FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Walkable Multi-User VR

The Effects of Physical and Virtual Colocation

PhD THESIS

submitted in partial fulfillment of the requirements for the degree of

Doctor of Technical Sciences

within the

Vienna PhD School of Informatics

by

Iana Podkosova, MSc
Registration Number 01128722

to the Faculty of Informatics

at the TU Wien

Advisor: Dr. techn. Hannes Kaufmann, Priv.-Doz.
Second advisor: Dr. techn. Peter Ferschin, Ass. Prof.

External reviewers:
Prof. Gregory F. Welch, PhD. University of Central Florida, Orlando, Florida, USA.
Prof. Alan Chalmers, PhD. University of Warwick, Coventry, UK.

Vienna, 1st December, 2018

Iana Podkosova | Hannes Kaufmann

Declaration of Authorship

Iana Podkosova, MSc
Am Heumarkt 23/6, 1030 Wien

I hereby declare that I have written this Doctoral Thesis independently, that I have completely specified the utilized sources and resources and that I have definitely marked all parts of the work - including tables, maps and figures - which belong to other works or to the internet, literally or extracted, by referencing the source as borrowed.

Vienna, 1st December, 2018

Iana Podkosova

Acknowledgements

First of all, I would like to thank my advisor Hannes Kaufmann for his continuous support, encouragement, patience, valuable advice and always inspiring feedback. Few PhD students are lucky to have such an attentive and kind advisor.

I would like to thank Prof. Christian Breiteneder for his support and interesting discussions, those regarding my PhD research and also those completely unrelated to it. Especially the latter type has always reassured me that it is OK to have multiple interests.

Greg Welch and Alan Chalmers kindly agreed to review this thesis. Thank you!

I would like to thank all members of the VR Group at TU Wien for valuable discussions, feedback and a great working atmosphere. Working on ImmersiveDeck together with Christian Schönauer, Emanuel Vonach, Khrystyna Vasylevska, Elisabeth Broneder and Peter Fikar was one of my most favorite times during the whole PhD. In addition, it has been great to share the office with Peter Kán and Minkyung Lee, always being able to spontaneously ask their advice and enjoying their positive mood. A special thanks to Ingrid Lissa for organizing so many things and always having an answer to any administration-related question.

My research was funded by PhD School of Informatics and the Doctoral College Computational Design of the Center for Geometry and Computational Design. I am grateful for this opportunity. The development of ImmersiveDeck was funded by our collaborators in Berlin, IllusionWalk KG. I would like to thank them for sharing and supporting our vision of multi-user VR.

This thesis included many user studies, and I am thankful to all the participants who took part in them, investing their time and traveling to sometimes far away locations in Vienna with little reward but a chance to walk around in a VR world.

I am grateful to my parents for teaching me the value of education and encouraging me to study a technical discipline, and, of course, for their love and support.

Finally, I would like to thank my friends Anna, Rolanda, Manon, Stevie, Flo, Fintan, Julian, Jonathan, Anja, Berni, Wagi and Aimee in Vienna and Sascha, Alex, Lena and Maria in Moscow for bringing so much joy into my life and for defying a popular opinion that PhD years are only good for work.

Abstract

The research presented in this dissertation focuses on multi-user VR, where multiple immersed users navigate the virtual world by physically walking in a large tracking area. In such a setup, different combinations of user colocation within the physical and the virtual space are possible. We consider a setup to be multi-user if at least one of these two spaces is shared. The dissertation starts with the classification of combinations of physical and virtual colocation. Four such combinations are defined: colocated shared VR, colocated non-shared VR, distributed shared VR and shared VR with mixed colocation. The characteristics of each of these four setups are discussed and the resulting problems and research questions outlined.

The dissertation continues with the description of ImmersiveDeck - a large-scale multi-user VR platform that enables navigation by walking and natural interaction. Then, four experiments on multi-user walkable VR developed with the use of ImmersiveDeck are described.

The first two experiments are set in colocated non-shared VR where walking users share a tracking space while being immersed into separate virtual worlds. We investigate users' mutual awareness in this setup and explore methods of preventing mutual collisions between walking users. The following two experiments study shared VR scenarios in situations of varied physical colocation. We investigate the effects that different modes of physical colocation have on locomotion, collision avoidance and proxemics patterns exhibited by walking users. The sense of copresence and social presence within the virtual world reported by users is investigated as well.

The experiments in the colocated non-shared VR setup show that HMD-based VR can produce immersion so strong that users do not notice others being present in their immediate proximity, thus making collision prevention the task of utmost importance. In our proposed method of displaying notification avatars to prevent potential imminent collisions between colocated users, the suitability of a particular type of notification avatar was found to be dependent on the type of scenario experienced by users. The general result of the experiments in shared VR is that physical colocation affects locomotor and proxemics behavior of users as well as their subjective experience in terms of copresence. In particular, users are more cautious about possible collisions and more careful in their collision avoidance behavior in the colocated setup compared to the real environment. In the distributed setup, conventional collision avoidance is often abandoned.

Contents

CHAPTER 1

Introduction

1.1 Introduction to multi-user virtual reality

Multi-user Virtual Reality is a Virtual Reality (VR) setup that can be used by multiple users at a time, providing each of them with an immersive virtual experience. The end goal of multi-user VR discussed in this dissertation is a setup that enables highly natural and intuitive interactions within and with the virtual world for all immersed users. Highly natural means, first, that users should be able to navigate the virtual world in an intuitive way - by walking. The possibility for natural interaction is closely connected to having a virtual body. In the real world, the majority of interactions are body-to-body interactions, i.e. our bodies provide vital information about our activities, goals and intentions. If the aim is to replicate this richness of communication, we have to provide users with virtual bodies that repeat the behavior of their real bodies.

While many sources associate the multi-user aspect with collaborative activities carried out by users, we consider a broader type of a VR setup that provides a technological foundation for possible collaborative scenarios but is not limited to them. A further assumption that is often associated with multi-user virtual environments (VEs) is that they bring geographically distributed users together in the shared virtual space. While this is an important use-case of multi-user VR we do not limit our discussion to it considering distributed and colocated scenarios alike.

This dissertation proposes a system that is our first attempt of reaching the goal of natural and intuitive multi-user VR. Further, we discuss the usage of such a highly immersive and intuitive multi-user system in different combinations of physical and virtual colocation with the special focus on experiences of users walking in VEs.

1.2 Motivation and contribution

This dissertation is concerned with a specific type of multi-user VR environments, those that achieve immersion through the use of head-mounted displays (HMDs) and allow their exploration by natural walking in a large physical space.

The specific focus on navigation by natural walking is motivated by previous research on locomotion techniques that has demonstrated multiple advantages of physically walking in VEs. To name a few, walking in VR is positively correlated with the illusion of presence [128, 131], it helps to correctly estimate the dimensions of the environment [123], improves spatial understanding and task performance [141, 115]. To navigate VR environments by walking, a dedicated physical space and a suitable tracking technology are necessary. The application potential of walkable multi-player VR covers multiple application domains: team training, rehearsals [103], collaborative visualization, entertainment. To allow a wide variety of applications where users can explore the VE by walking a large tracking area is beneficial. Team rescue training, team sports playthroughs and dancing rehearsals are examples of areas which can benefit from the use of multi-player VR but require large spaces for the practice to be effective. A number of techniques have been developed to allow walking in virtual spaces that are larger than the available physical tracking space [113, 137]. These techniques however generally impose a lower limit on the size of a required tracking area.

When we talk about a multi-user VR system that supports natural walking, two spaces are important: the physical space where the immersed users are situated (tracking space) and the virtual space where the experience happens. In the context of this dissertation, at least one of these two spaces has to be shared by users for the system to be considered multi-user. Furthermore, different experiences for users and therefore different research questions arise depending on which of the environments in question are shared: the tracking space, the virtual space or both of them.

The first contribution presented in this dissertation is a large-scale multi-user VR platform that allows its users natural interactions within and with a virtual world. Intuitive navigation is achieved by the ability to walk within a large tracking area. The possibility of natural interactions between users is achieved through the use of full body motion capture to present users with avatars that repeat their physical movements.

With the use of the developed platform, we investigate research problems that arise when users walk in presence of each other within shared spaces. First, we investigate how to ensure users' safety and comfort when they move within a shared tracking space without sharing a VE. The results of two experiments on mutual awareness and collision prevention contribute to this investigation. Further, we turn to the investigation of scenarios where the VE is shared by users. In this case, we investigate how locomotory patterns of users walking within a shared virtual space compare to those displayed in real spaces. Finally, we study the influence that relative physical location of users (shared or non-shared tracking space) may have on their locomotory behavior within shared virtual worlds and on the subjective experience of users being together in these worlds.

1.2.1 Resulting publications

The results presented in this PhD dissertation have appeared in the following peer-reviewed publications:

- Iana Podkosova, Khrystina Vasylevska, Christian Schoenauer, Emanuel Vonach, Peter Fikar, Elisabeth Broneder and Hannes Kaufmann. "Immersive Deck: A large-scale wireless VR system for multiple users". In the proceedings of the 9th IEEE Workshop on Software Engineering and Architectures for Realtime Interactive Systems (SEARIS), pp. 1-7, Greenville, SC, USA, 2016.
- Iana Podkosova and Hannes Kaufmann. "Mutual proximity awareness in immersive multi-user virtual environments with real walking". In the proceedings of the 25th International Conference on Artificial Reality and Telexistence and 20th Eurographics Symposium on Virtual Environments (ICAT-EGVE), pp. 109-116, Kyoto, Japan, 2015.
- Iana Podkosova and Hannes Kaufmann. "Preventing imminent collisions between co-located users in HMD-based VR in non-shared scenarios". In the proceedings of the 30th Conference on Computer Animation and Simulated Agents (CASA), pp. 1-10, Seoul, South Korea, 2017
- Iana Podkosova and Hannes Kaufmann "Mutual collision avoidance during walking in real and collaborative virtual environments". In the proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (i3D), pp. 1-9, Montreal, Canada, 2018.
- Iana Podkosova and Hannes Kaufmann "Proxemics and copresence during walking in shared virtual environments with mixed colocation". In the proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology (VRST), pp. 1-10, Tokyo, Japan, 2018.

1.3 Thesis organization

Chapter 2 presents a taxonomy of walkable multi-user VR scenarios based on the relative colocation of users in the real and virtual space. Research questions associated with each scenario are outlined.

Chapter 3 gives an overview of previous research and provides theoretical foundations in the areas that informed this dissertation. This includes examples of previously built VR systems with multi-player support as well as areas that are relevant to scientific questions that arise in each combination of the real and virtual colocation.

Chapter 4 describes the design and implementation of the walkable multi-player VR platform.

Chapter 5 describes the experiments in a setup with the shared tracking space and non-shared VEs.

Chapter 6 describes the experiments in shared virtual spaces in the colocated and distributed physical setup.

The dissertation is concluded and outlines for future work are given in Chapter 7.

CHAPTER 2

Taxonomy of Multi-user VR

Immersive walkable multi-user setups that are addressed in this dissertation result in different experiences for users depending on whether the real or virtual or both spaces are shared by the users. Figure 2.1 illustrates four different combinations of shared and non-shared environments.

We will call *colocated shared VR* a setup where several immersed users are located in the same tracking space and are immersed in the same VE where they can see each other. If the users share the same tracking space but do not see each other within the VE (they may be immersed into entirely different VEs) we will talk about *colocated non-shared VR*. An opposite situation is possible: each user is walking around in their own separate tracking area but several immersed users share a common VE. We will call such a setup *distributed shared VR*. Finally, a shared VE can be experienced in a situation of mixed physical colocation. In this case, some users may share a tracking space with a few others and some may be alone in their tracking space. The VE is however shared by all users. We will refer to such a setup as *shared VR with mixed colocation*, or *mixed shared VR* for short. The important property of this mixed setup is that a user is colocated with a bigger number of users in the VE than in the real environment.

Each of the described setups has its advantages for a specific application scenario. Shared VEs can be experienced in colocated or distributed setups depending on the priorities of the application and its users. The distributed setup offers the flexibility of participating in a shared virtual experience from any physical location, on the condition of it being equipped with a tracking system. The colocated setup however has an advantage of enabling direct communication and interactions between users which might be crucial in certain training scenarios. Apart from enabling direct interactions, the colocated setup might be quite simply enforced by the rare availability of an equipped tracking area, especially of large dimensions. As for the colocated setup where the virtual world is not shared by users, there are two possible scenarios for its use. The first one is the aforementioned rare availability of large tracking areas. It has been indeed suggested

