were set as fixed effects, and participant was set as a random effect.

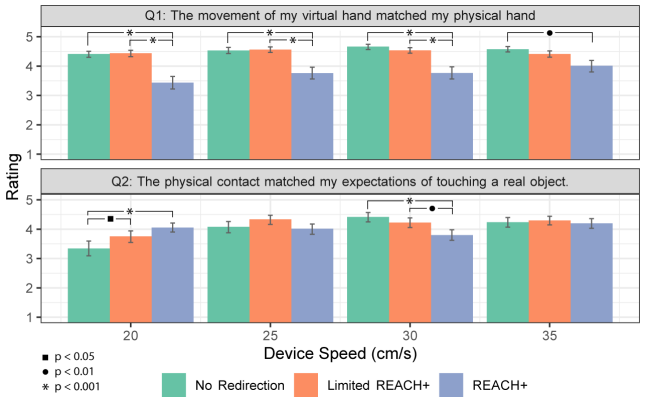

Figure 15. Average ratings for each question. 1 = Strongly Disagree, 5 = Strongly Agree. 95% bootstrap CIs shown.

Q1 was intended to measure the extent to which participants noticed any redirection during the trial, with a rating of 5 indicating no redirection was noticed. A type II ANODE revealed significant main effects of both device speed ($\chi^2(3, N = 2079) = 22.56, p < 0.001$) and redirection condition ($\chi^2(2, N = 2079) = 202.54, p < 0.001$). A significant interaction effect was also found ($\chi^2(6, N = 2079) = 22.97, p < 0.001$). Post-hoc comparisons using Tukey's HSD test revealed that REACH+ yielded significantly lower Q1 ratings ($p < 0.01$) than both Limited REACH+ and No Redirection, within each device speed. No difference was found between No Redirection and Limited REACH+ within any speed.

Q2 was intended to measure the subjective realism of each physical touch interaction, with a rating of 5 indicating an interaction that strongly matched expectations. A type II ANODE revealed significant main effects of both device speed ($\chi^2(3, N = 2079) = 53.96, p < 0.001$) and redirection condition ($\chi^2(2, N = 2079) = 7.05, p < 0.05$). A significant interaction effect was also found ($\chi^2(6, N = 2079) = 66.49, p < 0.001$). Post-hoc comparisons using Tukey's HSD test revealed that both Limited REACH+ and REACH+ yielded significantly improved Q2 ratings compared to No Redirection, within 20 cm/s device speed trials ($p < 0.01$).

Discussion

Overall, the increase in on-time arrival rate (OAR) and self-reported realism, measured via Q1 (visuo-proprioceptive coherence) and Q2 (physical realism), demonstrate the potential for REACH+ to improve the performance of slower EHDs through dynamic redirection. Without redirection, the 20 cm/s EHD yielded large visuo-haptic latencies, lower OAR, and lower Q2 ratings. Under REACH+, however, the same speed yielded less than half as much latency, a 25% increase in OAR, and a significant increase in Q2 rating. As expected, these improvements came at the expense of lower Q1 ratings, indicating users felt a noticeable offset in the positions of their physical and virtual hand. This emphasizes a key trade-off of our framework: the closer $\mathbf{p_p}$ is moved to the EHD's initial position, the less likelihood of device latency; conversely, the further p_p is shifted from the virtual target, the greater redirection needed, potentially detracting from the experience.

Thus, the results of the Limited REACH+ trials (where redirection was limited to 3 cm) are of particular interest. For the slowest device (20 cm/s), Limited REACH+ lead to comparable improvements as REACH+ – a 28% increase in OAR and a significant increase in self-reported physical realism (Q2) – while yielding virtually no noticeable difference between virtual and physical hand motion (Q1). These results suggests that device performance can be improved even when redirection is kept at or below noticeable levels. Additionally, from Figure 12, we see that Limited REACH+ did not yield significantly different reach times, reach speeds, or ETA prediction accuracy compared to No Redirection. In contrast, REACH+ (which often required larger redirection, see Figure 13), tended to yield slower reaches and larger negative errors in predicted ETA. These results suggest that setting a limit on redirection (here 3 cm) may be one way to achieve a balance of EHD performance improvement with mitigating the impact of redirection on user experience. It is worth noting participants experienced different redirections each trial (depending on the target and study condition), likely contributing to the reduced reach speed and performance of significantly redirected reaches. As users experience more consistent mappings, however, some adaptation can be expected [18, 30].

