

RESEARCH ARTICLE

Toward Full-Immersive Multiuser Virtual Reality With Redirected Walking

THOMAS VAN ONSEM^{1,2}, JAKOB STRUYE^{1,2}, (Graduate Student Member, IEEE),
XAVIER COSTA PEREZ^{3,4,5}, JEROEN FAMAËY^{1,2}, AND FILIP LEMIC^{1,3}

¹Internet and Data Laboratory (IDLab), University of Antwerp, 2000 Antwerpen, Belgium

²imec, 2000 Antwerpen, Belgium

³i2Cat Foundation, 08034 Barcelona, Spain

⁴NEC Laboratories Europe GmbH, 69115 Heidelberg, Germany

⁵ICREA, 08034 Barcelona, Spain

Corresponding author: Filip Lemic (filip.lemic@i2cat.net)


This work was supported in part by the MCIN/AEI/10.13039/501100011033/FEDER/EU HoloMit 2.0 Project PID2021—126551OB—C21; in part by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union—NextGeneration EU, in the framework of the Recovery Plan, Transformation and Resilience (PRTR) (Call UNICO I+D 5G 2021, under Grant TSI-063000-2021-6-Open6G Joint Open 6G Communications and Sensing); and in part by the Research Foundation—Flanders (FWO) under Grant G034322N and Grant 1SB0719N.

ABSTRACT We are witnessing a surge in the popularity and mass-availability of Virtual Reality (VR) technology. Toward enhancing its maturity, current research is focusing on achieving continuous full immersiveness of the VR users during their experiences in Virtual Environments (VEs), as well as on enabling multiuser experiences. In many cases, achieving the full immersiveness requires enabling the users' perception of walking naturally in unbound VEs. Redirected Walking (RDW) algorithms provide a solution to this problem by exploiting the limitations of the human spatial perception. In other words, the RDW algorithms are envisioned to unnoticeably steer the VR users for constraining them within a certain tracking space. Regardless of notable research efforts targeting RDW, there are still not many experimental evaluations of the performance of RDW algorithms, especially in multiuser settings and utilizing contemporary Head-Mounted Devices (HMDs). Toward addressing this issue, we have developed a modular framework that enables straightforward experimentation with single and multiuser RDW in different tracking spaces and for varying VEs. The capabilities of the developed framework were demonstrated by carrying out an extensive experimentation study to capture the performance and noticeability of a contemporary RDW algorithm for a varying number of users in different tracking spaces and VEs. The lessons learned during the study have been conceptualized in a form of a set of RDW design enhancements that can be generically applied to existing RDW algorithms. We show that the proposed enhancements significantly increase the overall performance of RDW and decrease the noticeability of RDW-supported experiences in VEs.

INDEX TERMS Full-immersive multiuser virtual reality, redirected walking, artificial potential fields.

I. INTRODUCTION

With the release of the HTC Vive, Oculus Rift, and PlayStation VR, Virtual Reality (VR) technology became widely accessible to the public. Taking as an argument the estimated market value of \$11.64 billion and the projected

The associate editor coordinating the review of this manuscript and approving it for publication was Jiachen Yang .

one of \$227.34 billion by 2029 [1], one can safely argue that VR technology is here to stay. To further strengthen this claim, entertainment, medicine, and education are only some industries that have taken advantage of the technology [2], [3], [4].

Currently, VR-related research is focusing on achieving the continuous full-immersiveness of the users during their Virtual Environments (VEs) [5]. A major aspect of achieving

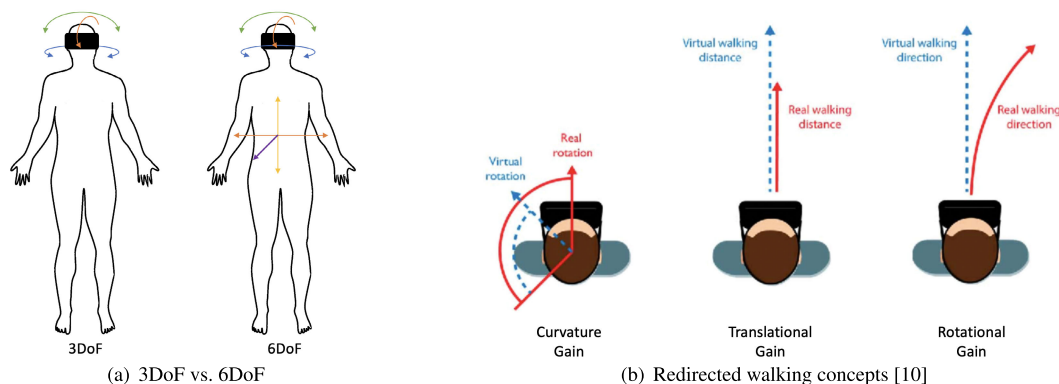


FIGURE 1. Main concepts of redirected walking enabled through 6DoF non-tethered wireless HMDs.

such immersiveness is the ability to simulate the experience of natural human movements (primarily walking) through unbound VEs. Hardware solutions such as treadmills and foot platforms have been suggested as a solution to this problem. However, they come with a high cost, possibility of motion sickness, and limited portability [6], and have, therefore, not seen widespread usage.

Another challenge limiting further research toward enhancing the immersiveness is the lack of open-source support for multiuser experiences [7], [8]. In such systems, the VR users are envisioned to jointly roam a constrained physical tracking space, while simultaneously being immersed in VEs. Intuitively, the immersion of the users in the VEs should be continuous and unbounded, while in the tracking space they should roam without a possibility of collisions with any obstacles, environmental boundaries, or other users [9].

Today, VEs are delivered to the users through tethered (i.e., wired) and standalone (i.e., wireless) HMDs. Tethered HMDs are usually wired to an external device and utilize its computational power, consequently featuring high visual quality and consistently high frame rate. The main downside is a result of tethering to an external device, which confines the users in their movement capabilities and often increases the noticeability [11]. Such HMDs were originally envisioned to support three Degrees of Freedom (3DoF)-like mobility of the users, where only their rotational movements are tracked, as shown in Figure 1a. Subsequently, tethered HMDs with six Degrees of Freedom (6DoF) visual support were proposed, featuring the ability to provide the users with the perception of their translational and rotational movements being tracked in VEs. However, the wiring to an external device limited the tracking of translational movements to the length of a cable, constraining the overall immersiveness of the user experience. In addition, the cable generally represents a tripping hazard and is, therefore, unwanted in full-immersive VR setups. This issue has been addressed through HMDs tethered to back-strapped laptops or similar portable devices, with the resulting systems approximating standalone wireless setups. Although such setups can enable the sensation of unrestricted walking for the VR users [12], they are also rather bulky and feature undesired wiring. Alternatively and more recently,

standalone HMDs with built-in support for delivering VEs have emerged. Such HMDs can deliver full-immersive VR experiences, are not bulky, and feature the ability to track 6DoF-like mobility (i.e., both translational and rotational) of the users without constraints. Therefore, the developments in this work are targeting this latest type of HMDs.

6DoF non-tethered wireless HMDs are envisioned to utilize Redirected Walking (RDW) algorithms to imperceptibly steer the users within constrained physical spaces [10], [13] (more details in Section II-A). To maintain the full-immersiveness in both single- and multiuser VR systems, there is a need for proper dimensioning of a tracking space, primarily accounting for the trade-offs between the physical sizes of the tracking space and the number of collocated users on the one hand, and the immersiveness, noticeability, and the overall user experience on the other [9], [14]. Toward addressing this issue, we have developed a modular framework for experimentation with RDW that can support different RDW algorithms and VEs, and account for varying tracking space sizes and number of collocated users. An extensive experimental study was carried out using the framework, yielding several lessons that can be generically applied to existing RDW algorithms. Our improvements include introducing latency-awareness in the design of RDW, accounting for and mitigating drifts and head rotations, and compensating for oscillations and border bouncing.

To evaluate the proposed generic improvements, we have introduced them to a well-known RDW algorithm, namely Artificial Potential Fields (APF) [7]. Using the framework, we have evaluated the performance enhancements of the instantiated RDW algorithm due to the introduction of our generic improvements. Additionally, we assessed the practicality of a large number of configurations distinguished based on the size of the tracking space, type of VEs, and number of collocated users. Our results indicate that the utilization of the proposed generic improvements yields both objective (i.e., decrease in the number of noticeable “resets”) and subjective (i.e., questionnaire-based assessment using Likert scale) benefits for the VR users. At the same time, the study was carried out to demonstrate the utility and modularity of the developed experimentation framework. The framework is

provided to the community to enable testing of other RDW approaches and configurations (see Section IV-B).

The paper is structured as follows. In Section II, we provide an overview of the related works. Section III presents the generic improvements intuitively applicable to existing RDW algorithms. In Section IV, we present the design of the experimentation framework, as well as the experimental setup for assessing the performance of the considered RDW algorithm. Section V describes the results of the performed experiments and analyses their implications. Finally, the paper is concluded in Section VI.

II. RELATED WORK

A. REDIRECTED WALKING FUNDAMENTALS

A straightforward approach for enabling full-immersiveness in VEs is through one-to-one mapping of the users' physical movements to the ones in the VEs. Intuitively, this approach limits the users' mobility in the VEs to the usually rather constrained tracking spaces, consequently negatively affecting the immersiveness. To address this issue, researchers have proposed various virtual locomotion techniques for supporting walking over large and often infinite VEs, while being constrained within relatively small tracking spaces.