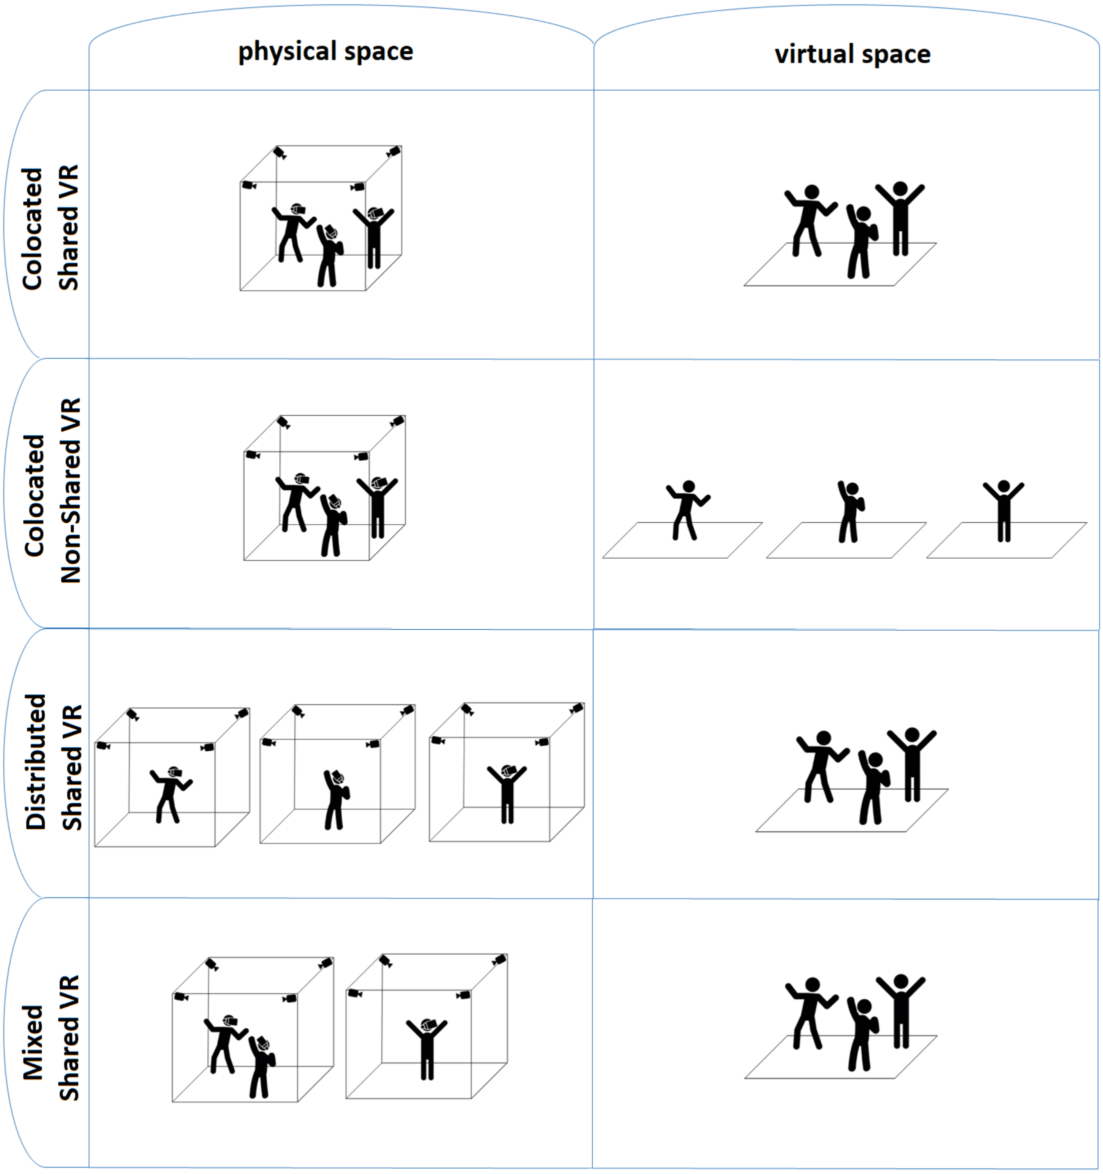

Figure 2.1: Taxonomy of multi-user VR scenarios based on different combinations of physical and virtual colocation.

that the same tracking space could be used for multiple users at a time, each of them immersed into a different virtual world. The second use case arises when several players immersed into the same VE loose the direct mapping between their positions in the tracking space and in the virtual world. For example, consider a situation where a team of colocated players start an exploration game or a training scenario in a building with multiple floors. In the beginning of the game, all players may be located in the same virtual room on the ground floor, perfectly visible to each other. With the progress of the game, some of them travel to different floors of the virtual building and stop seeing other players who stayed on the ground floor. For these players, the experience has transformed itself from the colocated shared VR to the colocated non-shared VR.

In this chapter, we discuss features inherent to each of the described multi-user VR setups and consequently research questions that arise when users are walking together within these setups.

2.1 Colocated Non-Shared VR

In the colocated non-shared setup users do not see each other while possibly being very close to each other in the tracking space, a situation that leads to two possible problems. The first one, and also the one causing the most concern is that users may collide with each other. The second problem is that users may experience breaks of presence because of close physical proximity of someone who they do not see in the VE. Therefore, enabling safe and comfortable exploration of VEs without the risk of collisions or breaks in presence is the major research task for colocated non-shared VR.

2.2 Shared VR scenarios

Shared VR where walking in large physical space is possible is well-suited for applications such as team training and rehearsals, for example for fire-fighters, first responders etc. Such applications require a lot of movement, possibly fast walking or even running, within a VE. Therefore, it is important to study how patterns of users' movement within shared VR compare to the real world and how they are influenced by the physical setup, to be able to plan the transfer of skills from VR to the real world and possibly account for differences between user behavior in VEs and in reality.

Users of the colocated shared VR setup *are there together*, both in the physical and virtual space. This setup can be a powerful tool for eliciting the sense of copresence as visual cues of the body of the other player presented through an avatar can be combined with haptic cues of the other's actual body. Direct spoken communication is also possible within the same tracking area. Haptic interactions can be augmented by bringing in physical objects for users to interact with and to pass to each other. Accurate distribution of players' movements is especially important in colocated shared VR as asynchronicity between haptic and visual cues can negatively contribute to the sense of presence and copresence.

<table>
<tr><th>Colocated Non-Shared VR</th><th>Colocated Shared VR</th><th>Distributed Shared VR</th><th>Mixed Shared VR</th></tr>
<tr><td rowspan="2">How to prevent collisions?

Do users cause breaks of presence for each other?</td><td colspan="2">How does physical colocation influence experienced co-presence?</td><td>Are physically colocated co-players perceived in the same way as distributed ones?</td></tr>
<tr><td colspan="3">How do users move in the presence of each other?
How do locomotory patterns compare to those observed in real spaces?</td></tr>
</table>

Figure 2.2: Research questions addressed in this dissertation, classified in respect with each of the described multi-user VR scenarios.

In distributed shared VR, users are immersed into the same VE while being physically located in different tracking spaces, possibly far away from each other. Therefore, it might be more difficult for users to suspend the disbelief and develop an illusion of being together in the virtual world. The lack of perceived copresence may lead to situations where users treat each other's avatars as merely computer-generated images, for example walking through them. Such a situation would possibly undermine the goals of a shared VR application, especially if it is designed to train users for some tasks that would later be performed in a real setting. Shared VR scenarios with mixed colocation might, on one hand, positively contribute to the illusion of copresence for these users who have physically colocated co-players. On the other hand, a setup in which some users are physically colocated and some are not might produce discrepancies in the perception of colocated and remote users negatively affecting perceived copresence of the latter ones.

To summarize, this dissertation attempts to contribute to the answers on the following questions:

How to provide comfortable and collision-safe exploration of virtual worlds?

How are movement patterns of users walking within virtual worlds affected by physical colocation? How does this effect contribute to the sense of being together in a virtual world?

Figure 2.2 presents the connection of these questions to the described combinations of physical and virtual colocation.

CHAPTER 3

Background and Theoretical Foundations

This chapter discusses previous research in several areas relevant to the work presented in the dissertation. These areas cover the key aspects of VR systems in question: multi-user, walkable and large-scale.

First, the review of previous multi-user and some single-user but otherwise relevant VR systems is given. Then, we turn to the research on how people manage space and avoid collisions when they walk in proximity of each other, informed by studies performed in the real world and in VR. The theory of collision avoidance is important for situations of both shared and non-shared VEs that require active measures for collision prevention. Finally, we discuss the notions of copresence and social presence that arise when virtual worlds are shared by users. Embodiment is discussed in this section as well as it necessarily affects user experience in shared VEs.

3.1 Multi-user VR systems

First multi-user VR systems date back to the time when VR itself was an emerging technology heavily constrained by hardware and computing limitations. The DIVE (Distributed Interactive Virtual Environment) [29, 30, 63] system presented in 1993 provided an architecture for developing interactive multi-user VR applications implementing a space-based model for user interaction [42]. The system allowed a large number of users to connect simultaneously to a virtual world over the Internet. Users and automated players were abstracted as actors with graphical embodiments. The 6 DOF position of an actor's embodiment could be controlled by a variety of input devices, from a conventional mouse to data gloves and fully tracked HMDs. The networking architecture of DIVE was based on the active replication protocol with peer-to-peer multicast communication but evolved to additionally use client-server connections to a name server [51]. Interactions and modifications to the state of a world (a shared virtual context) which was held in the form of a shared database were first applied locally and then replicated to all connected peers. This approach ensured the fastest response of the virtual world to the actions of the interacting user while all the other users received the changes with a delay. A concurrency control mechanism used dead reckoning [84] to achieve the consistency of local copies. Although the system tolerated some initial inconsistency between peers the equality over time was ensured. Additionally it provided mechanisms to divide a virtual world into sub-hierarchies that were only replicated to applications currently interested in them to limit network traffic. An actor was able to move between different worlds through "gateways", objects with embedded trigger functionality that allowed to query a new multicast group to join and subsequently move the actor with its graphical embodiment to a new world. While the DIVE platform itself was written it plain C, it offered three different ways of building specific applications: using C, using DIVE/TCL [50] scripting language to describe the behaviors of objects, or using the DCI interface [49] to communicate with external applications.

Another software platform for developing multi-user interactive environments was presented in 1995 under the name Spline [2]. It had a similar to DIVE networking architecture that used a mixed multicast and client-server approach. The major differences between DIVE and Spline were in the mechanisms used for partitioning a virtual world into sub-regions, as well as in the choice of a scripting language used for virtual objects creation. Spline was used to create a prototype virtual environment Diamond Park, a social VR environment themed around cycling [154].

MASSIVE [61] was a multi-user VR system designed specifically for teleconferencing. Users represented by T-shaped avatars could communicate with each other via a combination of typed messages, graphical gestures and live audio. The system was further developed into MASSIVE-2 [18, 60] and later MASSIVE-3 [62], a general-purpose collaborative VR platform. MASSIVE-3 adopted a distributed database model with the logically multicast but physically client-server based networking. Several primary-based consistency mechanisms which could be used for different objects were employed: data ownership transfer, centralized update or a combination of them. World structuring used an extension of the mechanism used in Spline. Similarly to DIVE, MASSIVE-3 allowed objects and interactions to be applied locally only.

AVANGO, initially named Avocado, was a framework that supported the development of distributed and interactive VR applications. It was presented in 1999 [144] and later revised in 2008 [77]. In contrast to the former presented frameworks it had a focus on rapid prototyping using Python as its scripting language and on shielding developers from the distribution-specific aspects. The underlying networking architecture, similarly to DIVE, was built on the concept of distributed shared memory with active replication and reliable multicast communication between peers.

The COVEN (Collaborative Virtual Environments) platform was developed within the scope of the COVEN project [104, 47]. The platform was based on two separate VR systems. The primary development system was Division's dVS [116], a commercial VR system at the time, and was intended as the final delivery system. DIVE was used as a research system for exploratory prototyping. DIVE was extended to include DiveBone [48], an application-level multicast backbone as well as live audio and video sharing support and individualized for each user world views. The COVEN platform was used in a longitudinal study of user behavior and computational demands of collaborative virtual environments [145].

Studierstube [142, 119] was an example of a collaborative Augmented Reality (AR) system presented in 1998. The system focused on local collaboration, with all collaborating users being in the same physical location while sharing a virtual environment (a virtual augmentation of a real room). Each user was equipped with a magnetically tracked see-through HMD that allowed to see virtual augmentations. The networking architecture used a client-server approach where modifications to the state of a shared virtual object performed by one user are sent to the server that subsequently distributes them to other users. Being an AR system for colocated users, Studierstube did not have to use any graphical representations of users or artificial concepts for user-user interaction which are among the biggest design challenges of multi-user VR systems. The system was used to create Construct3D [71], a collaborative application for geometry education.

Some of the later multi-user VR systems specifically focused on real walking in a large physical space. The HIVE (Huge Immersive Virtual Environment) system presented in 2007 [152] enabled users to naturally walk inside large virtual environments. The system used an infrared optical outside-in tracking solution working with eight cameras and a HMD attached to a rendering computer worn by a user. The tracking component of the system could track either head positions of several users or several body parts of one user in the area of 572 m^2. Tracking ran on a separate stationary machine and the position data was sent to the rendering system over a 802.11x wireless connection where it was fused with the orientation data obtained from an inertial device also worn by the user. Additionally, the setup included a separate graphics station that monitored the state of the virtual environment. The system used Vizard and Panda3D software that was extended with HIVE-specific API to interface various input devices.

Further work of the authors of the HIVE included the development of an untethered large-scale and low-cost walkable VR system for multiple users. Self-contained incremental inertial tracking corrected by periodic absolute measurements and redirected walking techniques [113, 68] were proposed to enable free explorations of very large virtual environments [6]. A practical implementation of the proposed concepts resulted in the WeaVR system [67] used in an outdoor environment. Inertial tracking units attached to users' feet providing relative position updates were combined with a GPS module used for absolute position corrections. All processing was done on a laptop attached to a backpack frame, with an HMD used for visual output. The experiments conducted with the system only demonstrated its use for a single user at a time, however multi-user capabilities were claimed to be easily provided by adding an ad-hoc network to the system.

None of the described VR setups supporting real walking featured full body tracking of multiple users. In the BEAMING project [136, 110], users' body movements were tracked with the use of a motion tracking suit and transmitted to a robot avatar representing the user at a remote location. Motion data of arms and hands was streamed in real time over the Internet. However, the robot's movements were not completely in sync but loosely coupled with the user's movements. In a projection-based system used for group to group telepresence [15], a cluster of depth and color cameras was used to capture movements of the users and reconstruct them at a remote location.

In case of live interaction between several users that share the same physical space, full sync of the users' avatars is needed in real time to enable free interaction, including the possibility of direct haptic interactions.

When we started working on ImmersiveDeck, none of the systems provided a combination of all the functionality that we considered necessary for our purposes. The first generation of multi-user VR systems relied on well-investigated networking concepts and provided thorough theoretical background on CVEs. However, the technology with which these concepts were implemented was outdated and a new core for the system had to be found. The modern systems that used the advantages of real walking were only in theory or in plans multi-user, with no record of multi-user experiments, not to mention full-body motion capture and its distribution. Shared object interaction, very much discussed and present in the early systems, was not present in any setups with real walking. Furthermore, our special challenge was in creating a system for colocated users (as a base system), the fact that meant pushing the tolerance to network-related delays to a minimum.

3.2 Walking with others: proxemics and collision avoidance in real environments and VR

3.2.1 Theory of proxemics and collision avoidance

In the real world, an implicit set of rules are at work when physical space is shared between individuals. The process of walking in the surroundings of others entails a constant effort of trajectory planning and collision avoidance governed by these rules.