More generally, these results also demonstrate the potential of minimum jerk-based estimation of reach arrival time computed on-the-fly. Figure 12c indicates predictions were reasonably accurate, though with some variance (IQR ≈ 0.5 s). In this study, a window of 100 ms from reach onset was used to fit the model and generate an estimate based on pilot testing; a longer window may lead to more accurate/precise estimates, however at the cost of delaying the onset of redirection. It is important to note that because redirection itself influences reach behavior (which we do not account for), this estimation of arrival time is perhaps most generally useful as a heuristic to estimate whether the EHD can arrive at the original target ($\mathbf{p_v}$) in time. The reduction in ETA accuracy for more highly redirected reaches indicates that motion models such as the minimum jerk model may require modification for more accurate use with redirected motion, as suggested in [16].

Perhaps most clear from these results is the influence of device speed on EHD performance. At speeds above 20 cm/s, no significant differences were found between redirection conditions in terms of visuo-haptic latency or on-time arrival. Overall, we see a plateauing trend for both device latency and OAR as speed is increased (Figure 14). Qualitatively, however, participants still tended to report a mismatch between their physical and virtual hands (Q1) in the REACH+ condition at higher device speeds. This is likely because negative errors in arrival time prediction (Figure 12c) propagate as extraneous redirection in the REACH+ conditions (Figure 13, right col.).

It is important to note that in practice, EHD performance is determined largely by the travel distance required to reach the target as well as the timing of the user's reach. For this study, maximum travel distance was fixed at 20 cm; if this value were larger (say, arm's length) it's likely that the faster EHDs would begin to suffer from similar latency issues as the 20 cm/s EHD, though further testing is needed to confirm these effects.

LIMITATIONS AND FUTURE WORK

In our study, we focused on the latency compensation component of REACH+, as latency compensation via redirection had not been previously explored in the literature. Further work is needed to assess the effectiveness of REACH+ to compensate for both arrival latency and limited workspace simultaneously. In addition, we assumed perfect knowledge of the user's intended target in order to best evaluate latency compensation in a controlled manner. While previous studies with similar target prediction algorithms demonstrated promising results [10], future work should directly compare the proposed algorithm with other multi-target prediction strategies, such as recursive Bayesian filtering [58] and inverse optimal control [37].

The findings of the evaluation are also dependant on the required travel distance of the EHD (20 cm) and user reach speed (trained to 45 cm/s). For example, a larger travel distance would have likely resulted in lower rates of on-time arrival even for the faster EHD speeds; we expect that REACH+ would be beneficial for those devices in such a scenario.

Additionally, a fixed redirection limit of 3 cm was used in the Limited REACH+ condition It may be more generalizable, however, to compute the maximum allowable redirection based on the relative position of the target to the hand as well as known or estimated perceptual thresholds, such as those reported by Zenner et al. [59]. Moreover, our tabletop robot EHD resulted in planar physical-virtual offsets and the evaluation of REACH+ should be extended to offsets in 3D space.

The present implementation of REACH+ faces also certain limitations. Due to the observed data required to make predictions about the user's intended target and arrival time, redirection is always started part-way through the reach. In most previous work, required offsets are known a priori and redirection is started at reach onset [4, 29, 59]. While Cheng et al. [10] also explore dynamically updating redirection based on predicted target, more work is needed to understand how mid-reach redirection influences user perception and task performance. Similarly, while we used a fixed data collection time window to make predictions, future work should further explore the trade-off between prediction accuracy and redirection delay.