The current state-of-the-art approaches support traversing the VEs by exploiting walk-like gestures (e.g., walking-in-place [15]), which was shown to provide the realistic impression of walking [16]. The reason for that comes from perception psychology research that indicates that the human vision usually dominates proprioception (i.e., the sense of perceiving location, movement, and action of the human body or its parts) and vestibular sensation (i.e., the movement, gravity, and sense of balance) if they are in disagreement [17]. In other words, humans are rather good in estimating momentary direction of self-motion and much less so in perceiving their paths of travel [18].

In the context of VEs, this implies that the users unknowingly compensate for small inconsistencies during walking in case of disagreement between the visual, proprioceptive, and vestibular sensations. This allows for their imperceptible redirections utilizing visual cues delivered through VEs, which can be considered as the defining feature of RDW. The redirections are accomplished through (cf., Figure 1b): i) curvature gains, i.e., rotations of the virtual scene, ii) translational gains, i.e., modifying linear movements in a VE, resulting in changes the user's physically traveled distances, and iii) rotational gains, i.e., introducing additional rotations to the already rotating user. Details on the mathematical formulations of these gains, their perception thresholds, as well as on their experimental derivations, are presented in [16] and [19]. It is also worth pointing out that the authors in [20] show that the VR users modify their speeds when significant translational gains are applied, regardless of whether the gains are perceptible. The authors followed by suggesting the need for a large-scale study to investigate the effects of combining different types of gains on the detection thresholds. We argue that the envisioned framework

can support and simplify such and similar experimentation requirements.

B. REDIRECTED WALKING ALGORITHMS

In the following, our intent is not to provide an exhaustive overview of RDW, but to make a case for utilizing APF as a contemporary algorithm for demonstrating the capabilities of the framework and the benefits of the generic improvements. An exhaustive overview of RDW techniques and algorithms can be found in [21]. The RDW decisions are based on context data produced by the HMDs, usually utilizing a combination of the following context information instances: the user's current and past physical locations, steering directions, and walking trajectories in the VEs [22]. Specifically, scripted RDW algorithms require full knowledge of the virtual paths of the users, allowing for redirection optimization along a pre-defined (i.e., scripted) pathway. Moreover, predictive RDW approaches utilize users' current locations in the tracking space and their potential walking trajectories in the VEs for selecting the steering directions. Finally, reactive RDW algorithms consider users' current and past physical locations and steering directions in the VEs for adjusting the redirection. In this work, we utilize a contemporary algorithm from the category of reactive RDW algorithms.

While the majority of RDW algorithms utilize the same re-steering techniques (cf., Figure 1b), they differ on the high-level strategy for selecting the steering directions (also called "controller" in some works from the literature, e.g., [10]). This results in variations in the minimum sizes of the tracking spaces required for the algorithms to operate well, making some of them more effective than the others. The three "steer-to" algorithms, proposed by Razaque [13], aim at steering the users toward a given target in the tracking space. The three approaches consist of: i) Steer-To-Center (S2C) where the algorithm attempts at steering the user toward the center of the tracking space, as this area is assumed to be furthest away from the walls and obstacles, hence least likely to incur a future collision; ii) Steer-To-Orbit (S2O) where the algorithm attempts at steering the user in an orbit around the center of the tracking space; iii) Steer-To-Multiple-Targets (S2MT) where the algorithm attempts to steer the user toward a number of targets sequentially.

A recent research effort by Strauss et al. [23] applies machine learning, specifically reinforcement learning, to the RDW problem, naming the approach the Steering algorithm learned via Reinforcement Learning (SRL). SRL consists of a learning agent interacting with the tracking space, where the agent is envisioned to receive both a reward and next state through the interactions with the tracking space. The system is applied with the goal of maximizing the virtual distance traveled between consecutive collisions, while unnoticeably adjusting the path of the user.

Artificial Potential Fields (APF) [7] is another RDW approach in which the user is repelled away from all obstacles in the tracking space and "pushed" into an open space. The algorithm accomplishes this by generating a force vector

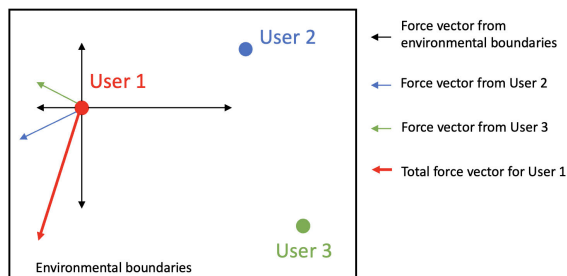


FIGURE 2. Determining the steering direction in APF.

that is directed away from the obstacles and has the length inversely proportional to its distance from each obstacle. The steering direction is determined by summing the individual force vectors for obtaining a total force vector that points toward an open space, as depicted in Figure 2.

The “steer-to” algorithms can be considered as a basic form of RDW, primarily designed to show the viability of the technique. Originally, they were envisioned to utilize solely the user’s physical location and steering direction as context information instances from the HMDs, as well as targets in the tracking space for making steering decisions. Moreover, these algorithms assumed collision-free environments and were not suited for tracking spaces without obvious centers (especially concave ones [13]) and obstacles [24]. Over the years, additions to the algorithms were suggested, mainly to improve the selection of the overall redirection angles resulting in the algorithms’ performance enhancements in terms of the overall number of perceivable resets and the average distances between them, as well as to provide multiuser support in tracking spaces without obvious centers [25], [26], [27]. Regarding the SRL algorithm, although it is a promising approach in redirecting the user, it is experimentally untested as it requires large amounts of data for training, and multi-user steering is not supported.

In APF [7], [28], the force vector-based approach significantly simplifies the computational complexity compared to other algorithms as there is no need to predict collisions or make explicit decisions for preventing them. Because the users constantly get repelled from obstacles and tracking space boundaries, the algorithm can utilize the full tracking area and can easily be used in tracking spaces without obvious centers [28], as long as there are no “bottlenecks” in the tracking space. For example, if the tracking space features different rooms connected by doors, the APF algorithm would experience issues in guiding the users through the doors, since the force field would be too high as a result of the bottleneck created by the adjacent walls. Additionally, APF can operate in the tracking spaces that feature obstacles, both stationary and moving, such as other users [28]. With live tests, [7] demonstrated that the APF algorithm improved upon S2C (for single user) in increasing the average distance between resets by 86% and decreasing the number of resets by 64%. This, of course, majorly improves the immersiveness, decreases nausea, and increases the overall experience of the users.

Based on the above discussion, we have selected APF as the RDW algorithm for our wireless multiuser setting with different tracking space sizes and VEs.

For completeness, it is also worth emphasizing that, even with an RDW algorithm such as APF in place, a collision is sometimes imminent and cannot be avoided with imperceptible re-steering. In such cases, the resetting algorithm is triggered and the user gets stopped in its tracks. Artificial Potential Fields - Resetting (APF-R) is a resetting algorithm proposed in [7] that uses the total force vector, calculated similarly to APF, to determine the angle the user has to physically turn toward. To keep the break in immersion as small as possible, the previously mentioned curvature gain method is used to perform a 2:1 turn. In the 2:1 turn, the user’s rotational speed gets increased, resulting in the user always turning 360° in the VE while in reality turning a smaller computed angle. Thus, the user physically gets directed toward a safe zone while virtually having the exact same direction as before [26]. In this study, we will utilize APF-R for resetting in the situations of an imminent collision.

C. EXPERIMENTATION WITH REDIRECTED WALKING IN FULL-IMMERSIVE MULTIUSER VR

Steinicke et al. captured the RDW noticeability threshold by empirically determining the noticeability thresholds for rotational, translational, and curvature gains [14], [16], [19]. In determining the thresholds, they utilized a 3DoF HMD tethered to a back-fitted laptop. The perception threshold of the rotational gain was captured by instructing the subjects to rotate their heads by 90° and letting them decide whether the physical rotation was larger than the visually simulated one. For deriving the translational gain threshold, the subjects were instructed to virtually walk along a straight line, followed by determining if the physically and virtually traveled distances were comparable. A similar procedure was followed for determining the curvature gain threshold, with the difference that the subjects were steered along curved physical paths. Similar experiments were also reported in [7], [26], and [29].

These experiments were largely carried out in a 25×45 m room at the Miami University called “The Hive” [30], which was adapted to support 6DoF-like experimentation with HMDs originally designed for supporting 3DoF-like mobility. Additionally, infrared sensors (from InterSense and WorldViz LLC) were utilized for capturing the user’s rotational and translational changes. This was done due to the inability of the HMDs to provide such data. Moreover, WiFi was utilized to send this contextual data to the user’s back-mounted laptop, which was then used by an RDW algorithm generating steering recommendations.