Human interactions in the spatial context have been widely studied in the field of sociology and psychology. Goffman's [58] theory on pedestrian interactions has been influential in this field. According to his observations, two processes govern the avoidance behavior of pedestrians, the externalization and the scanning. In the externalization, a pedestrian uses body gestures to inform the surrounding others about his intentions. At the same time, he continuously scans the environment to collect and interpret body signals sent by other pedestrians. This non-verbal interaction eventually results in the coordination of actions between two interacting pedestrians and potential collisions are resolved. Wolff [160] argues that the collaboration is an essential part of the interaction and pedestrians expect to be cooperative upon interacting with each other. Personal space is another important concept in interpersonal interactions. Sommer [133] describes the personal

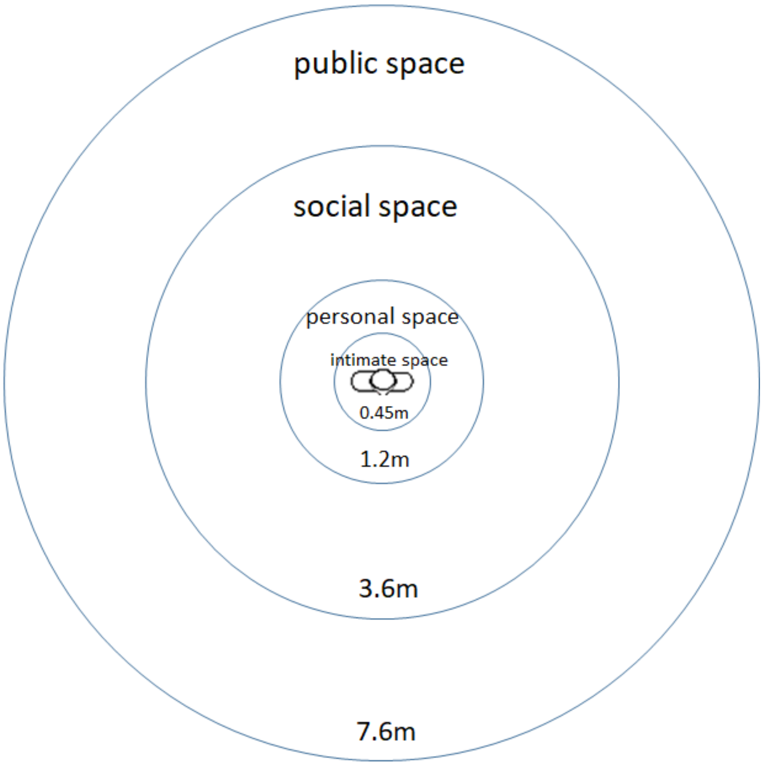

Figure 3.1: Interpersonal distances according to the proxemics theory of Hall [64].

space as the portable territory around an individual that others should not violate. It regulates the psychophysical distance that the individual needs to maintain in order to feel comfortable.

Proxemics first introduced by Hall [64] studies how humans manage space during interactions with the environment, including others present in their surroundings. Interpersonal distance plays an important role in proxemics. According to Hall [64], this distance can be used to determine the nature of the relationship between the individuals and decreases as the the level of intimacy increases. In his classification of the distances that people prefer to keep, Hall identifies four main zones: intimate, personal, social and public. illustrated in Figure 3.1. These zones provide a sense of socially and culturally associated acceptable distances for different types of people based on familiarity. The proposed interpersonal distances can vary significantly depending on the gender and the age of the interacting individuals, as well their ethical and cultural background. In addition, the reported distances were obtained in static observational conditions (e.g. people waiting on a train platform). The distances that people prefer to maintain while walking may differ from those proposed by Hall.

Since Hall's work, a considerable amount of research has focused on the interpersonal spatial interactions. Dabbs and Stokes [36] reported that standstill pedestrians grant more space to approaching male pedestrians than to female pedestrians. In contrast, Sobel and Lillith [132] reported that females were given more personal space than males. The contradiction could be due to the fact that Dabbs and Stokes studied situations where one of the pedestrians was stationary, whereas Sobel and Lillith focused on cases where both interacting pedestrians were moving. Caplan and Goldman [28] suggested that besides gender, the physical dominance can also affect the size of the personal space. In their observations, pedestrians invaded the space of short people more frequently compared to that of tall people. Hartnett et al. [65] studied the role of height and gender in a laboratory setting. In their experiment, participants were instructed to approach a standing or a sitting subject and to stop when they felt uncomfortable. Height and position as well as the interaction of gender and position were found to be significant determinants of personal space.

Controlled experiments have continued to be extensively used in the more recent research in the context of biomechanics, especially prompted by the development of positional tracking and motion capture technology. Gerin-Lajoie et al. studied locomotory adaptations of walking trajectories in the presence of stationary and moving inanimate obstacles [57]. In their interpretation, the personal space serves as a protective zone intended to provide a sufficient time delay in case of appearance of unexpected obstacles and allows to plan ahead the locomotor trajectory. The experiment showed that the personal space has an elliptic shape, being larger in front of an individual and smaller on the sides. Although there is large amount of work on locomotory adaptations that allow walkers to preserve their personal space when circumventing static or moving objects (for example [92, 35, 146, 43]), few studies have investigated mutual collision avoidance between two human walkers.

Ducourant et al. studied time and space-related characteristics of interpersonal coordination and leader-follower interactions between pairs of walkers [39]. The resulting movements of leaders and followers were highly correlated, as anticipated, and that temporal delays were found between the movements of a leader and a follower.

Basten et al. used motion capture technology to describe collision avoidance behavior in face-to-face collision situations focusing on the effect of gender and height [14]. Four metrics were proposed to characterize the collision avoidance strategy of pairs of walking subjects: collaboration, clearance, anticipation and synchronization. In the conducted user study, pairs of male participants displayed lower collaboration than pairs of female participants. Similarly, pairs of tall participants collaborated less than pairs of short participants or short and tall participants. Clearance was smaller between pairs of males than between a male and a female participant. Anticipation and synchronization used to describe the spatial and temporal aspects of deviating from the straight trajectory while avoiding a collision were similar among the experimental groups.

Olivier et al. also used motion capture to study pairwise interactions between walkers, however in situations of crossing trajectories [108]. The authors proposed to describe walkers' collision avoidance behavior by a mutual function of their states called minimal predicted distance, which is the anticipated crossing distance at every moment of time. The analysis of the motion data showed that walkers only modified their walking trajectories when the minimal predicted distance became as low as 1m. Furthermore, collision avoidance strategies were found to be role-dependent for participants walking on crossing trajectories [107]: the person giving way made larger adaptations to their locomotory trajectory by modifying both their heading and walking speed whereas the person passing first only reoriented their path.

All the research on interpersonal interactions in spatial context leads to similar conclusions: the formation of interpersonal distances as well as collision avoidance are complex and highly collaborative interactions. Planning the trajectory in the presence of another human is also different from merely avoiding a moving obstacle. For example, participants of an experiment [147] used a different strategy for avoiding a mobile robot in comparison to avoiding another human, preferring to give way to the robot even when they were likely to be able to pass first. These are the subtle cues of body-to-body interaction that allow humans to correctly interpret and act accordingly to the intentions of the others. In VR, even the high fidelity systems cannot (at least at present) accurately model such subtle cues. This is the reason why human-human interaction, also in spatial context, may follow somewhat different rules in VR and needs to be studied.

3.2.2 Studies on collision avoidance in VR

The studies on spatial interactions and collision avoidance in VR mostly focus on comparison of people's behavior in the real world and in VR. The idea that locomotory characteristics of people avoiding collisions in VR may differ from those of the real world comes from a number of considerations. Most importantly, locomotion in general and collision avoidance in particular involve heavy reliance on multi-sensory feedback from our visual, vestibular and proprioceptive systems. In VR, however, visual cues produced by a rendering system differ from those that we normally receive from our visual system, leading to differences in perception of virtual space. Multiple studies have demonstrated that distances are underestimated in VR [85, 66, 94, 159], a phenomenon referred to as egocentric distance compression. Early research argued against a causal connection between egocentric distance compression and poor quality of graphics [143] or a limited field of view (FOV) of HMDs [34]. Some follow-up studies found a certain influence of restricted viewing conditions [157] and the quality of graphics [78] on distance underestimation, however not sufficient to fully account for the magnitude of the effect observed with HMDs. Recently, an experiment using the HMDs of a modern generation (HTC Vive and Oculus Rift CV) has been conducted, coming to the conclusion that egocentric distance compression persists regardless of the improvements in HMD's FOV and resolution [112], also shown in a comparison of a high-end HMD and Oculus Rift DK2 [33]. Apart from mismatching visual cues, the novelty of the VR experience (especially in what concerns walking in VR) and mistrust of the system may contribute to differences of locomotory patters in VR.

Studies of collision avoidance behavior of users physically walking in VR involve mostly static obstacles. In an experiment of Fink et al. [44] participants walked to a stationary goal while avoiding a stationary obstacle in matched real and virtual environments displayed in an HMD. Several metrics were used to assess the similarity of locomotory trajectories in the real and virtual conditions: obstacle clearance (the shortest measured distance between a participant and an obstacle), walking speed, path curvature and maximum deviation from a straight line. Larger obstacle clearance and maximum deviation from a straight line as well as slower walking speeds were observed in the virtual condition. The decrease in walking speed and an increase in obstacle clearance when avoiding a collision with a static virtual object compared to its real counterpart was also demonstrated in a CAVE-like immersive projective environment [3]. Furthermore, participants of the study kept larger distances to an anthropomorphic obstacle than to an inanimate object. In the anthropomorphic condition, obstacle clearance was higher when the orientation of the obstacle was from profile compared to a front position, confirming the elliptic shape of the personal space [39]. The transfer of the concept of personal space from the real world to VR has been demonstrated in several experiments. Bailenson et al. observed that participants of their study maintained larger distance from virtual humans when approaching their fronts compared to their backs and moved away when virtual agents tried to invade their personal space [8]. Similarly, Wilcox et al. observed negative reactions expressed by participants on the violations of their personal space in stereoscopic viewing conditions [156].

Few experiments have investigated situations where users walking in immersive VR avoid moving obstacles or other users. Olivier et al. showed evidence that users are able to predict whether they would collide with walking virtual humans or not and correctly choose their collision avoidance strategy [109]. The experiment was however conducted in a desktop environment. In a further experiment conducted in a CAVE environment users could accurately estimate the risk of a collision with an animated virtual character and used similar strategies for avoiding real and virtual collisions [106]. Locomotion in the experiment was performed with the use of a joystick. The same setup was used in two further studies where experimenters were investigating the role that global and local visual cues [88] and gaze [87] play in orthogonal collision avoidance tasks. Both experiments showed that global body motion cues were sufficient for collision avoidance. It has however been demonstrated that proxemics behavior can be significantly altered by locomotion techniques [76], with users being able to adjust interpersonal distances more precisely when they navigate a VE by walking as compared to joystick-based navigation.

Another study compared user behavior while crossing a virtual road and avoiding moving virtual cars in a CAVE-like system and with an HTC Vive [90]. Participants from the HTC Vive group demonstrated more skillful but also somewhat riskier road crossing behavior. Naturally, direct comparisons with the real-world locomotory behavior are difficult to obtain in a pedestrian simulator scenario. In a small-scale collision avoidance experiment in a CAVE environment, participants preferred a collaborative collision avoidance strategy while interacting with a virtual human, preferring it to step aside to allow them more space for passing but also willing to modify their own trajectories [25]. A similar preference for virtual agents actively avoiding collisions was demonstrated in a fully immersive setup with navigation by walking [99].

Previous research on collision avoidance in VR observed similarities between users' tactics in the real world and in virtual environments. However, the few studies that aimed at studying collision avoidance with another human used a virtual human and not another participant. There is evidence that people immersed in VR may behave differently depending on whether they interact with virtual humans or other users represented with avatars, avatars producing stronger responses than virtual humans [8, 46] (details are given in further sections). While collaboration is considered to be crucial for collision avoidance in the real world, no experiments investigated collision avoidance of two simultaneously walking users. A large portion of research presented in this thesis aims at closing the present gap in understanding collision avoidance behavior of two or more simultaneously walking users.

3.2.3 Mutual collision prevention in VR

This section concerns a very specific situation that can happen in multi-player VR where real walking is possible: if several players share a tracking space but do not see each other in the virtual environment a danger of collisions arises. This situation has been described by a number of researches but few solutions have been suggested so far.

A recent exploratory study in room-scale VR suggested several methods of visual feedback to prevent possible collisions [118]. The authors of the experiment considered a somewhat different situation that can lead to potential collisions: when several colocated players use teleportation to travel further in a VE than the limits of a real room allow, positions of their tracking origins shift in respect to each other. As a result, players' avatars stop being colocated with their physical bodies which leads to potential collisions. The authors tested a simple human-like avatar, a bounding box and a camera overlay displayed at the supposed position of a second user. However, the second user in the study was simulated and static just as virtual objects in the experiment of Fink [44]. The results showed that the avatar and the camera overlay allowed the fastest walking time towards the goal while the bounding box performed the best at preventing collisions, however being the least favoured method for the participants. Lacoche and colleagues investigated a similar strategy of using a ghost avatar and a cylinder grid displayed at actual positions of colocated users [79]. Alternatively, safe for walking zones were displayed on the floor of the virtual environment. All three methods were compared to a baseline condition where the physical workspace was split into two zones, each one assigned to one of the test users. In a user study, the ghost avatar and the grid cylinder were as effective in preventing collisions as the space separation and were positively estimated by participants. Techniques introducing visual notifications in dangerous areas have been wildly used to ensure that walking users do not collide with static obstacles. For example, limits of the available walking space were presented as semi-transparent walls [45], a barrier tape [32] or a 3D grid [40]. These visual representations of the tracking space boundary only become visible when a user approaches sufficiently close. Current commercial systems (HTC Vive, Oculus Rift) also offer an option of displaying the limits of the tracking area.

Redirected walking-based methods for collision prevention do not require any visual representation for colocated users and do not limit the freedom of physical movement within the tracking space. Originally, redirected walking (RW) was invented to steer walking users away from the boundaries of the tracking space thus allowing unlimited walking possibilities [113, 137]. RW techniques are based on adding small rotational, translational or curvature gains to the simulated viewport of a walking user. As a result, the user walks in curves within the tracking space although their virtual path may look like a straight line. The manipulation can remain unnoticed provided a sufficiently large tracking space and small added rotations [138]. Alternatively, the layout of the virtual environment might contain overlapping spaces to fit the size of the available tracking space [140, 151].

When several users are immersed within one tracking space, it is proposed to use RW to steer them not only away from the physical boundaries but also away from each other [69, 5]. It has been demonstrated that a RW-based approach would be a good alternative to separating the available tracking space into two separate zones for two immersed users [5]. Although RW-based collision prevention techniques have powerful theoretical potential their robust implementation remains a hard problem. The proposed multi-user RW algorithms rely on the knowledge of the VE and the prediction of the future trajectories of users. Whereas short-term path prediction based on head tracking data has been demonstrated in several papers [100, 101], robust long-term prediction is more difficult to achieve. Approaches based on navigation meshes or connectivity graphs have been suggested for partially occluded environments [4, 162]. These methods have been shown to work for static and somewhat occluded environments.