One significant limitation is that of the limits of visuo-haptic redirection. Recent literature suggests that imperceptible redirection may only work for small displacements [59]. Beyond these thresholds, the redirection becomes noticeable and may impede natural human reaching. Further work is needed to investigate these thresholds and human reach performance beyond them. Overall, this limits the amount of spatio-temporal discrepancy that can easily be improved by REACH+.

The REACH+ framework focuses on addressing reachability issues caused by device workspace and speed limitations; however, another source of spatial discrepancy is positional inaccuracy. Compensating for positional inaccuracy is particularly challenging in this framework because it relies on pre-emptively correcting for predicted discrepancy, and positioning errors are often difficult to predict precisely. While we ensured the accuracy of our EHD implementation (3.5 mm) was suitable for the given task, different contexts and devices may result in different accuracy requirements. Abtahi et al. [2] presented one promising approach to inaccuracy compensation for a drone-based EHD, which involved applying a form of haptic retargeting paramaterized by hand displacement rather than absolute position, allowing the proxy point $\mathbf{p_p}$ to update during reaching without shifting the virtual hand. Another approach could be to correct for positional inaccuracy through redirected touching [28] during haptic exploration. Furthermore, while this work specifically focuses on the use of redirection to address reachability limitations of EHDs, a broader exploration of the trade-offs and benefits of traditional haptic retargeting versus encountered-type haptics approaches would provide valuable context for future work.

Another challenge facing REACH+ is that the MJ-based arrival time prediction does not account for the effect of redirection on reach speed profile. Our work and prior work [16] highlight that redirection can lead to slower reaches and more significant ETA error. In the future, better models of reaching motion under redirection can be incorporated into REACH+ to improve arrival time prediction under redirection and to determine the amount of redirection needed with higher accuracy. Future work might also consider more sophisticated EHD motion planning,incorporating multi-hand interactions with multiple EHDs, where device assignment, scheduling, and motion are optimized in parallel with redirection.

Finally, while the evaluated scenario was simple and constrained, we believe the benefits of REACH+ may become even more apparent in applications with a wider interaction space. For example, when interacting with a large control panel (Figure 1), users may need to quickly interact with different widgets that span an entire desktop. Similarly, in VR games users may wish to interact with objects in front of them, beside them, or behind them. Supporting dynamic encountered-type haptics for such person-scale workspaces is highly demanding, even for multiple EHDs, and further exacerbates latency issues. By reducing the required travel distance of EHDs and enabling them to render objects beyond their reachable limits, REACH+ can potentially better support wide-area haptic interactions in VR. Future work will directly explore the extent of these potential benefits.

CONCLUSION

Encountered-type haptic devices (EHDs) can enable dynamic and versatile physical interaction in VR. However, kinematic and speed constraints can contribute to spatial discrepancies between the virtual content and physical device. We have proposed REACH+, a framework for extending the reachability of EHDs through dynamic reach redirection, thereby improving their perceived performance. We synthesize and build upon existing work in encountered-type haptics and visuo-haptic illusions to predict user action and pre-emptively correct for expected spatial discrepancies between the virtual content and EHD. The results of our evaluation with a desktop mobile robot EHD suggest that REACH+ can effectively improve the performance of task-challenged EHDs, yielding a >50% increase in on-time arrival rate and an increased sense of physical realism for a 20 cm/s device. Moreover, when redirection was limited to 3 cm, comparable improvements were observed without redirection becoming significantly noticeable.

ACKNOWLEDGEMENTS

We would like to thank Elyse Chase and members of the Shape Lab for piloting this work and providing valuable feedback. We also thank all the volunteers who participated in our study.

REFERENCES

[1] Parastoo Abtahi and Sean Follmer. 2018. Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 150.

[2] Parastoo Abtahi, Benoit Landry, Jackie Junrui Yang, Marco Pavone, Sean Follmer, and James A Landay. 2019. Beyond The Force: Using Quadcopters to Appropriate Objects and the Environment for Haptics in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 359.

[3] Bruno Araujo, Ricardo Jota, Varun Perumal, Jia Xian Yao, Karan Singh, and Daniel Wigdor. 2016. Snake Charmer: Physically Enabling Virtual Objects. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 218–226.