In contrast to the above-discussed efforts utilizing either setups with HMDs tethered to an external device (thus, with translational mobility constrained to the length of a cable) or the ones with HMDs tethered to back-strapped laptop (thus, featuring bulkiness and limited portability), in our

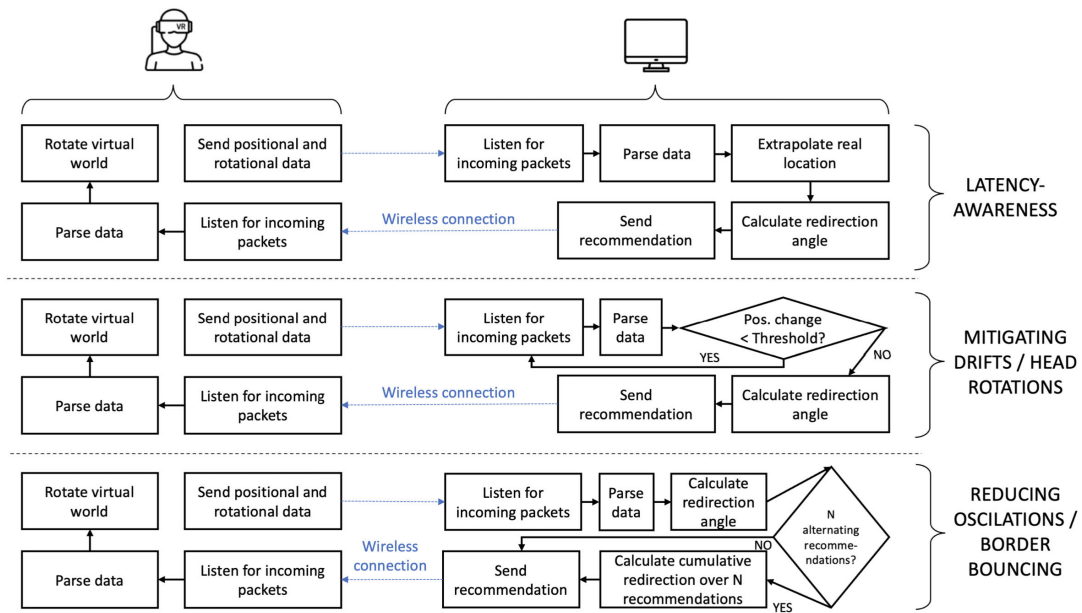


FIGURE 3. High-level designs of the RDW improvements.

experiments we utilize the Meta Quest 2. The Meta Quest 2 is a standalone HMD supporting 6DoF with high refresh rate and increased processing power. As such, it improves the immersiveness by providing better visual quality of VEs and removing the need for a back-strapped laptop. There is also no need for environmental sensors capturing positional and rotational changes as the HMD is able to produce contextual data and communicate it wirelessly to a central node using on-board sensors and inside-out tracking. Thus, we argue that our experiments feature a higher degree of realism than the ones from the literature along the lines of technology maturity and enhanced immersiveness. Finally, in contrast to the majority of the existing efforts focusing on single-user systems, we report on the RDW performance and user experience for multiuser scenarios.

III. DESIGN OF REDIRECTED WALKING ENHANCEMENTS

In this section, we discuss the RDW design enhancements that can be generically applied to existing RDW algorithms (cf., Figure 3). We consider the enhancements as generically applicable as they envision modifications in either raw input data (e.g., extrapolating a more precise user’s location) or in the post-processing of the algorithm’s recommendations (e.g., mitigation of small alternating positive and negative recommendations), while an employed RDW algorithm can be considered as an unmodified black-box approach.

A. LATENCY-AWARE RDW

Due to the fact that wireless HMDs were utilized in our experimental study, their communication with a central entity to determine steering recommendations incurred certain delays. The end-to-end delays in such a scenario intuitively depend on both the runtime performance of a RDW algorithm and the

underlying wireless infrastructure. Due to the fact that each instance of context information is produced at the HMD and sent wirelessly to the central entity, the received contextual information is outdated by the time it is utilized by the central entity for providing steering recommendations.

In our improved design, we account for the delays of data transmission from the HMD to the central entity, as shown in Figure 3. The experimental procedure for capturing this delay is discussed in Section IV-A. Based on this communication delay, we extrapolate the location of the user at the time of calculation of the steering recommendations. This extrapolation is based on the location reported by the HMD, previous locations of the user, an estimate of the delay, and the user’s rotation and speed.

B. MITIGATING DRIFTS AND HEAD ROTATIONS

It is established that accelerometers and gyroscopes in HMDs are imperfect and, therefore, provide noisy data [31]. This noise can propagate to the system used to locate an HMD, causing one or more of its components to drift. Most of the HMDs minimize this issue by using their on-device camera to locate fixed reference points in the tracking space. However, drifts cannot be completely removed, hence in our improved design of this aspect of redirected walking we aim at accounting for them. In other words, changes in positional data received by the server are discarded if they are smaller than the predefined drifting threshold, as shown in Figure 3. Determining this drifting threshold is envisioned as a device-specific procedure that involves letting the HMD to lay still and level, capturing the changes in the position, followed by establishing the threshold to recognize drifting.

Moreover, when referring to the position of the user, one actually refers to the position of the HMD mounted on the

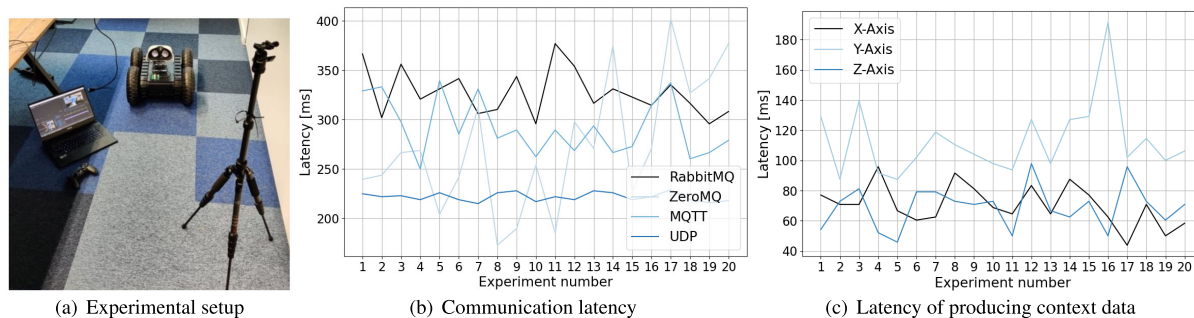


FIGURE 4. Capturing the end-to-end latency for instantiating latency-aware RDW.

user's head. Consequently, a change in rotation of the head will often result in small changes in position of the HMD, and consequently on the RDW algorithm reacting to those changes. This is undesirable and should be ignored since the user does not actually move. Mitigating this change is performed in the same fashion as for the positional drift mitigation described above. Specifically, when the positional data is received, it is matched with a predetermined threshold and either discarded or passed to the RDW algorithm.

C. REDUCING OSCILLATIONS AND BORDER BOUNCES

In some situations, the RDW algorithm recommended small alternating positive and negative rotations that resulted in an oscillating behavior of a VE and consequently of the user's movement trajectory, as also reported in [25]. To reduce the effect of this behavior, a history of sent rotations was kept and compared with real-time recommended rotations. When a series of small alternating negative and positive rotations were recommended, the alternating data is envisioned to be averaged to remove the oscillations, followed by providing a rotation recommendation (cf., Figure 3).

The same preliminary tests also exposed a specific and, to the best of our knowledge, not yet documented problem in multiuser scenarios, where a user would be directed into a reset loop. In other words, when two users are located close to both an environmental boundary and each other, the algorithm would reset the users toward an empty space and immediately redirect them back towards the wall, landing in a bouncing pattern with the walls. This resulted in a very nauseating experience with a large number of observable resets. To tackle this issue, a cool-down timer was implemented that is envisioned to be initiated after triggering a reset. This way, the user would be able to walk away from the wall and towards the empty space for a short period, thus maximizing the whole space and reducing the number of resets.

IV. EXPERIMENTATION SETUP AND METHODOLOGY


A. CONFIGURATION OF RDW ENHANCEMENTS

Four popular protocols were tested as a means for communication between the HMD and the central entity. Three of those protocols conform to the Publish/Subscribe (Pub/Sub) pattern that provides a structured exchange of messages between publishers and subscribers. This pattern relies on

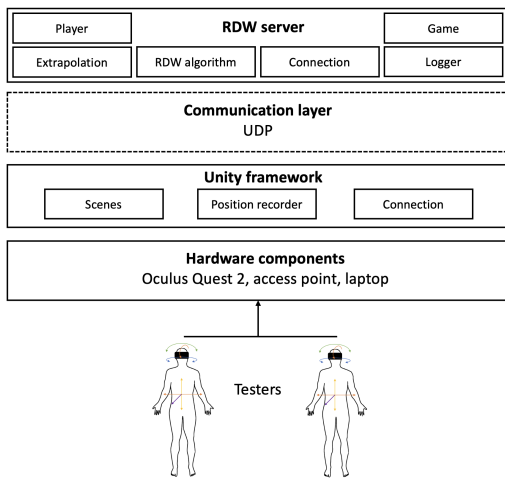
a message broker to relay messages from a publisher to all its subscribers. The following protocols conforming to the Pub/Sub pattern were selected: i) PahoMQTT: a lightweight Pub/Sub messaging protocol regarded as a Message Queuing Telemetry Transport (MQTT) client implementation, mainly designed for IoT applications in low-bandwidth and unstable network environments. ii) RabbitMQ: an open source message broker that focuses on reliability, flexible routing, and highly available message queues. iii) Zero-MQ: an asynchronous messaging library that does not require a broker and uses a message queue together with sockets that can represent a many-to-many connection between endpoints. In addition, User Datagram Protocol (UDP) as a connection-less transport layer protocol not conforming to the Pub/sub pattern was selected. UDP uses a simple transmission method without reliability guarantees and with only optional checksum-based data integrity support.