In the few publications exploring the potential of RW for multi-user collision prevention, computer simulations were used to test the efficiency of the method. The success of the method is then measured in the amount of occurrences where simulated walking users had to be stopped to avoid a collision that the method could not resolve. It was suggested to stop walking users and blank their HMDs if the case of failure of the collision prevention method [69]. However, it is unclear whether such last-minute collision prevention measure would work well in a real application, especially if users are immersed in a highly interactive and involving scenario. Apart from being potentially unsafe, stopping users and blanking their HMDs mid-experience could lead to confusion and breaks in presence [54]. We have not integrated RW techniques in our multi-user experiments.

3.3 Embodiment and copresence in VR

3.3.1 Embodiment through virtual avatars

The notion of *embodiment* has been used in two senses in the research literature. The first use refers to provision of a user with a virtual body representation that serves as a substitute for their own body in VR. According to Benford and colleagues who first addressed the issue of graphical representation of users in collaborative virtual environments (CVEs), the goal of a virtual body is to represent a user to others but also to themselves [16, 17]. The primary goal of an embodiment was to convey user's presence in the CVE, to specify their location and to manifest their identity [17]. It was also suggested that graphical representations of users would implement behaviors conveying their current activity, availability and the degree of presence in the system, thus providing an adequate substitute for the rich spatio-social context provided by our biological bodies in everyday life. Viewpoints and actionpoints were suggested as means of representing the areas of a VE where the user is looking at and where their manipulations of the VE takes place. Design ideas introduced in these papers were implemented in representations of users used in DIVE [63] and MASSIVE [61], the early multi-user VR platforms. Different graphical representations ranging from simple geometric "blockies" to more human-like virtual bodies controlled by 3D input devices were used to represent users interacting by means of different equipment.

Current research literature usually uses the term *avatar* or *self-avatar* to refer to a virtual body controlled by a human user. *Virtual agents* are virtual bodies whose behavior follows a computer algorithm, although they may look human-like and indeed be designed to display social behaviour and possess a "personality". Virtual agents are often used in experiments on social aspects of VR that are addressed in the next section.

The second and more commonly used meaning of the term embodiment, or the sense of embodiment (SoE), is the illusion of occupying and using a virtual body just as if it were one's own biological body. The research community agrees that the illusion of embodiment is composed of three elements: self-location, agency and body ownership [73]. *Self-location* refers to the location of self and is usually the location of the body, physical or virtual. It is however not necessary to perceive the self as being located within the body. Reports of out-of-body experiences demonstrate a case where the location of the self differs from the location of the body [41]. Similarly, it is possible to become embodied into a character that is not colocated with one's physical body, for example in desktop-based computer games. Most VR setups with head tracking however use avatars that are colocated with users' physical bodies as seeing a virtual body from the first-person perspective alone is a powerful stimulus for the embodiment illusion [91]. *Agency* refers to the attribution of authorship of our actions to ourselves and is produced as a result of a match between the predicted and actual sensory outcomes [38]. In VR, agency is strongly connected to motor control, i.e. being able to move the virtual body or its parts at will. Visuomotor synchronization (having both physical and virtual bodies move synchronously) produces especially strong contribution to the embodiment [75, 13]. It is however possible to experience agency over actions of a virtual body seen from the first perspective even without performing these actions in reality; for example, seated users can experience the illusion of walking [74]. *Body ownership* is a more complex phenomenon for which three factors are important: acceptance of the virtual body as one's own body, control closely related to agency and change in self-perception [114]. The Rubber Hand Experiment is the most famous demonstration of the induced feeling of ownership of an artificial body part [26]. In the experiment, a demonstrator simultaneously strokes a rubber hand placed in front of a participant in a position in which their real hand could be while their actual hand is occluded by a screen. When a threat is introduced to the rubber hand the participant moves away their real arm. The effect is attributed to multisensory integration between synchronous tactile and visual signals and can also be observed in VR [125, 161]. Moreover, it has been shown that users immersed into a virtual environment can experience ownership of additional body parts, for example a tale [139] or a third arm [80] as well as of virtual bodies with a different gender [127], age [12], race [111], or body shape [102]. Visuotactile and visuomotor synchrony are bottom-up processes that contribute to both agency and body ownership. Agency and ownership are correlated in a complex way; it has however been shown that they represent distinct cognitive processes and can be dissociated [70]. Top-down factors like similarity of form and appearance also contribute to the illusion of body ownership [91], prompting the use of virtual mirrors to induce and study the sense of body ownership, especially in non-immersive or semi-immersive systems [82, 83]. Detailed questionnaires aimed at assessing the subjective judgments on self-location, agency and body ownership have been developed [59, 114].

In single-user VR applications, providing a user with an avatar is not mandatory although multiple benefits of having a virtual body in connection with the illusion of presence and the perception of the virtual space have been demonstrated [95, 96, 134]. In shared VR experiences, however, users need to have avatars, at least with a goal of seeing each other. Furthermore, successful embodiment is considered to be one of the building blocks of copresence (as we shall see in the next section) which is the goal of creating shared virtual experiences.

3.3.2 Theory of copresence and social presence

While the goal of single-user VR experiences is in eliciting the sense of presence, i.e. the sense of "being there" in the virtual place, shared virtual environments are meant to produce the illusion of "being there together" simulating not only presence in a place but also in a company of others.

Slater and colleagues define presence as “a state of consciousness, the (psychological) sense of being in the virtual environment” [131] and classify presence into personal presence and shared presence (or copresence) [129, 124]. Personal presence relates to the subjective feeling of “being there” in the virtual environment, leading to an experience of “places visited, rather than images seen” [131]. Shared presence (copresence) has two aspects: that of feeling that the other participants in the VE actually exist and are really present in the environment, and that of feeling of being a part of a group. Biocca and colleagues [22, 21] make a distinction between these two aspects and suggests the term *copresence* to refer to the sense of perceiving others in the same place and *social presence* to refer to the feeling of being in a social situation with others. According to Blascovich, social presence is "a psychological state in which the individual perceives himself or herself as existing within an interpersonal environment” and “the degree to which one believes that

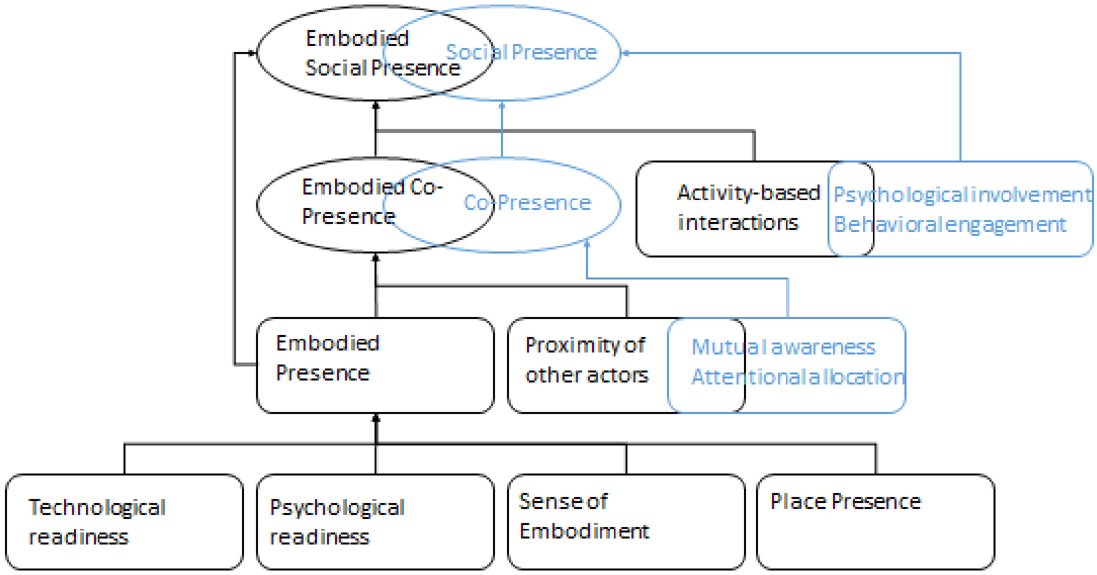

Figure 3.2: Components leading to Embodied Social Presence theory of Mennecke and colleagues (adapted from [93], in black) and social presence according to Biocca and colleagues (adapted from [22], in blue)

he or she is in the presence of, and dynamically interacting with, other veritable human beings "[23]. In this dissertation, we will follow this choice of terminology.

Research has shown that copresence and social presence are connected to embodiment, as social presence can only begin to occur when people feel that they have an access to "intelligence" of others conveyed by bodily action and expression [20]. Biocca and colleagues [22] state that copresence is based on sensory awareness of an embodied other, and the body of the other, however its representation may vary, is a key medium for the sense of copresence. Copresence contributes to social presence which however includes more factors. These further factors are psychological involvement and behavioral engagement [22]. The authors suggest that these factors can be further broken down into smaller contributing factors. Their factor analysis showed that copresence is affected by mutual awareness and attentional allocation, psychological involvement is based on mutual understanding and empathy, and behavioral engagement comprises behavioral interdependence and mutual assistance. A detailed questionnaire is suggested to assess these presented components. Mennecke and colleagues present their own theory of Embodied Social Presence (ESP) [93]. In developing the ESP theory, the authors explore the role that shared space and embodiment play in the illusion of presence and copresence. The ESP theory suggests that communication in a virtual environment builds on the embodied sense of self and is realized through coparticipation in a particular context that is partially defined by the symbolic meaning associated with the shared space. To experience ESP, the actor (in context of our research, a user of immersive VR) first needs to achieve the illusion of embodied presence and copresence at sufficient levels. When multiple users share a virtual space they have an opportunity to experience embodied copresence, provided that the virtual environment produces adequate visual, audio and possibly other stimuli. According to the authors, six factors are necessary for ESP: technological readiness of the system, psychological readiness of the user, virtual embodiment, illusion of being in the virtual place (place presence), proximity of other social actors and activity-based interactions between actors. Figure 3.2 illustrates how theoretical dimensions of copresence and social presence suggested by Biocca and colleagues [22] can be matched with the process of the elicitation of ESP proposed by Mennecke and colleagues [93]. Proximity of other actors in the shared virtual space leads to mutual awareness and attentional allocation that are necessary for invoking the sense of copresence. Activity-based interactions between embodied users can create psychological involvement and behavioral engagement that contribute to the sense of social presence alongside copresence. It is worth noticing that psychological readiness that is necessary for the illusion of embodied presence according to the ESP theory is likely to be important for establishing psychological involvement and behavioral engagement through activity-based interactions.

While the ESP theory considers the necessity of place presence for the elicitation of copresence and social presence several studies have demonstrated more complex and bidirectional connections between place presence, copresence and social presence. For example, in the study of Casanueva performed in a desktop VE built on DIVE [31], when users experienced stronger copresence fulfilling a highly collaborative task they also reported a stronger sensation of place presence than in a condition with low collaboration and consequently lower copresence. The study however used a single quantitative score for copresence and social presence related questions making it difficult to analyze further the individual impacts of copresence and social presence. In a later study, place presence, copresence and social presence have been found to affect user satisfaction in a virtual world [27], social presence having the strongest effect. Moreover, immersive tendencies of participants were related to the experienced levels of place and copresence but not to social presence. We will have a closer look at copresence in our study described in Chapter 6.

3.3.3 Effects of avatars on embodiment and copresence in shared virtual environments

One of the main questions that arise when designing shared VEs is which kind of avatars to provide for users, in terms of both appearance and behavior. Research has focused on the effects that an avatar produces for a user having it as a self-avatar and for others that see it as a virtual representation of another user.

Appearance of avatars, and especially its realism has received a lot of attention in embodiment research. Personalized avatars created with a fast full-body scanning procedure [1] have been shown to have a positive impact on embodiment [153]. In an experiment comparing a machine-like, a cartoon-like and a human avatar with full freedom of body motion [86] the non-human avatars elicited a slightly greater sense of body ownership that the human one, possibly due to the Uncanny Valley effect [97]. In another experiment, however, realistic avatars evoked a stronger acceptance in terms of body ownership component of embodiment [83], although some indications of a potential Uncanny Valley were found.

A recent experiment studied whether the sense of embodiment is greater in a shared VR experience compared to a single-player one [52]. The results indicate similar levels of experienced embodiment in both scenarios but greater engagement and enjoyment produced by the shared game.

Appearance of avatars has also received a lot of attention in the research on social presence. An experiment of Nowak and Biocca [105] examined the influence of anthropomorphism of avatars and perceived agency on presence, copresence, and social presence in a VE. Participants responded socially to avatars as well as virtual agents, and interacting with less anthropomorphic images resulted in stronger perceived copresence and social presence. The authors conclude that interacting with very realistic images set high expectations leading to a lower copresence. In an experiment of Bailenson et al. avatar realism was found to increase copresence but decrease self-disclosure [11]. A recent large-scale study with a virtual agent of varying rendering style has shown that realism of appearance had a positive impact on user experience [165].

Garau suggests that behavioral realism of an avatar can be more important for social interactions than photo-realism of avatars [53]. For example, an avatar displaying adequate gaze behavior could induce a stronger effect of social presence. Similarly, when interacting with virtual agents users responded socially if virtual agents themselves displayed responsiveness [55], the effect however being less pronounced among the experienced computer users. Another study however showed that avatars with gaze behavior inferred from a participant's voice were only effective in increasing social presence if they had semi-realistic appearance [56]. Lower-realism avatars were adversely affected by realistic gaze, showing an interaction effect between the realism of appearance and behavior. Bailenson et al. came to a similar conclusion that form realism and behavioral realism of avatars should be considered in connection with each other, and a large mismatch in appearance and behavioral realism leads to low copresence [9]. Steed and Schroeder argue that requirements to the avatars used in shared VEs should be formulated taking into account spatial context and task that users are solving in the virtual world [135]. Avatars need to be able to convey user movements in tasks requiring joint attention while their detailed facial expressions might be less crucial. In VEs primarily designed for social interaction, on the contrary, the details of avatar appearance may have more importance.