[4] Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D Wilson. 2016. Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proceedings of the 2016 chi conference on human factors in computing systems*. ACM, 1968–1979.

[5] Yuki Ban, Takashi Kajinami, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2012a. Modifying an identified angle of edged shapes using pseudo-haptic effects. In *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 25–36.

[6] Yuki Ban, Takashi Kajinami, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2012b. Modifying an identified curved surface shape using pseudo-haptic effect. In *2012 IEEE Haptics Symposium (HAPTICS)*. IEEE, 211–216.

[7] Yuki Ban, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2014. Displaying shapes with various types of surfaces using visuo-haptic interaction. In *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology*. ACM, 191–196.

[8] Gordon Binsted, Romeo Chua, Werner Helsen, and Digby Elliott. 2001. Eye–hand coordination in goal-directed aiming. *Human movement science* 20, 4-5 (2001), 563–585.

[9] Mattias Bratt, Christian Smith, and Henrik I Christensen. 2007. Minimum jerk based prediction of user actions for a ball catching task. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2710–2716.

[10] Lung-Pan Cheng, Eyal Ofek, Christian Holz, Hrvoje Benko, and Andrew D Wilson. 2017. Sparse haptic proxy: Touch feedback in virtual environments using a general passive prop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3718–3728.

[11] Kiran Dandekar, Balasundar I Raju, and Mandayam A Srinivasan. 2003. 3-D finite-element models of human and monkey fingertips to investigate the mechanics of tactile sense. *J. Biomech. Eng.* 125, 5 (2003), 682–691.

[12] Massimiliano Di Luca, Benjamin Knörlein, Marc O Ernst, and M Harders. 2011. Effects of visual–haptic asynchronies and loading–unloading movements on compliance perception. *Brain research bulletin* 85, 5 (2011), 245–259.

[13] William R Ferrell. 1966. Delayed force feedback. *Human factors* 8, 5 (1966), 449–455.

[14] Tamar Flash and Neville Hogan. 1985. The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of neuroscience* 5, 7 (1985), 1688–1703.

[15] James J Gibson. 1933. Adaptation, after-effect and contrast in the perception of curved lines. *Journal of experimental psychology* 16, 1 (1933), 1.

[16] Eric J Gonzalez, Parastoo Abtahi, and Sean Follmer. 2019. Evaluating the Minimum Jerk Motion Model for Redirected Reach in Virtual Reality. In *The Adjunct Publication of the 32nd Annual ACM Symposium on User Interface Software and Technology*. ACM, 4–6.

[17] Mar Gonzalez-Franco and Jaron Lanier. 2017. Model of Illusions and Virtual Reality. *Frontiers in Psychology* 8 (2017), 1125. DOI: http://dx.doi.org/10.3389/fpsyg.2017.01125

[18] Dustin T Han, Mohamed Suhail, and Eric D Ragan. 2018. Evaluating remapped physical reach for hand interactions with passive haptics in virtual reality. *IEEE transactions on visualization and computer graphics* 24, 4 (2018), 1467–1476.

[19] Zhenyi He, Fengyuan Zhu, Aaron Gaudette, and Ken Perlin. 2017. Robotic haptic proxies for collaborative virtual reality. *arXiv preprint arXiv:1701.08879* (2017).

[20] Koichi Hirota and Michitaka Hirose. 1995. Simulation and presentation of curved surface in virtual reality environment through surface display. In *Proceedings Virtual Reality Annual International Symposium'95*. IEEE, 211–216.

[21] Matthias Hoppe, Pascal Knierim, Thomas Kosch, Markus Funk, Lauren Futami, Stefan Schneegass, Niels Henze, Albrecht Schmidt, and Tonja Machulla. 2018. VRHapticDrones: Providing Haptics in Virtual Reality through Quadcopters. In *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia*. ACM, 7–18.

[22] Hsin-Yu Huang, Chih-Wei Ning, Po-Yao Wang, Jen-Hao Cheng, and Lung-Pan Cheng. 2020. Haptic-go-round: A Surrounding Platform for Encounter-type Haptics in Virtual Reality Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–10.