We have built a C++ program that runs natively on the HMD using the VrApi provided by the Meta Quest 2. The program was continuously sending the current coordinates of the HMD to the central entity using the previously mentioned protocols. When a movement was detected, the screen turned green to enable us to capture the communication delay. The HMD was mounted on a Rover Robotics 4WD to easily simulate and detect harsh movements. The server was set up on a laptop (MSI GS66) with a 120 Hz display as a Python 3.9 program that constantly checked and displayed received packets. To film the sequence of events, we opted to record the experiment at 480 fps with the camera of a OnePlus 6T smartphone, featuring an error of at most 2083 μ s. Given that there was a need to simultaneously film the HMD display turning green and the laptop showing a received packet, the camera was mounted on a tripod, as seen in Figure 4a.

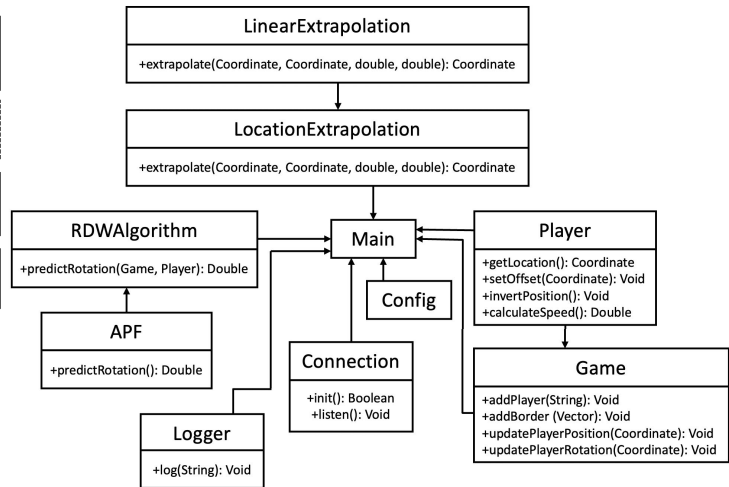
The experiment consisted of the following actions: i) start server, ii) connect Meta Quest 2 wirelessly to server, iii) start recording, iv) drive robot forward at full acceleration, v) stop movement and recording once server prints received packets. This experiment was repeated 20 times to provide enough data and possibly detect a difference in position polling. Measurements were obtained by counting the number of frames between the HMD screen turning green and the server showing the incoming packets. The delay in milliseconds



(a) Implementation process diagram



(b) Software/hardware architecture representation



(c) Class diagram

FIGURE 5. Overview of the developed framework for experimentation with RDW in scenarios with multiple collocated users.

could then be calculated from this data as $Delay = Frames \times 1000/480$. For UDP, the results showed an average delay of 214.90 ± 7.29 ms, as can be seen in Figure 4b. The error margin was calculated by taking half of the time between frames of all the used devices. For the used HMD, this was 8.33 ms, the laptop has 4.17 ms between frames and the camera 2.08 ms. The Pub/Sub protocols were tested with their default settings and produced worse results compared to UDP. This is because the Pub/Sub messaging is generally deployed at the application layer on top of the transport layer protocols such as Transmission Control Protocol (TCP) or UDP, resulting in an increased delay compared to UDP (cf., Figure 4). Nonetheless, we report on these results because such distributed messaging is an intuitive candidate for supporting communication in full-immersive multiuser VR settings.

To provide a full picture of the delay between the user movement and the central entity getting that information, we measured the number of frames that passed between the robot moving and the HMD detecting that movement. Because it is hard to exactly determine at which frame the robot started moving, the results served as an indication and had a large error margin. The measurements gave in comparable average delays of 50.23 ± 20.82 ms and 48.14 ± 20.82 ms for respectively the Z (forwards and backwards) and X-axes (left and right), and a larger one of 92.55 ± 20.82 ms for the Y-axis (up and down), as shown in Figure 4c. There

are two hypotheses that can be put forth for the comparably higher average delay observed along the Y-axis. We believe it is either caused by the HMD trying to save power by infrequently polling a lesser used axis or a wrongly calibrated sensor. Nonetheless, it was not an important element for our experiments as it was not utilized as an input for the considered RDW algorithms.

Given that an increasing delay in communication between the HMD and central entity decreases the performance of RDW algorithms as the utilized location becomes obsolete, we have decided to utilize UDP as a communication means in the remainder of this work. The results presented above provide an indication on how long it took to produce a context instance, send it to the central entity, and receive a steering recommendation: $Latency = (50.23 + 48.14)/2 + (2 \times 214.90)$. Hence, in the remainder of the paper we will utilize 478.98 ms as the latency value in latency-aware RDW.

B. EXPERIMENTATION FRAMEWORK

To the best of our knowledge, there are currently no open-source frameworks supporting experimentation with RDW in scenarios with multiple collocated users and 6DoF-enabled wireless HMDs. Consequently, following the implementation process diagram in Figure 5a we have developed a modular C++ framework¹ that features three main

¹ https://github.com/ThomasVanOnsem/VR_RW_Noticeability_Public

software components: RDW server, UDP-supported communication, and Unity-based HMDs (cf., Figure 5b). As depicted in Figure 5c, the core of the developed framework enables communication with the HMDs through a listening function. Additionally, it is envisioned to receive contextual data, parse it, and convert into data structures for managing the different users. The aim of the separated RDW algorithm module is to abide to the modular design of the framework, so that the introduction of a new RDW algorithm does not require changes in the framework itself.

As the HMD only requires rotation recommendations from the RDW algorithm, a single parent function named *predictRotation()* was created. The function takes as an input the user's physical location, its expected virtual movement trajectory, and the locations of the obstacles and potentially other users. Based on that and accounting for the noticeability constraints for rotational and longitudinal movements of the users, it provides a recommended rotation. Moreover, it is also used for triggering an observable reset in case of an imminent collision. The experimenter can override that function for introducing a new algorithm to the framework.

As argued before, we consider the extrapolation of the user's actual position based on the obsolete location produced by the wireless HMD as a potential improvement of an RDW system in general. The *LocationExtrapolation* module, implemented as linear extrapolation to demonstrate the feasibility of the approach, was developed in a fully modular way for straightforward replacement based on the needs of an experiment. Moreover, the *Tools* module consists of the system tools to log the received and produced data, and enable the experimenter to offset the users' position or invert their orientation. Finally, the device-specific drift mitigation thresholds have been established by letting the two utilized HMDs to lay still and measuring the variations in the reported positional data for one minute.

Because the RDW computation was fully performed at the central entity, we have made only limited modifications to the HMD side. Specifically, the VEs were built using Unity, which meant that the built-in C socket module was utilized for sending and receiving the required data via UDP. An additional Unity functionality also enabled the rotation of the VEs based on the received recommendations. These two functionalities are offered as an optional Unity package that can be imported into new projects for the purposes of reproducibility and further research.

C. EXPERIMENTAL SCENARIOS AND SETUP

The experimental scenarios were distinguished based on the size of the tracking space. Specifically, we considered the environment sizes of 5×5 , 10×10 and 15×15 m. This was done with the aim of roughly determining the minimum environmental size needed for maintaining the unnoticeability of RDW, while also abiding to the practical limitations of the future tracking spaces (e.g., their potential deployment in homes and residences with size constraints). Ideally, RDW would work in practically feasible room sizes for enabling

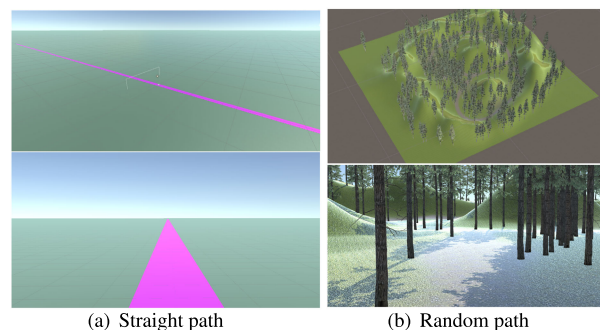


FIGURE 6. Experiences in unbound VEs.

mass-utilization of such systems. Large tracking spaces like “The Hive” [30] are not abundantly available, consequently limiting the practicality of RDW.

To accommodate different environmental sizes, we utilized a 106×65 m outdoor field near the Middelheim campus of the University of Antwerp. The environmental boundaries were precisely measured and marked with traffic cones to allow the observers (who were ensuring the safety of the testers and documenting the system behavior) to detect that the users collided with the boundaries of the tracking space without the algorithm resetting them. Moreover, although we did not aim at exhaustively capturing the performance difference in Simultaneous Localization and Mapping (SLAM)-based tracking provided by the HMD in outdoor and indoor environments, our tests of the users walking along fixed indoor and outdoor paths indicated no significant disagreement, as long as the HMD was not exposed to direct sunlight.