Subjective questionnaires are commonly used to assess the levels of presence, copresence and social presence experienced by users of immersive VEs. The weaknesses of self-report measures have however been discussed by multiple researches and using behavioral measures in addition to self-report has been suggested [7, 19, 9]. In a longitudinal study of collaborative VEs, self-report measures did not change over time not reflecting important changes in user behavior [10]. Bailenson and colleagues suggested to use interpersonal distances maintained by users as behavioral markers of social presence connected to the notion of personal space [7]. Personal space has since been used in multiple experiments involving avatars or virtual agents. For example, "personality" of virtual humans has been shown to affect the size of the personal space of an immersed user: users stopped approaching angry-looking virtual agents at larger distances [24]. It has also been shown that users can be primed to experience stronger social presence. In a recent experiment, participants who witnessed two virtual agents interact with each other before engaging into the interaction of the experimental task reported higher scores of copresence [37].

3.3.4 Effects of the VR setup on user experience and copresence in shared virtual environments

Schroeder, referring to copresence simply as "being there together", states that both presence and copresence depend on the extent of mediation [120] and gives an example of how an immersed person would sometimes easily walk through an avatar of another person and sometimes maintain conventional interpersonal distance, depending on the type of the system used and on how habituated the person is to this type of mediation. In a paper on the taxonomy of copresence, Zhao distinguishes between two ways of understanding it: as a mode of being with others and as the feeling of being together with others [164] (in the sense of social presence defined by Biocca and colleagues [21]). The mode of copresence is defined by a combination of physical and electronic proximity at the site of interaction, be it virtual or real. Different forms of technological mediation thus reflect different modes of copresence, and shared VR is one of them. According to Zhao, the mode of copresence can influence the extent of the experienced social presence.

The research on the effect of mediating technology usually compared different types of systems enabling shared VEs. Slater et al. describe an experiment with three participants at a time in which participants had to solve riddles in VR and during a real meeting [126]. Two participants in each team performed the task in a desktop environment and the third one with a HMD. The immersed participants tended to be seen as group leaders even though other participants from a group did not know which type of technology the leader used. Schroeder and colleagues demonstrate that using an immersive CAVE-like system for collaborative solving of a spatial riddle (Rubik's cube) can be just as good as collaborating on the task face-to-face [121]. The same task however takes longer to solve is a desktop system is used.

While physical proximity is one of the important aspects of mediating technology, almost no research has addressed its possible impact on copresence and social presence. One study that directly compared a colocated and a distributed VR setup in a two-player game setting found no differences in social presence experienced by players [89]. The study used a Google Cardboard as a display device and participants were playing seated rather far away from each other, although in the same room. In the distributed condition, the same pair of participants were playing in neighboring rooms while talking via intercom. Seeing that no physical movement through the shared or individually occupied physical space was possible, the mode of communication comprised the major difference between two experimental conditions.

Physical colocation or the absence of thereof gains importance in situations where users are not merely stationary within a shared VE but actively and simultaneously move through it. Whereas the connection between walking through VEs and the sense of presence has been established [131] virtually no research has addressed the question of how users walking in the same or different physical space experience copresence in a virtual space. We intend to get a step closer to understanding how people collectively navigate walkable virtual spaces by studying the experience of users not only "being there together" but "walking there together".

CHAPTER 4

ImmersiveDeck: A Walkable Large-scale Multi-user VR Platform

This chapter describes ImmersiveDeck, a multi-user platform that supports natural walking and intuitive interaction in large physical and virtual environments. ImmersiveDeck provides a base platform for further experiments on walking in multi-user VR setups described in this dissertation.

The chapter starts with an overview of the developed platform where we describe various requirements that were taken into account while designing the system and the resulting hardware and software solutions that we decided to use based on these requirements. Afterwards, the implementation details of user and object tracking are presented, followed by the description of our networking solution. Then, we describe a scenario that illustrates the possibilities of using the developed platform for the creation of non-shared VEs. The chapter is concluded with the details of the initial evaluation of the system and the overview of its evolution.

4.1 System overview

When we started designing ImmersiveDeck, the vision that we had was of a flexible, scalable (in terms of walking area and the number of users) and a very high-fidelity VR system that would allow natural and intuitive interactions. Our goal of achieving high fidelity meant that users would be enabled to explore VEs by walking, equipped with HMDs providing the largest available FOV, and that they would have avatars repeating their movements in VR. Functional and non-functional requirements reflecting our main design goal were formulated.

4.1.1 Requirements

- *Users should be able to move freely in a large area.*

  This requirement means, first, an untethered setup. Secondly, by free motion we understand the possibility to make the full range of movements, from upright walking to running to sitting down or even crawling on the floor. Therefore, our tracking solution should be able to handle mutual and self-occlusions.

- *The system should support at least three users at a time.*

  The system should implement the distribution of user motion and interaction data across several clients.

- *Users should experience as little latency as possible.*

  The change of the viewport position and the result of interaction with nearby objects or other users should be seen immediately.

- *The cost of the system should be low.*

  We intended to use easily accessible and affordable hardware.

Additionally to these core requirements, several other desirable features were considered at the design stage.

- The system should allow voice communication between users.
- The system should allow collaborative interaction with virtual objects.
- The system should include the possibility of having haptic feedback while interacting with some virtual objects.
- If other design decisions allow, the system should be scalable and portable.

4.1.2 Software and hardware design

This section explains our choices regarding the components of the developed platform. While each component is described separately, the most important elements such as user tracking, networking architecture and the rendering engine had to be considered simultaneously as the choice of one component put limitations on the possible choices regarding the other ones. Thus, our design process entailed the analysis of combinations of elements rather than of separate components. Finally, our choices were strongly influenced by the considerations of available input and output devices as well as computational resources. For example, the use of inside-out head tracking would not be possible without the availability of small and light-weight cameras. Similarly, we had to equip users with portable computers to achieve an untethered setup since wireless solutions for HMDs did not exist (and cannot be used in colocated multi-user VR until now).

User tracking

The requirement of free movement with full-body tracking combined with the need of a low-cost solution has limited the choice of tracking solutions. Most tracking systems providing full-body tracking work with an array of cameras installed on the perimeter of a tracking space that track retro-reflective markers attached to a user. Such systems are usually expensive as the number of cameras required to cover the designated tracking area grows with the increase in the size of a desired workspace. Yet occlusions are a big problem for outside-in tracking systems, especially if users come very close to each other.

Based on the above considerations, we opted to use an inside-out tracking solution for head tracking. Our head tracking solution is based on marker tracking. Initially, we used marker tracking provided by Metaio[1] that we largely extended for our purposes. A small tracking camera with a fish eye lens is attached to a user's HMD tracking markers on the ceiling of the tracking room. Body tracking is done separately with the use of an inertial motion capture suit. The detailed implementation of user tracking is described in Section 4.2.

Game engine

We chose to implement the system on the basis of Unity3D (version 4.x at the time of development) game engine, as it is commonly used the development of virtual worlds and provides a real-time rendering engine as well as an easy integration of multiple input devices.

Architecture

ImmersiveDeck has a client-server architecture where each client's application runs on a separate machine (laptop carried by a user). Input tracking data of each user is applied locally to the virtual camera and the avatar in the user's application. Motion and interaction data is then synchronized to the applications of other users through a central server. The server and clients communicate over a local wireless network.

This particular network architecture was chosen as a result of multiple considerations. As already mentioned, users had to be provided with laptops to render the images for HMDs. If an outside-in tracking system is used, it is possible to collect tracking data of all users on one workstation and send out tracking data for each user to their respective rendering application where it is applied to the virtual camera and the avatar. Such a solution would provide a simpler networking architecture since data would be send only in one direction, from the server (a tracking workstation) to clients. In this case, however, possible network delays could result in the perceivable lag in the pose updates of a user's virtual viewport and avatar. This was the main argument for tracking input to be applied directly in the client application. In addition, our tracking solution could

[1] Metaio GmbH was purchased by Apple in 2015. Metaio SDK used for tracking of fiducial markers is no longer available.

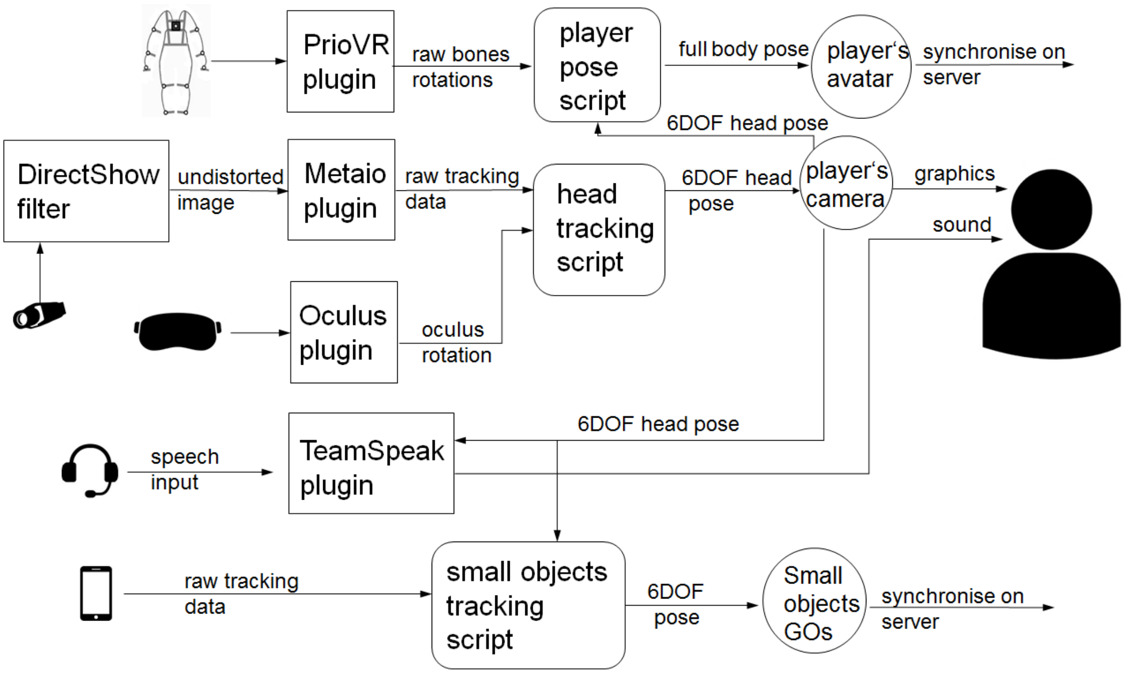

Figure 4.1: Dataflow in a client ImmersiveDeck application. Circular components represent GameObjects that can be manipulated by a player: their own avatar and virtual camera and tracked physical objects.

only be used in circumstances of a tracking camera being attached directly to the client's laptop. In our solution, since tracking and rendering are performed locally each user receives updates of their own movements with a minimal latency.

Figure 4.1 illustrates the data flow in each client application of ImmersiveDeck. Tracking input from each hardware device is received by Unity3D through an integration plugin. Input data from the used hardware is processed on a client and used to position Transform nodes of Unity3D.GameObject containers (further referred to as GameObjects for simplicity) corresponding to the player's avatar and virtual camera. The full 6DOF pose of the player's avatar and the poses of virtual objects that are being manipulated by this player are then synchronized to the server. The details of data distribution are described in Section 4.4.

Object tracking

ImmersiveDeck implements tracking of selected physical objects used for haptic feedback during interactions with virtual objects. Our implementation of object tracking is marker-based. We implemented two different approaches to object tracking. Objects of large size, such as furniture, are tracked by a camera attached to the server machine and their positions are streamed to clients. Since such objects are likely to be heavy or bulky, it is improbable that users would move them quickly enough to notice possible lag in position

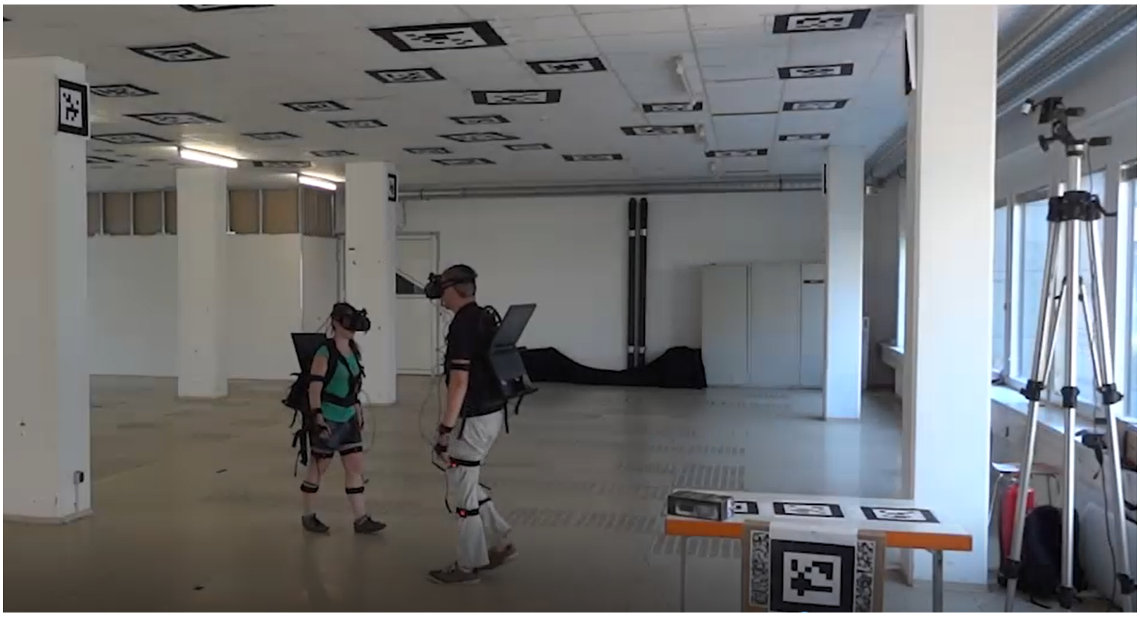

Figure 4.2: ImmersiveDeck being tested by two users.

updates. Smaller physical objects are tracked by smartphones fixed in front of users' HMDs and connected to users' laptops. As a result, these objects can be manipulated when they are in the tracking range of a user's mobile phone. The tracking data of these objects is applied to the pose of the corresponding virtual objects directly on the client, in the same way as user tracking data. The detailed implementation of object tracking iis described in Section 4.3.

Voice communication

Users of ImmersiveDeck are able to talk to each via headset microphones. We use TeamSpeak[2] to enable voice communication. Further details can be found in Section 4.4.