[23] Brent Edward Insko, M Meehan, M Whitton, and F Brooks. 2001. *Passive haptics significantly enhances virtual environments*. Ph.D. Dissertation. University of North Carolina at Chapel Hill.

[24] Hiroo Iwata, Hiroaki Yano, Hiroyuki Fukushima, and Haruo Noma. 2005. CirculaFloor [locomotion interface]. *IEEE Computer Graphics and Applications* 25, 1 (2005), 64–67.

[25] Mitsuo Kawato. 1996. Trajectory formation in arm movements: minimization principles and procedures. *Advances in motor learning and control* (1996), 225–259.

[26] G Drew Kessler, Larry F Hodges, and Neff Walker. 1995. Evaluation of the CyberGlove as a whole-hand input device. *ACM Transactions on Computer-Human Interaction (TOCHI)* 2, 4 (1995), 263–283.

[27] Roberta L Klatzky, Susan J Lederman, and Catherine Reed. 1987. There's more to touch than meets the eye: The salience of object attributes for haptics with and without vision. *Journal of experimental psychology: general* 116, 4 (1987), 356.

[28] Luv Kohli. 2010. Redirected touching: Warping space to remap passive haptics. In *2010 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 129–130.

[29] Luv Kohli, Mary C Whitton, and Frederick P Brooks. 2012. Redirected touching: The effect of warping space on task performance. In *2012 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 105–112.

[30] Luv Kohli, Mary C Whitton, and Frederick P Brooks. 2013. Redirected Touching: Training and adaptation in warped virtual spaces. In *2013 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 79–86.

[31] Anatole Lécuyer. 2009. Simulating haptic feedback using vision: A survey of research and applications of pseudo-haptic feedback. *Presence: Teleoperators and Virtual Environments* 18, 1 (2009), 39–53.

[32] Susan J Lederman and Lynette A Jones. 2011. Tactile and haptic illusions. *IEEE Transactions on Haptics* 4, 4 (2011), 273–294.

[33] Chang-Gyu Lee, Gregory Lynn Dunn, Ian Oakley, and Jeha Ryu. 2017. Visual Guidance for a Spatial Discrepancy Problem of in Encountered-Type Haptic Display. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* (2017).

[34] Chang-Gyu Lee, Ian Oakley, Eun-Soo Kim, and Jeha Ryu. 2015. Impact of visual-haptic spatial discrepancy on targeting performance. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 46, 8 (2015), 1098–1108.

[35] Thomas H Massie, J Kenneth Salisbury, and others. 1994. The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, Vol. 55. Chicago, IL, 295–300.

[36] William A McNeely. 1993. Robotic graphics: a new approach to force feedback for virtual reality. In *Proceedings of IEEE Virtual Reality Annual International Symposium*. IEEE, 336–341.

[37] Mathew Monfort, Anqi Liu, and Brian Ziebart. 2015. Intent prediction and trajectory forecasting via predictive inverse linear-quadratic regulation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.

[38] Roberto A Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. 2017. Erg-O: ergonomic optimization of immersive virtual environments. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, 759–771.

[39] Sebastiaan FW Neggers and Harold Bekkering. 2002. Coordinated control of eye and hand movements in dynamic reaching. *Human movement science* 21, 3 (2002), 37–64.

[40] C Prablanc, JF Echallier, E Komilis, and M Jeannerod. 1979. Optimal response of eye and hand motor systems in pointing at a visual target. *Biological cybernetics* 35, 2 (1979), 113–124.

[41] Irvin Rock and Jack Victor. 1964. Vision and touch: An experimentally created conflict between the two senses. *Science* 143, 3606 (1964), 594–596.

[42] Ken Shigeta, SATO Yuji, and Yasuyoshi Yokokohji. 2007. Motion planning of encountered-type haptic device for multiple fingertips based on minimum distance point information. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*. IEEE, 188–193.

[43] Adalberto L Simeone, Eduardo Velloso, and Hans Gellersen. 2015. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 3307–3316.