During the execution of the experiments, the server that was running the RDW algorithm was a MSI GS66 laptop with an Intel i7 processor, 16 GB of RAM and WiFi 6, running Windows 10. Moreover, the utilized HMDs were Meta Quest 2 with Qualcomm Snapdragon XR, 120 Hz refresh rate and 6 GB of RAM, running Android. Finally, an Access Point (AP) (i.e., a mobile hotspot using a Samsung Galaxy S8) was used for providing wireless connectivity between the server and the HMDs.

We designed two experiences in unbound VEs utilizing Unity (v.2020/03) to test the noticeability of the considered RDW approaches, as shown in Figure 6. In multiuser VR setups, three different types of user coexistence can be distinguished: i) the users sharing solely the tracking space, ii) the users sharing only the VE, iii) the users sharing both the tracking space and VE. In this study, both of the designed experiences abide to the first category. In the “straight path” experience, the users were instructed to follow a straight path during the full duration of the VE. This was considered as the worst case scenario given that the RDW algorithm was intuitively expected to have the most difficulties to unnoticeably redirect the users. In the “random path” experience, the users walked in an unbound VE and were encouraged to follow a randomly curved path. Hence, the curvature introduced by the RDW algorithm was expected to be less noticeable. Conceptually, the experiments consisted of the users walking

TABLE 1. Main experimentation parameters.

Parameter	Value
Experimental setup	
Number of testers	9
Experiment duration	300 [s]
Redirected walking (APF)	
Exponential fall-off due to other users	1.4
Arc radius or redirected walking	7.5 [m]
Maximum rotational rate	15 [°]
Velocity threshold	0.1 [m/s]
Design enhancements	
Delay in latency-aware RDW	478.98 [ms]
Cool-down period	1 [s]
Drifting threshold	5 [cm]

in the unbound VEs while being confined to the restricted tracking space. The positional data of the users was sent from the HMD to the server, where the RDW algorithm steered the users inside the confined physical environment for collision avoidance. Using this setup, we captured the performance of two APF instances (i.e., with and without the proposed improvements). This was done by simply overriding the *predictRotation()* function in the *RDWAlgorithm* module with different APF instances and introducing corresponding communication delays in the *LinearExtrapolation* module (i.e., no delay and delay of 478.98 ms for the instances without and with improvements, respectively).

We were assisted by nine testers with different professional backgrounds and levels of VR experience. Two of the testers are women, one tester has background in information technology, while the others practice law, pharmacology, economy, chemistry, or medicine. Their ages range from 21 to 52 years, with seven of them being younger than 30, and two older than 50. Six testers expressed no knowledge and prior experience in VR, while only one of them suggested extensive prior experience.

At the start of each experiment, the testers were instructed to simply follow a predefined path and enjoy the experience. They were also informed that a reset could happen due to the recommendations from the RDW algorithm. In case of the reset, the users would see a stop sign followed by the world rotating and guiding them in a recommended direction. Since the results regarding noticeability could decline after several immersions in the VEs, each tester was involved in at most four experiments (with an experiment being defined as one up to 5 minutes long immersion in a VE), in which the experiences were ordered from the expected least to the most noticeable one. The duration of each experiment was set to 5 minutes because, as there are no distractions and interactions in the VEs, the users would lose interest afterwards. The main parameters pertaining to the experimentation itself, the utilized RDW algorithm, and the proposed design enhancements are summarized in Table 1.

D. PERFORMANCE METRICS

1) ALGORITHM PERFORMANCE

Several metrics have been utilized to measure the performance of RDW, as highlighted in [7], [9], [19], [29], and [32].

Based on these efforts, we utilize the number of resets to capture how well an algorithm utilizes the available physical space. Furthermore, the average distance between resets (in m) is used to capture how long the users can traverse the space without noticeable interruptions. We also report on the physical walking trajectories of the users in the VEs.

2) USER EXPERIENCE

In the existing literature, the noticeability in VEs is measured utilizing two psychometric methods [33], [34], [35], [36], i) the Likert scale and ii) the Two-Alternative Forced Choice (2AFC) test. The Likert scale is a five point scale-based method used for allowing individuals to express their agreement or disagreement with a given statement. It is easy to set up and implement, but sensitive to different types of biases. Subjects can, for example, have the tendencies to change their answers in an attempt to be “good” test subjects [37]. Two-Alternative Forced Choice (2AFC) is a test to measure the sensitivity of a subject to an input. It assesses the noticeability of testers to certain actions by subjecting them to cases where the action is active and cases where it is not. This would indicate the subjects actually recognizing the action by correctly identifying the two cases.

We utilize the Likert scale due to the complexity of the 2AFC and more frequent use of the former in the literature. The lack of knowledge the testers had about the purpose of the experiments and the absence of any positive, negative, or socially acceptable answer allowed us to reduce the bias related to the usage of the Likert scale. Moreover, we use two performance metrics for characterizing the user experience, specifically the noticeability of VEs and nausea-ness of the users after each of them. This is done using a 5-point Likert scale, which is similar to the simulator sickness questionnaire used in [7]. In addition to the questionnaire that was given to the testers, exhaustive data was collected during the experimentation. Specifically, the user’s positions, rotations, numbers of collisions, and algorithm’s recommendations were recorded to analyze the real world noticeability and performance.

V. EXPERIMENTATION RESULTS

A. PERFORMANCE OF THE APF ALGORITHM

The average number of resets that occurred during one runtime, as well as the average distance traveled between such resets, are depicted in Figure 7 for both random and straight path VEs in a single-user system. As visible in the figure, the random path experience yielded better performance across all environmental sizes compared to the straight experience. The random VE consists of approximately 200 m long randomly generated path with curved sections. Those curved sections cause the user to experience the rotations from the VE itself, in addition to the recommended rotations of the APF algorithm. This results in a more curved physical path and, consequently, a lower number of resets and larger distances between them compared to the straight path VE.




FIGURE 7. RDW performance for a single-user system in different VEs and for different sizes of the tracking space without introducing the proposed generic design enhancements.

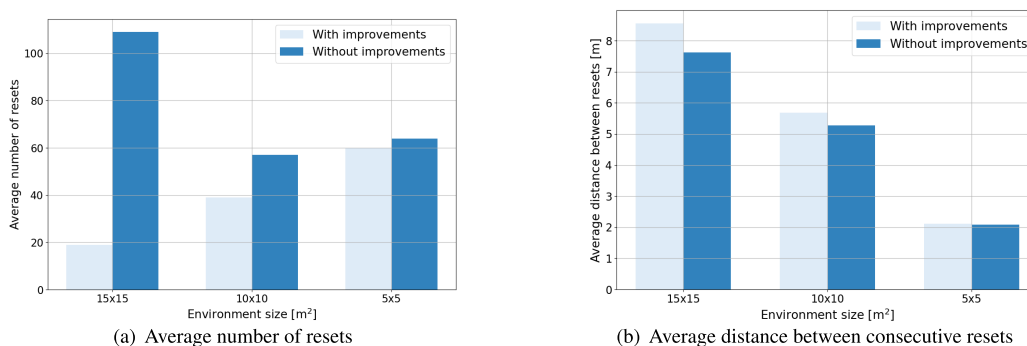


FIGURE 8. RDW performance improvement for a single-user system, straight VE, and different sizes of the tracking space due to the introduction of the proposed generic design enhancements.

For the 5 × 5 m environment, the difference between the experiences is negligible (3.23% in the number of resets and 0.17% in the distance between them). This is because such a small environment does not provide the algorithm with enough space for proper steering, resulting in the comparable performance for the two VEs. The difference between the RDW performance for the two VEs starts being visible in the 10 × 10 m environment. In terms of the number of resets, a decrease of around 40% was observed for the random VE compared to the straight VE. Same as before, we assume the additional rotation provided by the curved path in the random VE to be the cause for this improvement.

Surprisingly, our results showed an increase in the average number of resets at the 15 × 15 m environment compared to the smaller configurations. Intuitively, this number should be lower than for the 10 × 10 m environment, since a larger space can accommodate more uninterrupted movements. This was also confirmed by the RDW simulator [9]. The reason behind the unexpected behavior in the 15 × 15 m environment lies in the high frequency of positional updates (nb., which has been addressed with the proposed generic enhancements). Specifically, in some cases when a reset is triggered, it takes a couple of seconds for the user to stop and rotate itself while the HMD keeps sending positional updates to the server. This results in the APF perceiving the user as not reacting and, as a consequence, continuing to recommend a reset. In reality, the

user visually only got one reset, while the algorithm recommended multiple consecutive and identical ones. We suspect that this only occurs in the 15 × 15 m environment because of the APF resonating redirection set to 7.5 m.

Using the same parameters, we analyzed the performance improvements due to the introduction of the proposed design enhancements of the RDW in general and the APF algorithm in particular. As visible in Figure 8, the introduced generic enhancements contribute to the better performance of the APF algorithm for the straight VE compared to the baseline. Similar results were observed for the random VE and are, therefore, omitted. Our improvements result in roughly 30% and 80% reduction in the average number of resets for the 10 × 10 and 15 × 15 m environments, respectively. As for the average distance between resets, the performance in the 10 × 10 and 15 × 15 m environments is increased by 15% and 10%, respectively. The discrepancy between the 15 × 15 and 10 × 10 m environments in terms of the average number of resets without improvements is again observed and resolved with the introduced enhancements. However, in the 5 × 5 m environment the APF algorithm maintains an impractically large number of resets, even with the improvements introduced, suggesting practical infeasibility of this setup.