Hardware

The work on the development of ImmersiveDeck started in 2014, about at the same time as Oculus Rift DK2 became available. It was the obvious choice being the only consumer-priced HMD with a wide FOV.

Each client application runs on a XMG Schenker laptop equipped with two NVIDIA GTX 980M graphics cards and an Intel Core i7 quadcore processor.

The server runs on a PC with an Intel Core i7 processor.

[2] https://www.teamspeak.com/

The wireless network connecting the server and clients is set up with an ASUS RT-AC87U router using the 802.11ac protocol connected to the server. The router supports simultaneous dual band operation, which results in a maximum throughput of 1.73 Gbps on 5 GHz band (4 x 433 Mbps streams). The client laptops have Intel Dual Band Wireless-AC 7260 cards.

Head tracking is implemented with the use of IDS cameras UI-3251LE streaming monochrome images via a USB3 connection at 60 fps. The maximum resolution of the camera image is 1600 x 1200 pixels. Attached to the cameras are fish eye lenses with 175° to 190° FOV.

Body tracking is done with the beta version of PrioVR motion capture suits with 11 inertial sensors.

Finally, noise-canceling headsets (ASUS Vulcan Pro and Logitech G35) are used for voice communication between players.

Figure 4.2 shows two fully equipped users during a test of the platform.

4.2 User tracking

4.2.1 Head tracking

Our tracking is an extension of a commercially available marker-based method from Metaio. We have chosen to use ID markers, since they offer the most robust and fastest tracking compared to different marker types that we tested. We have experimented with marker sizes and the map density. The resulting setup uses about 80 square markers of the size 550mm distributed evenly on the approximately 200 m^2 (30x7 m) large ceiling area. The height of the ceiling in the original tracking lab is 3.2m. The tracking lab with the markers arranged on the ceiling can be seen in Figure 4.2. The tracking data of each marker is streamed directly into Unity3D using the Metaio plugin for Unity. This data is then processed to calculate the final camera pose.

Our inside-out tracking solution allows users to be tracked in an arbitrary large area as long as it is covered with sufficiently many markers. Furthermore, fish eye lenses allow to track the marker map even when the camera is not pointing at the ceiling directly. Finally, there is no occlusion problem even when users are very close to each other.

Image Undistortion and Lens Calibration

Fish-eye lenses produce very wide angle images, so markers can be seen in the camera image even if a user is looking down or lying on the floor. These images are however strongly distorted and cannot be used as tracking input as-is. We have implemented a separate software component for runtime undistortion of camera images.

First, we obtain intrinsic camera calibration parameters by using the OCamCalib toolbox by Scaramuzza et al. [117] whose method is especially suitable for omnidirectional and

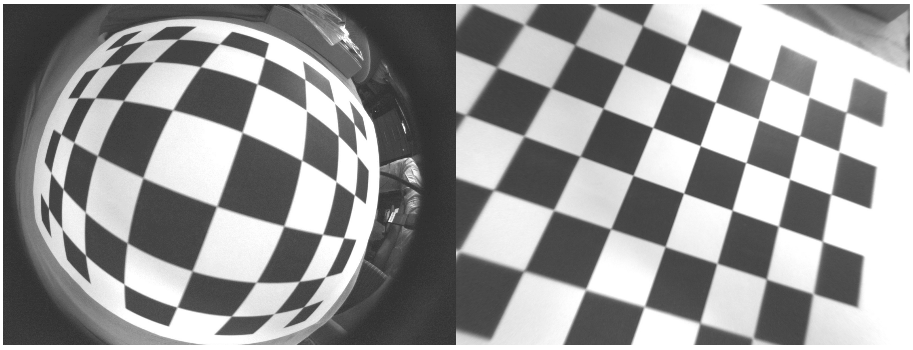

Figure 4.3: Results of the lens undistortion procedure. Left: an input image produced with 175° lens FOV, right: the output of the undistortion algorithm. The undistorted image is clipped.

fisheye lenses. To do this, we take 15 to 20 pictures of a 12 by 9 field chessboard pattern with square 80 mm patches covering as much of the view as possible at close range from varying angles. Calibration parameters are then obtained from these images using the method of Scaramuzza et al. [117] and saved in a calibration file. A lookup table in created in which every pixel in the undistorted image is matched with the corresponding pixel position in the original distorted image.

The runtime undistortion algorithm is implemented as a DirectShow source filter and uses GPU-accelerated methods in OpenCV for efficient processing. The DirectShow filter loads parameters for the omnidirectional camera model from the calibration file upon startup.

At runtime we use OpenCV to retrieve the camera image of our IDS camera. Subsequently, the image is transferred to the GPU where we use CUDA to efficiently remap the image using the lookup table. Once downloaded from the GPU the undistorted image is made available to the tracking software. An example of an input and the corresponding output image of the undistortion procedure are shown in Figure 4.3.

Undistortion is not the only necessary calibration procedure. Metaio requires another set of calibration parameters for successful tracking. To obtain them, we take another 7 to 10 pictures of the same chessboard pattern, apply the above undistortion method to them and then use the camera model of Zhang [163] on the undistorted images. The resulting calibration parameters are applied at runtime by the Metaio tracking process directly.

Marker map generation

The user's head camera is tracked relatively to the map of markers attached to the ceiling. Distances and rotations of the markers relative to each other need to be known in order to use tracking data from all visible markers for the final camera pose calculation. As measuring those distances and rotations manually or arranging markers on the ceiling precisely does not seem feasible, we have developed a scanning procedure that calculates

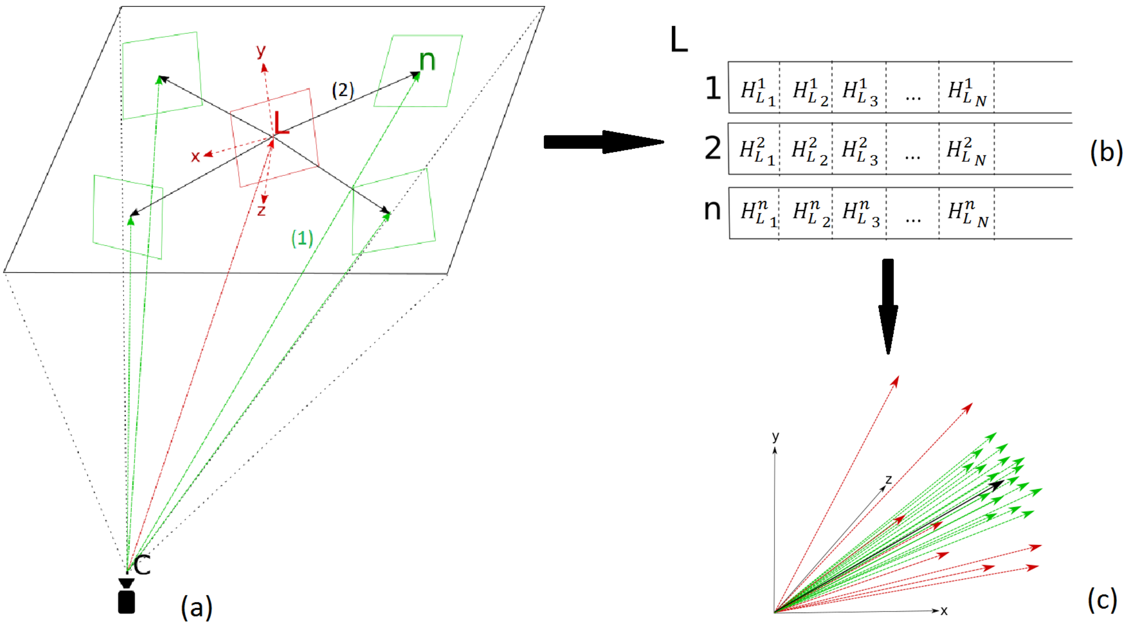

Figure 4.4: Calculations during the scanning procedure: in each frame during data collection (a) and in the post-processing step (b), (c).

relative positions and rotations of the markers. For a successful scanning, the markers need to be attached to a flat surface and must not be mechanically distorted.

During the scanning, the camera is kept pointing directly at the ceiling and is moved along the whole tracking area. We use a wheeled platform with a tripod with the camera mounted on it to perform the scanning. One of the markers is chosen to be the center of coordinates for the marker map.

The calculations of the scanning procedure are illustrated in Figure 4.4. In each camera frame, the marker that is the closest to the image center is chosen as a temporary coordinate center. This marker L is shown in red in Figure 4.4 (a). Raw tracking data provides poses

$$H_C^n = [R_C^n T_C^n] \tag{4.1}$$

of all recognized markers in camera coordinates, where H_C^n is 4x4 pose matrix, R_C^n is the rotation matrix of the marker n in the camera coordinates and T_C^n its translation vector. An inverse transformation

$$H_L^n = [R_L^n T_L^n] = \left[R_C^L T_C^L\right]^{-1} [R_C^n T_C^n] \tag{4.2}$$

with the coordinates of the central marker in the camera frame is calculated for each marker to obtain its coordinates in the coordinate frame of the temporary center L.

Moving the camera along the tracking space results in a set of stored temporary centers together with sets of coordinates of neighbouring markers in respect to them as shown

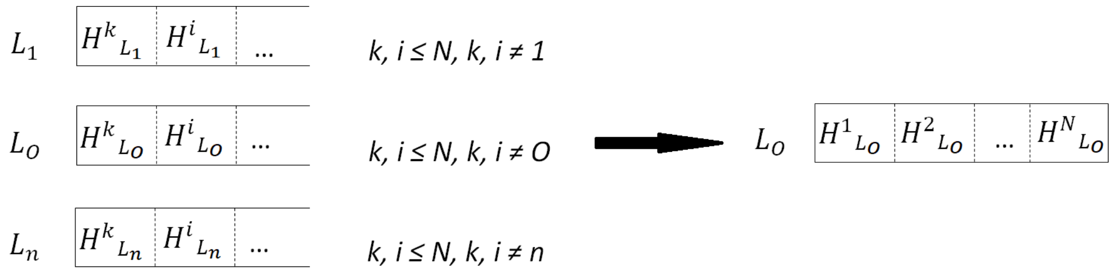

Figure 4.5: Calculations .

in Figure 4.4 (b). After real-time tracking data has been collected, the final local pose relatively to a temporary center is calculated for each marker in a post-processing step. This is done by finding the median pose from the stored values and calculating an average pose from the values around the median. This calculation is illustrated in 4.4 (c). Here, all green and red arrows represent all collected poses of marker n in the coordinate frame of marker L. These collected poses are slightly different because of the possible tracking jitter and the finite precision of calculations. Poses represented by red vectors strongly deviate from the central tendency so they are discarded as outliers. Poses represented by green vectors are taken as input for the calculation of the 3D-median pose (shown as a black vector).

After this step, we have a set of local origins, including the market L_O chosen to be the global origin of the marker map, each with a corresponding list of markers the coordinates of which are known in the local origin. The breadth-first search is performed on the resulting set of local coordinate centers to calculate the poses of all Nmarkers relatively to the marker L_O, the global coordinate center. Figure 4.5 illustrates this process.

Marker Tracking

The tracking process uses the marker map generated in the step described above. As in the scanning procedure, the tracking process streams poses $H_C^n = [R_C^n T_C^n]$ of each marker n visible in the camera frame relatively to the camera. Known markers' poses $H_O^n = [R_O^n T_O^n]$ in the marker map are used to calculate the position

$$T_O^C = T_O^n + R_O^n [T_C^n]^{-1} \tag{4.3}$$

and rotation

$$R_O^C = R_O^n [R_C^n]^{-1} \tag{4.4}$$

of the camera in the coordinate frame of the origin O of the marker map. This way, a number of camera poses are calculated for every camera frame. The final pose is a result of median filtering.

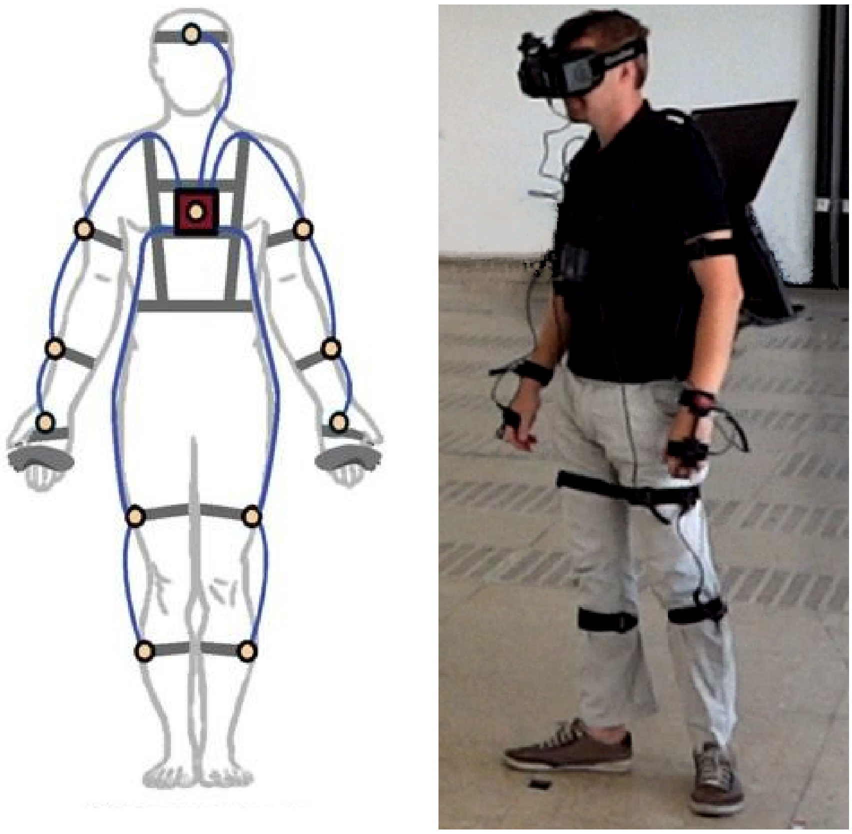

Figure 4.6: Left: placement of inertial sensors composing the PrioVR motion capture suit. Right: a user wearing the suit.

Only the translational component of the obtained pose is used for setting a user's viewport position, while the rotation is provided by the Oculus Rift tracking data. The global coordinates of the marker map define the final position of the player avatar in the virtual world.

4.2.2 Body tracking

A PrioVR suit consists of 12 inertial sensors placed on a user's limbs which are connected to a hub on the user's chest. Figure 4.6 demonstrates the placement of the sensors.