[44] Alexa F Siu, Eric J Gonzalez, Shenli Yuan, Jason B Ginsberg, and Sean Follmer. 2018. Shapeshift: 2D spatial manipulation and self-actuation of tabletop shape displays for tangible and haptic interaction. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 291.

[45] Barton A Smith, Janet Ho, Wendy Ark, and Shumin Zhai. 2000. Hand eye coordination patterns in target selection. In *Proceedings of the 2000 symposium on Eye tracking research & applications*. ACM, 117–122.

[46] Christian Smith and Patric Jensfelt. 2010. A predictor for operator input for time-delayed teleoperation. *Mechatronics* 20, 7 (2010), 778–786.

[47] Ryo Suzuki, Hooman Hedayati, Clement Zheng, James L Bohn, Daniel Szafir, Ellen Yi-Luen Do, Mark D Gross, and Daniel Leithinger. 2020. RoomShift: Room-scale Dynamic Haptics for VR with Furniture-moving Swarm Robots. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.

[48] Susumu Tachi. 1994. A construction method of virtual haptics space. In *Proc. of the ICAT'94 (4th International Conference on Artificial Reality and Tele-Existence)*. 131–138.

[49] Ingrid MLC Vogels. 2004. Detection of temporal delays in visual-haptic interfaces. *Human Factors* 46, 1 (2004), 118–134.

[50] Emanuel Vonach, Clemens Gatterer, and Hannes Kaufmann. 2017. VRRobot: Robot actuated props in an infinite virtual environment. In *2017 IEEE Virtual Reality (VR)*. IEEE, 74–83.

[51] David Wang, Kevin Tuer, Mauro Rossi, Liya Ni, and Joseph Shu. 2003. The effect of time delays on tele-haptics. In *The 2nd IEEE Internatioal Workshop on Haptic, Audio and Visual Environments and Their Applications, 2003. HAVE 2003. Proceedings.* IEEE, 7–12.

[52] Yuntao Wang, Zichao Chen, Hanchuan Li, Zhengyi Cao, Huiyi Luo, Tengxiang Zhang, Ke Ou, John Raiti, Chun Yu, Shwetak Patel, and others. 2020. MoveVR: Enabling Multiform Force Feedback in Virtual Reality using Household Cleaning Robot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[53] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. 2018. Haptic revolver: Touch, shear, texture, and shape rendering on a reconfigurable virtual reality controller. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.

[54] Motohiro Yafune and Yasuyoshi Yokokohji. 2011. Haptically rendering different switches arranged on a virtual control panel by using an encountered-type haptic device. In *2011 IEEE World Haptics Conference*. IEEE, 551–556.

[55] Yasuyoshi Yokokohji, Ralph L Hollis, and Takeo Kanade. 1999. WYSIWYF display: A visual/haptic interface to virtual environment. *Presence* 8, 4 (1999), 412–434.

[56] Yasuyoshi Yokokohji, Junji Kinoshita, and Tsuneo Yoshikawa. 2001. Path planning for encountered-type haptic devices that render multiple objects in 3d space. In *Proceedings IEEE Virtual Reality 2001*. IEEE, 271–278.

[57] Yasuyoshi Yokokohji, Nobuhiko Muramori, Yuji Sato, and Tsuneo Yoshikawa. 2005. Designing an encountered-type haptic display for multiple fingertip contacts based on the observation of human grasping behaviors. *The International Journal of Robotics Research* 24, 9 (2005), 717–729.

[58] Andrea Maria Zanchettin and Paolo Rocco. 2017. Probabilistic inference of human arm reaching target for effective human-robot collaboration. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6595–6600.

[59] André Zenner and Antonio Krüqer. 2019. Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 47–55.

[60] Yiwei Zhao and Sean Follmer. 2018. A Functional Optimization Based Approach for Continuous 3D Retargeted Touch of Arbitrary, Complex Boundaries in Haptic Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 544.

[61] Yiwei Zhao, Lawrence H Kim, Ye Wang, Mathieu Le Goc, and Sean Follmer. 2017. Robotic assembly of haptic proxy objects for tangible interaction and virtual reality. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces*. 82–91.