Figure 9 depicts the physical path a randomly selected tester followed, both without and with the introduction of the generic RDW design enhancements. As visible, without




FIGURE 9. User movement trajectory during a single experiment without and with introducing the proposed generic design enhancements.

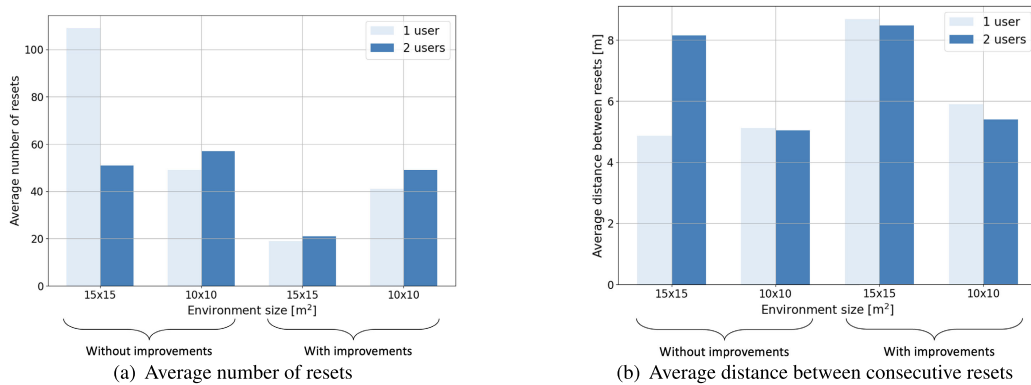


FIGURE 10. RDW performance degradation due to the introduction of the second user in the straight VE without and with the introduction of the proposed generic design enhancements.

the enhancements the user is sometimes able to “escape” the physical environment and wander outside its boundaries, as indicated with grey circles in Figure 9a. This was sometimes caused by the user not immediately reacting to a reset. Moreover, the APF algorithm utilizes a distance threshold for triggering the reset mechanism. This threshold sometimes was not properly adjusted to the small environments, resulting in the algorithm triggering the resets too late. In the enhanced design, at every reset point the algorithm directs the user away from the walls, followed with a short “cool down” period. This way we were able to circumvent the triggering of two consequent resets that were degrading the algorithm’s performance and user experience.

Given that the 5 × 5 m environment yielded poor performance in the single-user experiments discussed above, it was not considered in the two-user experiments, with results as depicted in Figure 10. As visible, for both of the considered metrics, the introduction of the second user in the 15 × 15 m environment has insignificant effect on the performance of the APF in the scenario with introduced generic improvements. However, without the introduction of the proposed generic design enhancements in the 15 × 15 m environment, the introduction of the second user actually benefits the average number of resets and the distance between them.

The reason for that can be found in the above-discussed resonating radius of APF, which was set to 7.5 m. When the second user is introduced to the system, it acts as an additional force vector and, therefore, introduces additional rotations to the first user and “breaks” the resonance, causing the improvements in the observed metrics. Moreover, the algorithm’s performance in terms of the number of resets significantly degrades (i.e., around 30%) due to the introduction of the second user in the 10 × 10 m environment. Interestingly, the introduction of the second user in the system has only minor effects on the average distance between resets (i.e., less than half a meter), as visible in the figure.

Finally, the effects of introducing the proposed generic enhancements in a two-user system is depicted in Figure 11 for the straight VE. Similar trends have been observed for the random VE and these results are, thus, omitted. As visible in the figure, the proposed design enhancements significantly reduce the number of observable resets experienced by the users, as well as the distance between them. For example, in the 15 × 15 m environment the number of such resets in a two-user system is reduced to from more than 50 to around 20. At the same time, the average distance between consecutive resets has been increased by roughly 0.5 m for both sizes of the tracking space.

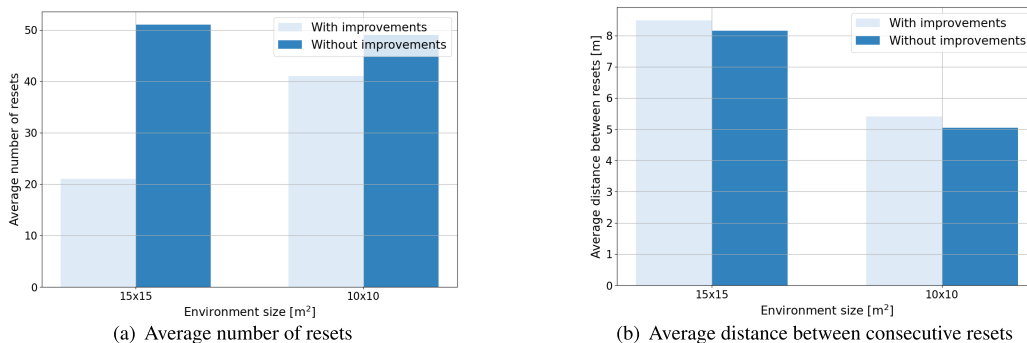


FIGURE 11. RDW performance improvement for a two-user system, straight VE, and different sizes of the tracking space due to the introduction of the proposed generic design enhancements.




FIGURE 12. User experience for a single-user system in different VEs and for different sizes of the tracking space without introducing the proposed generic design enhancements.

B. USER EXPERIENCE

In the following figures, the median is used instead of the mean value due to the fact that qualitative numeric data is being depicted. First, different environments were compared with regard to noticeability and nausea levels for a single-user system without introducing the proposed generic design enhancements, as depicted in Figure 12. As visible from the figure, for all environment sizes and for both the random and straight VEs the testers reported only mediocre noticeability levels. In the 15 × 15 m environment, this is presumably due to the above-discussed issue with the resonating redirection radius. Moreover, the curved path in the random VE only affected the noticeability in the 15 × 15 m environment compared to the straight VE. Apart from that, the APF algorithm generally struggled to unknowingly steer the users in both VEs, regardless of the size of the environment.

At the same time, the testers reported highly comparable and acceptable nausea levels in both 15 × 15 and 10 × 10 m environments, presumably due to the insufficient additional rotation introduced by the curved path in the random VE. This is with an exception of a minor variation in the reported nausea levels across VEs in the 15 × 15 m environment, where the more complex straight VE yielded better performance, presumably due to the subjectiveness of the testers’ experiences. The 5 × 5 m environment was shown to be too small for comfortably accommodating one user and the results on

the noticeability and nausea confirmed this insight. The high number of resets (around 60 during a five minute runtime) meant the user had to drastically turn every five seconds, resulting in most of the test subjects getting nauseous before the end of an experiment. At this point, noticeability became irrelevant because the users had to stop their VEs.

Second, the effectiveness of our enhancements with regard to the noticeability and nausea were studied. Figure 13 shows the results for the straight VE. One can observe that the enhancements substantially reduced the noticeability, while having a minor positive affect on the reported nausea levels. The achieved reduction in noticeability varied from 1 point (5 × 5 and 10 × 10 m environments) to 2 points (15 × 15 m environment). Again, for the 5 × 5 m environment the nausea level was rather high, rendering the noticeability as irrelevant.

In addition, we assessed the change in the user experience due to the introduction of a second user, as depicted in Figure 14. As visible in the figure, the introduction of the proposed generic enhancements mainly affects the noticeability level of the VEs. For example, the noticeability in the 15 × 15 m environment is decreased by two points due to the introduction of the proposed enhancements and regardless of the number of users in the system. Previously, we have concluded that the introduction of an additional user does not lead to the significant degradation of the RDW performance in scenarios with the introduced generic improvements. Sim-

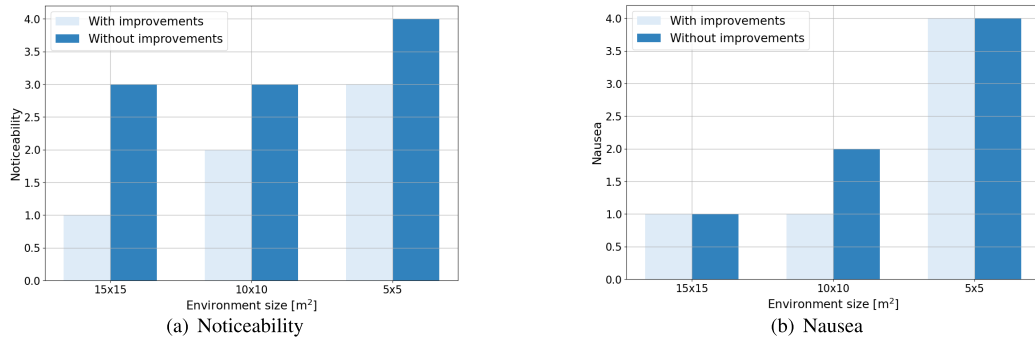


FIGURE 13. Improvement in user experience for a single-user system, straight VE, and different sizes of the tracking space due to the introduction of the proposed generic design enhancements.