Relative rotations of the sensors are streamed from the hub to a wireless USB receiver connected to the laptop with the transmission latency of 5-8 ms. This data is streamed into Unity3D by the integrated SDK provided by PrioVR in each frame, where it is used to recalculate a player's avatar root position and joints' rotations in accordance with the already calculated head camera pose.

In the beginning of each VR session, each suit needs to be calibrated. For this purpose, the user takes a pre-defined pose (for example, T-pose) that corresponds to the default pose of their avatar. The calibration procedure takes about three seconds. It can be repeated at runtime.

4.2.3 Tracking data fusion

Two GameObjects need to be correctly positioned for each user: the virtual camera and the user's avatar. The avatar consists of multiple hierarchically arranged GameObjects. Positioning of the root GameObject leads to the whole avatar being moved with its parts maintaining relative positions and rotations to the root.

The player's camera is kept separately and not attached to the avatar GameObject. The camera is positioned directly according to the position calculated by the tracking script and the rotation of the HMD. Rotations of the bones in the hierarchy of the avatar are determined by the PrioVR script that reads in body tracking data. During the fusion of the tracking data, the root of the avatar is repositioned to coincide with the position of the player's body.

To do this, an empty (not containing a mesh) GameObject at the position of an HMD is introduced into the skeleton of an avatar. If the avatar is positioned correctly according the camera pose, the global position of this auxiliary GameObject should be the same as the camera position. At the same time, the local position of the HMD GameObject within the skeleton is known from the body tracking data. This allows us to traverse the skeleton to calculate the global position of the root that would bring the HMD GameObject to the position of the camera.

4.3 Object tracking

In our system, users can have haptic feedback while interacting with some of the virtual objects present in the scene. We achieve this by having real objects tracked in the real environment and overlaying them with virtual representations. Object tracking also allows users to perform collaborative interaction with virtual items in combination with haptic feedback. For example, two users can simultaneously lift a real tracked table and see themselves lifting a virtual table. Or they can pass each other a smaller item (we use an example of a small box).

We divide objects into two categories depending on their size (we call them large objects and portable objects). Different tracking methods are used for each of the two categories. Further, different network distribution approaches are used for large and portable virtual objects.

4.3.1 Large objects

Big objects have typical furniture sizes. A table and a chair are used in our test setup.

For tracking of big objects, we use multi-marker tracking based on the Metaio solution. Several ID markers are attached to the table and the chair and are tracked by a camera

Figure 4.7: Portable (a) and large (b) tracked objects and the example of users interacting with them.

connected to the server machine. We used an IDS camera with the resolution of 2048x2048 pixels and markers of the size 287x287 mm. The camera was attached to the server machine and placed statically at a know position in the tracking room.

Positions of the corresponding virtual objects are calculated on the server and streamed to clients. With this type of setup, the area in which users can interact (touch and move) with big objects is limited by the field of view of one or multiple cameras used for tracking. In our case, this area was about 3x3 m. Figure 4.7 demonstrates the examples of tracked objects used in our test setup.

4.3.2 Portable objects

Portable objects are of a size that users can easily hold them in their hands, pick them up and transport them across the tracked space.

The tracking of smaller objects is implemented locally with an Android smartphone mounted on a user's HMD. For our prototype we employed the Huawei Ascend Y300, which only weighs 130 g and provides a 1 GHz dual-core CPU and a 5 MP camera.

An application based on the Metaio SDK for Android was developed. It runs directly on the smartphone, taking the most of the required processing load for the additional tracking process from the client laptop. The employed tracking is marker-based. In this case, we chose image markers since they are robust to occlusions caused by holding the objects. The size and look of the markers can be chosen appropriately for the tracked object. For our prototype we designed six markers for a box with the dimensions of 244 x 142 x 84 mm. Rendering on the smartphone is disabled to achieve a tracking rate of up to 30 Hz. The acquired pose of a tracked object is streamed to the laptop over a virtual network over a USB connection using the Android Debug Bridge (ADB). The ADB on the client PC is instructed to forward any data on a specific port to the localhost. This allows a network socket in the implemented Unity3D application to connect to the server on the smartphone and to read the streamed tracking data. As soon as an object is detected on any client, its position and rotation are updated for other users as well. The details of this synchronization process are described in the following section.

4.4 Network

The challenge in designing a network architecture for a multiuser VR system is to develop an approach that can distribute a large amount of tracking data while keeping the latency perceived by each user as low as possible. This requirement is critical in situations when users touch each other or pass a virtual (and a tracked real) object to each other. In such cases, a user should not notice latency between the haptic feedback when another user touches them and seeing that user's avatar touching their own virtual body. Similarly, the latency between physically manipulating tracked real objects and seeing their virtual representations change should be minimal. Taking into account these requirements, we carefully select which changes of virtual objects and their properties are synchronized over the network, as well as the type of data distribution.

We use the built-in Unity3D networking as the core of our implementation. Unity3D uses a server-client networking architecture with a star topology: each client communicates only with the server and never directly with other clients. The networking basics in the version 4.x of the game engine used for the first prototype of ImmersiveDeck were significantly different from those of the later engine versions. The most important changes affecting our implementation of networking were the transition to a strictly authoritative server approach and the introduction of limitations on which kinds of GameObjects with their associated network roles can send data over the network. The following sections describe the general concept of our networking approach as well as the details of the initial networking implementation in the version 4.x of Unity3D. The most important changes that were made with the transition to Unity3D 5.x are then described as well.

Generally, our server implementation is non-authoritative. Input of tracking devices is applied to the pose of a player's avatar and their virtual camera directly on the client. The pose of the avatar is then distributed to the server and from the server further to other clients via unreliable state synchronization This way, the delay between the player's input (in this case, movement) and the visual response of the environment is minimized. The same concept of enabling direct visual response is applied to portable objects. When a smartphone attached to a player's HMD is tracking a portable object in the tracking space the tracking input is directly applied to the virtual replica of the object. However, a separate mechanism for the management of object manipulation rights is necessary to enable all the participating players to interact with portable objects. Such a mechanism provides one player at a time the exclusive right to manipulate the object, provided that a special condition enabling object manipulation is fulfilled. Our implementation of the management of manipulation rights is described in the further sections.

Figure 4.8 illustrates tracking input and data exchange between the server and a client application in ImmersiveDeck. The data flow presented in the diagram applies to all the subsequent versions of the platform independently of the details of a particular networking implementation.

4.4.1 Distribution of user movement

In Unity3D version 4.x, networking is managed by NetworkView components that are attached to every GameObject that needs to be distributed. The avatar of each player is network-instantiated by its client application. This way, the server and every client application has a GameObject representing this player's avatar, but changes in this GameObject's state (6DOF pose) can only be made on the client that instantiated it. Head and body tracking runs on the client. The tracking data is first applied to the client avatar and then distributed to the corresponding avatar copies in the server and other clients' applications. The avatar's pose distribution is implemented by synchronizing each avatar bone rotation and the avatar root position and rotation.

The synchronization of each avatar bone is implemented via a Unity3D script that fires a serialization call sending variables as a binary stream at a rate of 60 times per second. State synchronisation of a NetworkView component attached to each bone is set to unreliable. This means that Unity3D sends packets without checking that they have been received, following the UDP protocol. In case of a package loss, pose data coming with the next package is applied to a bone that is being synchronized. In case of an avatar pose synchronization, it is more important to send updates frequently than to receive all packages since the pose is changing constantly and the loss of single packages would not lead to visible errors in the pose. In case of reliable synchronization, the delay caused by packets' receipt verification could lead to disturbing latency.

4.4.2 Object distribution

Our object distribution approach focuses on a direct interaction-feedback loop for each user. Here, it is illustrated on the example of portable tracked objects. However, interaction with door handles, light switches and similar objects that each user can manipulate to change the state of the VE visible to all the other users can be handled following the described concept, independently of whether these objects provide haptic feedback or not.

Our exampled of a portable object is a box that can be held by users. The server and all of the clients have a virtual representation of the box. It is a GameObject that is network-instantiated by the server and therefore controlled by it. Each client has also a local GameObject representing the box. This local copy is created in the client application and is not distributed over the network. When all the users are far away from the virtual box, its pose is set by the server. Users can see the server-controlled copies of the virtual box in their applications but cannot manipulate them. Local copies stay invisible. This situation is illustrated in Figure 4.9 (a) where the blue box represents the client version of the server-instantiated object and the light orange box represents the local object that is currently invisible in the client application. When certain interaction

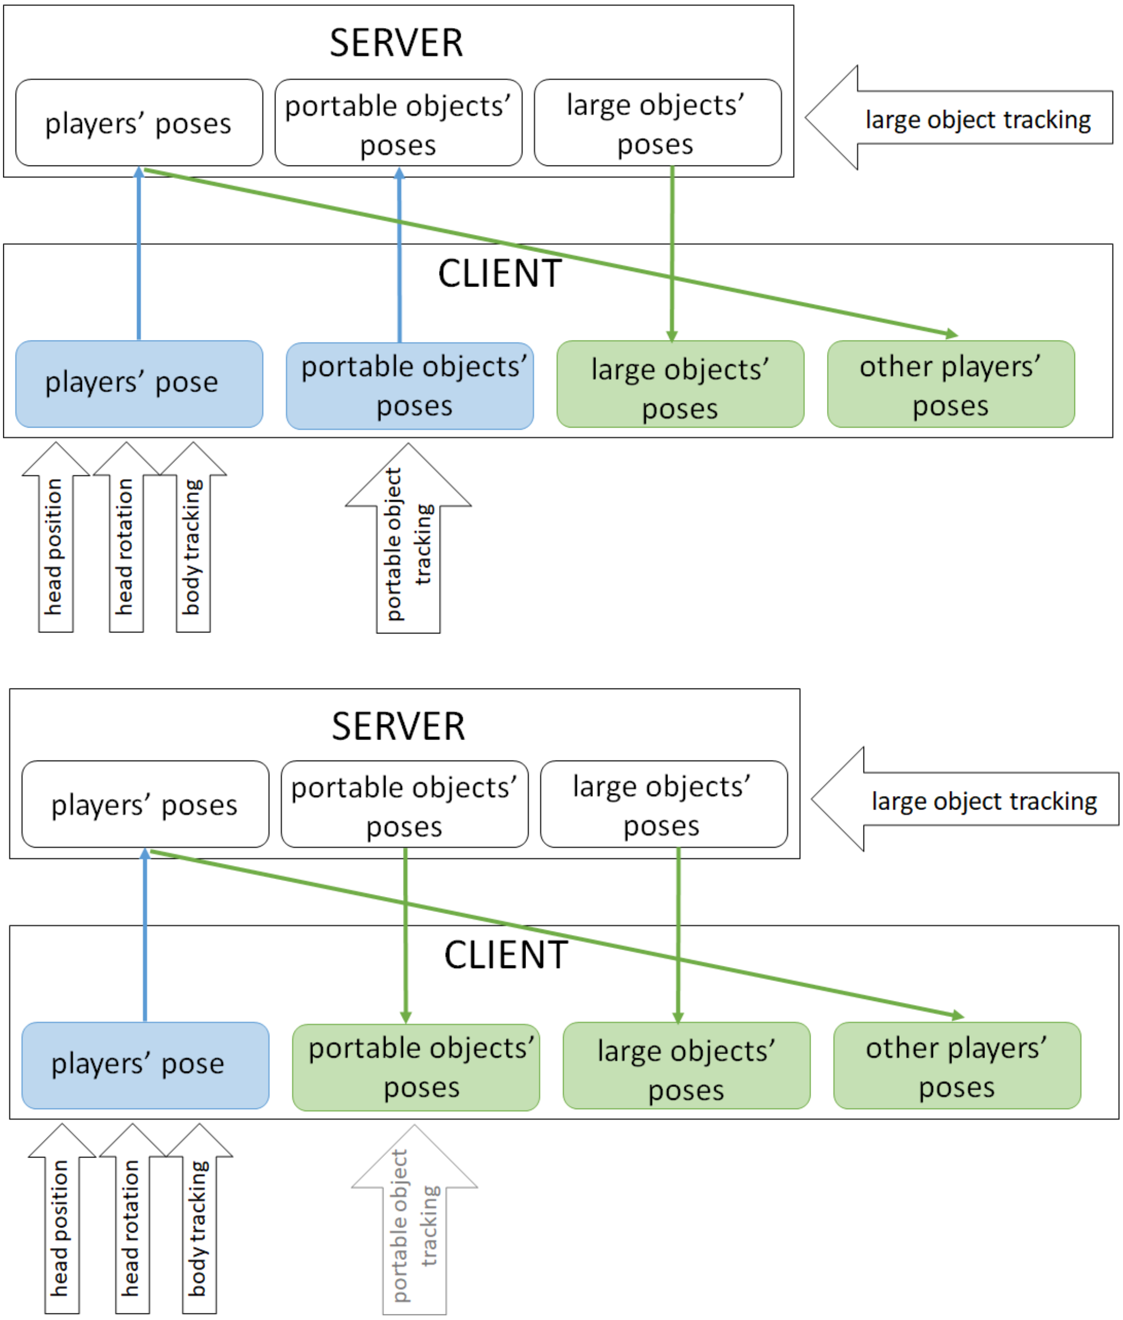

Figure 4.8: Data flow between a client and the server when the tracking of portable objects is active (up) and not active (down) on the client.

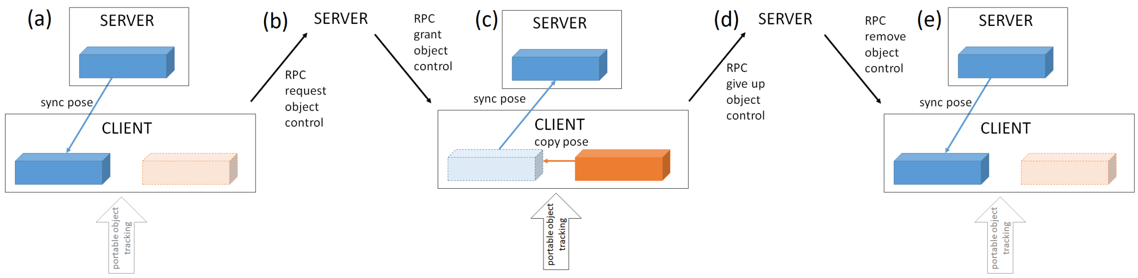

Figure 4.9: Management of object manipulation rights on the example of a tracked portable object.

conditions are fulfilled, i.e. when the real box is tracked by a user's smartphone attached to the HMD, the client application sends an object control request to the server (Figure 4.9 (b)). The server grants control to the client and the local copy of the virtual box gets visible to the interacting user.