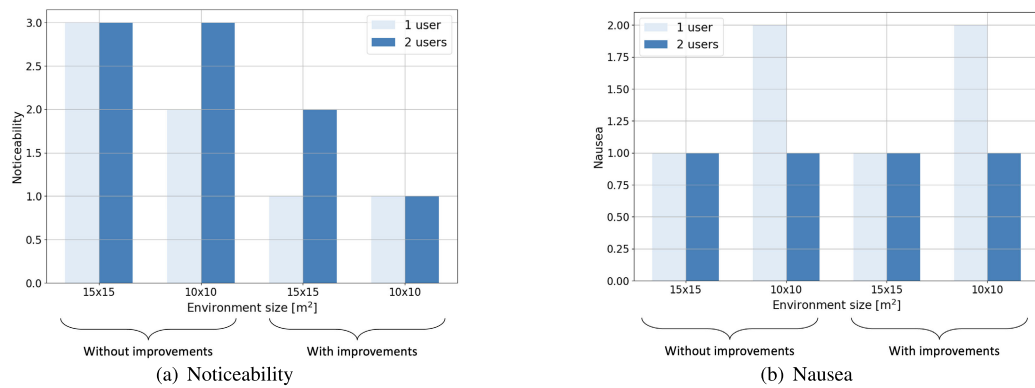


FIGURE 14. Change in the user experience due to the introduction of the second user in the straight VE without and the introduction of the proposed generic design enhancements.

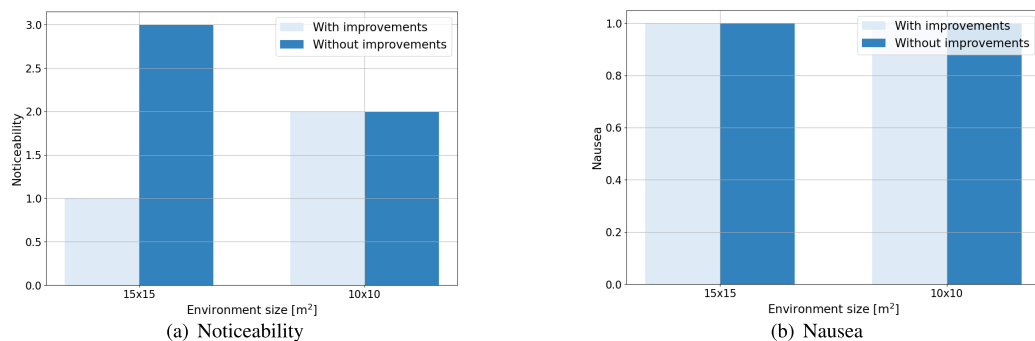


FIGURE 15. Improvement in user experience for a two-user system, straight VE, and different sizes of the tracking space due to the introduction of the proposed generic design enhancements.

ilar observations can be made by looking at the change in the noticeability due to the introduction of the second user. In the 15 × 15 m environment, the increase in the noticeability might be due to the additional steering by the APF algorithm to push the users away from each other, which introduces more opportunities to notice the redirection and break the immersion. However, we did not observe an increase in noticeability in the 10 × 10 m environment. We hypothesize the reason for that coming from the fact that the movement of users changes across experiments, resulting in a variation of the reported experiences. This might also explain the small discrepancy in

the nausea levels depicted in Figure 14b, where the introduction of the second user seemingly reduces the nausea. Finally, Figure 15 depicts the change in the user experience due to the introduction of the generic improvements in the RDW design. As visible, a two point improvement in the noticeability level has been observed for the 15 × 15 m environment, presumably due to the above-discussed consecutive and identical recommendations yielded by the algorithm in this environment. Similar to the one-user system, no significant change in the nausea level has been observed due to the introduction of the improvements.

VI. CONCLUSION

Based on an extensive set of real-life experiments and observations, we have developed a series of generic design enhancements of contemporary Redirected Walking (RDW) algorithms. The enhancements pertain to the mitigation of negative effects of drifts and head rotations, reduction of border bouncing and trajectory oscillations, and introduction of latency-awareness in RDW. The proposed enhancements were tested across a number of scenarios distinguished based on the number of users, type of Virtual Environment (VE), and size of the tracking space. We have demonstrated that our enhancements improve the RDW performance by reducing the average number of resets and increasing the average distance between them. In addition, our improvements were shown to decrease the noticeability of VEs, resulting in enhanced user experience. Our results also indicate that the APF performance and user experience are affected more by the size of tracking spaces and type of VEs than by the introduction of a new user in a Virtual Reality (VR) system. Due to the limited number of testers and experiment iterations in our study, the user experience-focused performance of the APF algorithm reported in this work should at this time be considered only as an indication.

To support the experimental study for capturing the RDW performance, we have developed a modular framework for straightforward experimentation with different RDW algorithms, varying number of VR users, and different VEs. We have demonstrated the capabilities and validated the usefulness of our framework by performing experiments with two RDW algorithms (i.e., APF with and without improvements), two types of VEs (i.e., the users walking along straight and random virtual paths), and three sizes of tracking spaces. We argue that, despite the limited number users in our multiuser-focused experiments, our results demonstrate the utility of the framework and the usefulness of the proposed generic design enhancements in full-immersive multiuser VR setups with RDW.

The results presented in this work are limited to single- and two-user systems primarily due to the lack of necessary hardware and time restrictions. Future work will focus on extending the experimental study to more test subjects for enhancing the reliability of our findings, primarily focusing on the user experience metrics that feature higher variability. In addition, the benefits of the proposed generic enhancements are currently assessed jointly, hence the individual contribution of each instance is lacking and a target of our future efforts. Similarly, each design enhancement instance was parameterized utilizing average values obtained through experimentation (e.g., latency in latency-aware RDW and motionless HMD-specific drift). Clearly, a more precise parameterization accounting for the environmental and temporal variability of the relevant parameters could intuitively further enhance the benefits of the generic enhancements and is considered as a part of our future efforts. Moreover, we will be extending a set of possible experiences in VEs to encompass the ones in which solely the experience is shared, as well

as the ones in which both the experiences and the tracking space are shared by the users. Finally, we will be introducing new RDW algorithms (nb., primarily targeting predictive and scripted RDW approaches), communication protocols, and position extrapolation methods to the framework.

REFERENCES

- [1] *Share & COVID-19 Impact Analysis, By Application (Imaging and Lighting, Medical, satellite, Electronics & Semiconductors), Regional Forecast, 2020–2027*, Xenon Market, Bengaluru, India, 2021.
- [2] T. Joda, G. O. Gallucci, D. Wismeijer, and N. U. Zitzmann, "Augmented and virtual reality in dental medicine: A systematic review," *Comput. Biol. Med.*, vol. 108, pp. 93–100, May 2019.
- [3] S. Bialkova and M. S. Van Gisbergen, "When sound modulates vision: VR applications for art and entertainment," in *Proc. IEEE 3rd Workshop Everyday Virtual Reality (WEVR)*, Mar. 2017, pp. 1–6.
- [4] L. Bozgeyikli, A. Raij, S. Katkooi, and R. Alqasemi, "A survey on virtual reality for individuals with autism spectrum disorder: Design considerations," *IEEE Trans. Learn. Technol.*, vol. 11, no. 2, pp. 133–151, Apr./Jun. 2018.
- [5] J. Struye, F. Lemic, and J. Famaey, "CoVRage: Millimeter-wave beamforming for mobile interactive virtual reality," *IEEE Trans. Wireless Commun.*, early access, Dec. 20, 2022, doi: 10.1109/TWC.2022.3229134.
- [6] J. J. LaViola Jr., "A discussion of cybersickness in virtual environments," *ACM SIGCHI Bull.*, vol. 32, no. 1, pp. 47–56, 2000.
- [7] E. R. Bachmann, E. Hodgson, C. Hoffbauer, and J. Messinger, "Multi-user redirected walking and resetting using artificial potential fields," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 5, pp. 2022–2031, May 2019.
- [8] J. Struye, F. Lemic, and J. Famaey, "Towards ultra-low-latency mmWave Wi-Fi for multi-user interactive virtual reality," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2020, pp. 1–6.
- [9] F. Lemic, J. Struye, and J. Famaey, "User mobility simulator for full-immersive multiuser virtual reality with redirected walking," in *Proc. 12th ACM Multimedia Syst. Conf. (MMSys)*, Jun. 2021, pp. 293–299.
- [10] N. C. Nilsson, T. Peck, G. Bruder, E. Hodgson, S. Serafin, M. Whitton, F. Steinicke, and E. S. Rosenberg, "15 years of research on redirected walking in immersive virtual environments," *IEEE Comput. Graph. Appl.*, vol. 38, no. 2, pp. 44–56, Mar. 2018.
- [11] T. D. Schepper, B. Braem, and S. Latre, "A virtual reality-based multiplayer game using fine-grained localization," in *Proc. Global Inf. Infrastruct. Netw. Symp. (GIIS)*, Oct. 2015, pp. 1–6.
- [12] T. Nescher, Y.-Y. Huang, and A. Kunz, "Planning redirection techniques for optimal free walking experience using model predictive control," in *Proc. IEEE Symp. 3D User Interfaces (3DUI)*, Mar. 2014, pp. 111–118.
- [13] S. Razzaque, *Redirected Walking*. Chapel Hill, NC, USA: The Univ. of North Carolina at Chapel Hill, 2005.
- [14] F. Steinicke, G. Bruder, T. Ropinski, and K. Hinrichs, "Moving towards generally applicable redirected walking," in *Proc. Virtual Reality Int. Conf. (VRIC)*, 2008, pp. 15–24.
- [15] J. Feasel, M. C. Whitton, and J. D. Wendt, "LLCM-WIP: Low-latency, continuous-motion walking-in-place," in *Proc. IEEE Symp. 3D User Interfaces*, Mar. 2008, pp. 97–104.
- [16] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe, "Estimation of detection thresholds for redirected walking techniques," *IEEE Trans. Vis. Comput. Graphics*, vol. 16, no. 1, pp. 17–27, Jan./Feb. 2009.
- [17] J. Dichgans and T. Brandt, "Visual-vestibular interaction: Effects on self-motion perception and postural control," in *Perception*. Berlin, Germany: Springer, 1978, pp. 755–804.
- [18] M. Lappe, F. Bremmer, and A. V. van den Berg, "Perception of self-motion from visual flow," *Trends Cogn. Sci.*, vol. 3, no. 9, pp. 329–336, Sep. 1999.
- [19] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe, "Analyses of human sensitivity to redirected walking," in *Proc. ACM Symp. Virtual reality Softw. Technol.*, Oct. 2008, pp. 149–156.
- [20] A. Nguyen, F. Cervellati, and A. Kunz, "Gain compensation in redirected walking," in *Proc. 23rd ACM Symp. Virtual Reality Softw. Technol.*, Nov. 2017, pp. 1–4.
- [21] Y.-J. Li, F. Steinicke, and M. Wang, "A comprehensive review of redirected walking techniques: Taxonomy, methods, and future directions," *J. Comput. Sci. Technol.*, vol. 37, no. 3, pp. 561–583, Jun. 2022.

- [22] F. Lemic, J. Struye, and J. Famaey, "Short-term trajectory prediction for full-immersive multiuser virtual reality with redirected walking," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2022, pp. 1–7.
- [23] R. R. Strauss, R. Ramanujan, A. Becker, and T. C. Peck, "A steering algorithm for redirected walking using reinforcement learning," *IEEE Trans. Vis. Comput. Graphics*, vol. 26, no. 5, pp. 1955–1963, May 2020.
- [24] J. Messinger, E. Hodgson, and E. R. Bachmann, "Effects of tracking area shape and size on artificial potential field redirected walking," in *Proc. IEEE Conf. Virtual Reality 3D User Interfaces (VR)*, Mar. 2019, pp. 72–80.
- [25] E. Hodgson, E. Bachmann, and D. Waller, "Redirected walking to explore virtual environments: Assessing the potential for spatial interference," *ACM Trans. Appl. Perception*, vol. 8, no. 4, pp. 1–22, 2008.
- [26] M. A. Zmuda, J. L. Wonsler, E. R. Bachmann, and E. Hodgson, "Optimizing constrained-environment redirected walking instructions using search techniques," *IEEE Trans. Vis. Comput. Graphics*, vol. 19, no. 11, pp. 1872–1884, Nov. 2013.
- [27] D.-Y. Lee, Y.-H. Cho, D.-H. Min, and I.-K. Lee, "Optimal planning for redirected walking based on reinforcement learning in multi-user environment with irregularly shaped physical space," in *Proc. IEEE Conf. Virtual Reality 3D User Interfaces (VR)*, Mar. 2020, pp. 155–163.
- [28] C. Hirt, M. Zank, and A. Kunz, "PREWAP: Predictive redirected walking using artificial potential fields," in *Proc. IEEE Conf. Virtual Reality 3D User Interfaces (VR)*, Mar. 2019, pp. 976–977.
- [29] E. Hodgson and E. Bachmann, "Comparing four approaches to generalized redirected walking: Simulation and live user data," *IEEE Trans. Vis. Comput. Graphics*, vol. 19, no. 4, pp. 634–643, Apr. 2013.
- [30] D. Waller, E. Bachmann, E. Hodgson, and A. C. Beall, "The HIVE: A huge immersive virtual environment for research in spatial cognition," *Behav. Res. Methods*, vol. 39, no. 4, pp. 835–843, Nov. 2007.
- [31] G. Tsaramirsis, S. Buhari, M. Basher, and M. Stojmenovic, "Navigating virtual environments using leg poses and smartphone sensors," *Sensors*, vol. 19, no. 2, p. 299, Jan. 2019.
- [32] E. Hodgson, E. Bachmann, and T. Thrash, "Performance of redirected walking algorithms in a constrained virtual world," *IEEE Trans. Vis. Comput. Graphics*, vol. 20, no. 4, pp. 579–587, Apr. 2014.
- [33] C.-B. Ciumedean, C. Patras, M. Cibulskis, N. Váradi, and N. C. Nilsson, "Mission impossible spaces: Using challenge-based distractors to reduce noticeability of self-overlapping virtual architecture," in *Proc. Symp. Spatial User Interact.*, Oct. 2020, pp. 1–4.
- [34] E. Kruijff, J. Orlosky, N. Kishishita, C. Trepkowski, and K. Kiyokawa, "The influence of label design on search performance and noticeability in wide field of view augmented reality displays," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 9, pp. 2821–2837, Sep. 2019.
- [35] J. Xu and B. W. Wah, "Consistent synchronization of action order with least noticeable delays in fast-paced multiplayer online games," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 13, no. 1, pp. 1–25, Feb. 2017.
- [36] S. Ghosh, L. Winston, N. Panchal, P. Kimura-Thollander, J. Hotnog, D. Cheong, G. Reyes, and G. D. Abowd, "NotifiVR: Exploring interruptions and notifications in virtual reality," *IEEE Trans. Vis. Comput. Graphics*, vol. 24, no. 4, pp. 1447–1456, Apr. 2018.
- [37] A. Joshi, S. Kale, S. Chandel, and D. K. Pal, "Likert scale: Explored and explained," *Brit. J. Appl. Sci. Technol.*, vol. 7, no. 4, p. 396, 2015.



JAKOB STRUYE (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in computer science from the University of Antwerp, in 2015 and 2017, respectively. He is currently a Ph.D. Researcher in the field of wireless networking with the Internet and Data Laboratory (IDLab), research group (University of Antwerp) and the IMEC Research Institute, Belgium. His research interest includes leveraging extremely high frequency wireless networks in the millimeter wave bands to improve the performance of truly wireless virtual and augmented reality. He has experience in applying artificial intelligence to time series prediction problems. He has authored several publications in these fields. He also participated in the DARPA Spectrum Collaboration Challenge, in 2019, with his team participating in the finals.



XAVIER COSTA PEREZ received the M.Sc. and Ph.D. degrees in telecommunications from the Universitat Politècnica de Catalunya (UPC). He is currently an ICREA Research Professor, the Scientific Director of the i2cat Research Center, and the Head of 6G Networks Research and Development with the NEC Laboratories Europe. His research interests include the digital transformation of society driven by the interplay of mobile networks and AI. His team generates research results which are regularly published at top scientific venues, produces innovations which have received several awards for successful technology transfers, participates in major EC research and development projects and contributes to standardization bodies. He has served on the committees of several conferences and has published papers of high impact and holds tenths of granted patents. In 2021, he served as an Editor for IEEE TRANSACTIONS ON MOBILE COMPUTING, IEEE TRANSACTIONS ON COMMUNICATIONS, and *Computer Communications* (Elsevier) journal.



JEROEN FAMAHEY received the M.Sc. and Ph.D. degrees in computer science engineering from Ghent University, Belgium, in 2007 and 2012, respectively. He is currently a Research Professor with the Internet and Data Laboratory (IDLab), research group (University of Antwerp) and imec, Belgium. He leads a team of researchers with IDLab, focusing on future wireless network technologies and protocols. His specific research interests include network modeling and optimization of high-rate low-latency communication protocols, mmWave and THz communications, and energy harvesting wireless systems. He has been involved in more than 25 international, national, and bilateral research projects, both fundamental and industry-driven in nature. His research has led to the publication of over 160 peer-reviewed journal articles and conference papers and seven granted patents.



FILIP LEMIC received the B.Sc. and M.Sc. degrees from the University of Zagreb, in 2010 and 2012, respectively, and the Ph.D. degree from the Technische Universität Berlin, in 2017. He was a Marie Curie Postdoctoral Researcher with the University of Antwerp (2018–2022) and the Universitat Politècnica de Catalunya (2020–2022). He was also affiliated as a Senior Researcher with imec (2018–2022). He was a Visiting Researcher with the University of California at Berkeley, Berkeley (2015–2016) and Shanghai Jiao Tong University (2019–2020). He is currently a Senior Researcher with i2Cat Foundation. He has coauthored more than 60 peer-reviewed research articles and was involved in various international research projects, notably EU MSCA ScaLeITN, EU EVARILLOS, EU eWine, NIST's PerfLoc, and UC Berkeley's beyond-5G.



THOMAS VAN ONSEM is currently pursuing the master's degree in software engineering with the University of Antwerp, Belgium. Currently, he assists companies with their digital transformations as a Technical Consultant with Exellys. His research interests include influential and impactful technologies.