An object control request and the request approval are implemented as Remote Procedure Calls (RPCs). Data in an RPC is sent reliably, i.e. it is checked whether every data packet is received using Unity3D's internal acknowledgment procedures. If a packet is dropped, no later packets will be sent until the dropped packet is resent and received.

The user can manipulate the real box and see the changes in the position of the virtual box (the local client copy) corresponding to their actions immediately. The pose of the local virtual box (orange in Figure 4.9 (c))is copied to the currently invisible server-instantiated version of the box (light blue in Figure 4.9 (c)) and streamed to the server. Other users continue seeing the network-instantiated versions of the virtual box mirroring the movement of the server object. The server-controlled object is distributed through serialization calls fired 60 times per second with state synchronization set to unreliable. When the user that is currently manipulating the box sets it aside, another RPC exchange about object manipulation rights takes place as demonstrated in Figure 4.9 (d). The virtual representation of the box switches to the network-instantiated copy again (Figure 4.9 (e)). The rights for the manipulation of the box can now be given to the next client.

4.4.3 User communication

Users of the ImmersiveDeck can talk to each other via headset microphones. While doing this, they hear correctly spatialized sound, i.e. it comes from the direction where a speaking co-player is in the VE and its loudness varies depending on the distance between the speaker and the listener.

The voice communication is accomplished via the TeamSpeak3 plugin for Unity3D. TeamSpeak creates its own server that is running independently of the Unity3D application. The TeamSpeak server gets started with the start of the main server application. With the start of a client application, client TeamSpeak scripts get initialized on a GameObject representing a user's head. Position and rotation values of the user's head are forwarded to

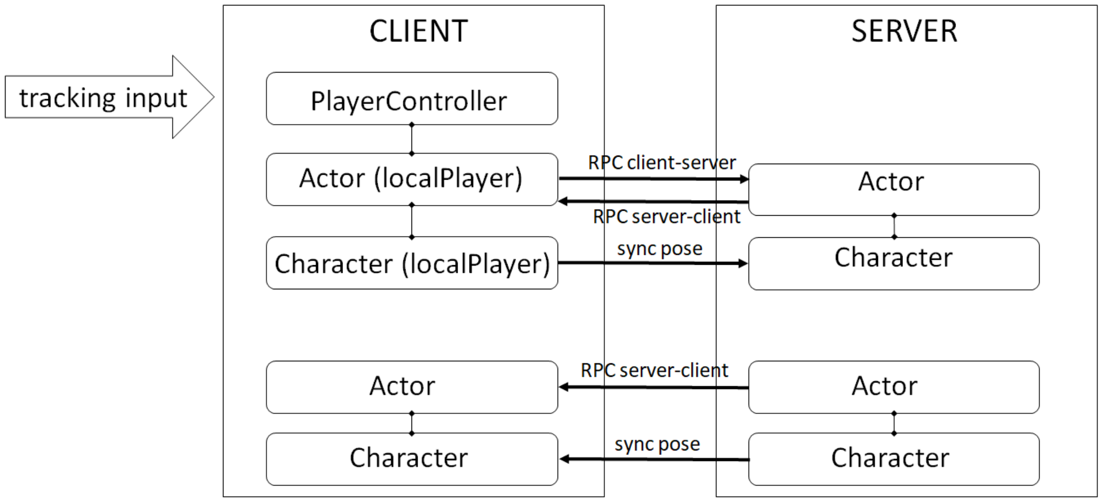

Figure 4.10: Avatar control in the networking implementation in the version 5.x of Unity3D and higher.

the TeamSpeak-program that modifies speech input produced by the users in accordance with their respective positions.

4.4.4 Networking implementation in Unity3D 5.x and higher

The transition from version 4.x to 5.x marked important changes in the general networking approach of Unity3D primarily oriented on the support of an authoritative server approach. In particular, in the versions 5.x and higher only the server application can network-instantiate GameObjects, including players' avatars. Only one GameObject representing the player can be set up in a networked project. When a client application joins the networked game, this GameObject is spawned by the server on all the participating clients and is marked as locally controlled on the newly joined client. The locally-controlled player can receive input from the client application whereas all the other network-instantiated GameObjects are controlled by the server. Since only one GameObject can be spawned as the main player for all instances of the application a special mechanism is necessary to enable the use of different avatars for different players.

In our implementation, the main player is set to an empty GameObject that has an Actor script in its components. The Actor script allows to configure an actual avatar to be used on each client machine. When invoked on the locally controlling client, the Actor script sends an RPC to the server machine with the information of which Character GameObject to spawn as its avatar. This Character GameObject gets attached as a child to the GameObject with the invoking Actor script on all machines. On the local player, the player GameObject gets attached to PlayerController that exists in the client locally and processes tracking input from the the involved hardware devices. The pose calculated by PlayerController is passed through by the Actor script to the Character where it

Figure 4.11: A user standing on the haptic platform (left) to travel along the virtual cliff between two levels of the test scene (right).

is applied to the avatar GameObject and synchronized to the server. This process is illustrated in Figure 4.10. Motion data of not locally controlled Characters is received from the server.

A restriction on client-server communication is introduced in Unity3D 5.x. GameObjects can send information to the server only if they are locally controlled by the client application. Any networked object can however send information from the server to clients. The consequence for our application is that requests of object control can be sent only in the Actor script, or any other scripts of the player GameObject. However, functionality for the assignment of object rights is introduced allowing to transfer the control of network-instantiated GameObjects (like our portable objects) from the server to the localPlayer and back. This way, two copies of the same GameObject (network-instantiated and local) are no longer necessary.

4.5 Non-shared spaces

ImmersiveDeck includes an example scenario that enables transitioning between colocated shared VR and colocated non-shared VR.

This scenario is implemented in an outdoors scene consisting of two levels separated by a cliff. Players can travel between the levels with a virtual elevator that looks like an open platform. We use a real vibrating platform to augment the impression of being in an elevator [148]. The real wooden platform is positioned within the tracking area at a

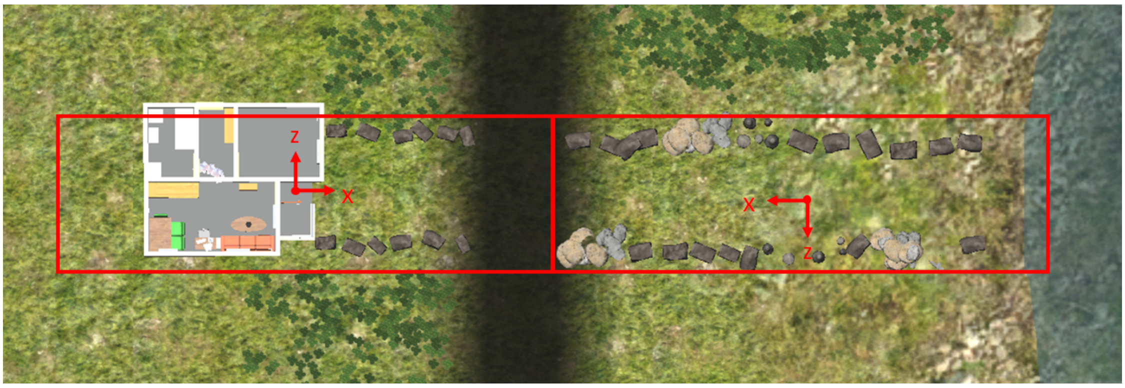

Figure 4.12: Marker map placements within the virtual world positioning a user on the top of the cliff (left map) or below it (right map).

spot coinciding with the position of the virtual elevator. Once a user starts moving on a virtual elevator the platform starts vibrating. The virtual cliff separating two walkable levels and a user traveling on the elevator can be seen in Figure 4.11.

Placement of a user within the virtual world depends on the placement of the marker map within the virtual world. Therefore, our two walking levels correspond to two different placements of the marker map that are illustrated in Figure 4.12 showing the top view of the cliff scene. The transition between two placements is achieved through an offset along the y-axis (up axis) in the Unity world coordinates and a 180° rotation along the y-axis. The translation is applied during a ride that a user takes on the virtual elevator. The rotation is performed at the end of the elevator ride forcing the user to rotate around on the elevator platform to face a new walkable area. This way, a user can traverse the whole tracking area walking from the house to the cliff, step on the elevator platform, rotate around when the platform arrives to the lower lever and traverse the available tracking area again while walking towards the waterfront.

When users are walking on different levels within this outdoors scene they are not supposed to see each other within the virtual world. To prevent possible collisions, we display "ghost" avatars at the real-space positions of those users who are on different level whereas users at the same level see each other as humans. A ghost avatar is simply a semi-transparent box of an average person height. The use of ghost avatars is a preliminary measure for collision prevention; the further investigation of collision prevention measures in colocated non-shared VR is presented in the next chapter of this dissertaion.

4.6 Initial evaluation

This section describes the initial technical evaluation of the ImmersiveDeck platform. This evaluation focused on the performance of user tracking, rendering performance and the analysis of network load.

In the analysis of head tracking, we could not evaluate the absolute accuracy as such an evaluation would require a robust external tracking solution or a measurement system to produce ground truth measurements that we did not have. Instead, tracking jitter was evaluated as it is a characteristic of a tracking system that has direct influence on user experience.

4.6.1 Head tracking

To evaluate static jitter of our tracking solution, we placed a camera on a tripod near the center of our tracking space and recorded its calculated 3D position over 450 consecutive frames. The camera was facing the ceiling, the distance from the camera to the ceiling was 1.5 m. This corresponds to the head position of a user who is 1.70m tall. 5 to 8 markers were tracked in each camera frame. We calculated standard deviation σ of the calculated camera pose $P = (P_x, P_y, P_z)$. Resulting precision is: P_x: $\sigma = 0.4mm$, P_y: $\sigma = 0.9mm$, P_z: $\sigma = 1.2mm$, P: $\sigma = 0.6mm$.

To evaluate dynamic jitter, we used a method similar to the one described by Mossel et al [98]. We rigidly attached two cameras to a metal bar at a fixed distance from each other. During runtime, we calculated the Euclidean distance between the 3D positions of the cameras in every frame, to estimate its standard deviation σ_{dyn} . We placed the bar with the cameras on a tripod at the distance of 1.5m from the ceiling and moved it along the tracking space at normal walking speed. The cameras were facing the ceiling. The resulting $\sigma_{dyn} = 1.5cm$.

4.6.2 Rendering performance and update rate

We measured the overall update rate of the running Unity3D application on the client laptops in two test scenes. The first virtual scene was very simple and contained a 3D model of our tracking room without any textures, 3D models of big (a chair and a table) and small (a box) tracked objects and a light source. The second scene contained a large terrain with plants, several houses and a water surface with enabled reflections. There were two light sources in this scene. This was the scene with the virtual elevator that we used as a test of colocated non-shared VR. The update rate measured on a laptop working on battery power was 23 fps in the first scene and 20 fps in the second scene (averaged among three laptops with identical configurations).

We chose laptops which did not use Nvidia Optimus technology to avoid lower clock speeds when running on battery. However, when running on battery power the GPU memory clock was reduced (not the GPU core clock as with Nvidia Optimus) which still led to a much lower rendering performance. Connecting the laptop to a power source increased the update rate by the factor of 2.5. A fully charged battery provided about 40 minutes of immersion time.

Number of clients	1	2	3
Sent, Mbps	2	1.9	1.9
Received, Mbps	0.55	2.4	4.2

Table 4.1: Network traffic measured on a client laptop (averaged data from three laptops, $\sigma = 0.1$ Mbps).

In Unity3D, a rendering call happens after the update functions of all scripts have terminated, including the update function of the tracking script. Therefore, tracking speed may influence the rendering speed. The speed of Metaio tracking decreased with the increase of the number of markers used in the marker map, since the tracking algorithm iterated through all possible marker IDs when trying to recognize a marker in the camera image. Our performance tests were conducted with the marker map containing 80 markers which we needed to cover the area of 200 m^2.

4.6.3 Network

We measured network traffic on each client laptop when one, two and three users were connected to the server. Average results are shown in Table 4.1. These numbers include head tracking data and body tracking data. Measurements were performed without activated object tracking. Networking update rate was set to 60 times per second in Unity3D. The 5GHz router that we used was theoretically capable to manage a total throughput of 1.73 Gbps. This way, only a very small percentage of the available bandwidth was used by every client.

4.7 Current state

4.7.1 Evolution of the platform

Upgrade to later versions of Unity3D

ImmersiveDeck has been gradually upgraded to newer Unity3D versions, starting with version 5.x and up to version 2018.x.

The most important changes for ImmersiveDeck included networking (described in Section 4.4.4) and the improved support of VR.

Marker tracking improvements

The upgrade of the platform to Unity3D 5.x was performed together with tracking improvements. We switched from the Metaio solution to another algorithm of marker-based tracking. Our new tracking algorithm works on GPUs. The algorithm is an extension of the work of Shibata and Yamamoto on the tracking of fiducial markers with subpixel accuracy [122]. The new tracking solution does not iterate through all present IDs of markers when recognizing a particular marker in a camera frame. This way, the increase in the scale of the tracking area and thus in the amount of markers does not lead to slower tracking. This tracking upgrade allowed us to achieve faster framerates (45 fps at the moment of use of Unity3D 5.x and 60 fps currently).

Integration of new devices

Several input devices have been integrated into ImmersiveDeck.

PrioVR motion capture suits were replaced with Perception Neuron[3]. Using an inertial motion capture suit in combination with a user-worn computer is not ideal since inertial sensors suffer from drift in the presence of sources of electro-magnetic radiation. However, Perception Neuron suits have proven more robust in performance than beta versions of PrioVR that we used in the initial setup, disconnecting much less frequently and requiring a fewer number of re-calibrations during a test run.

Further, new HMDs have been integrated. The current version of the platform can be used with Oculus Rift CV or with HTC Vive. Oculus Rift CV is used with marker-based tracking running on GPU whereas HTC Vive is used with its Lighthouse technology. Furthermore, an inverse kinematics based body tracking utilizing HTC Vive trackers placed on players' limbs can be used for motion capture in combination with HTC Vive and Lighthouse.

LeapMotion[4], a near-range depth tracking device is integrated and can be used for hand tracking.

Hardware improvements

The overall weight of the equipment worn by each user of the initial ImmersiveDeck prototype was 7.4 kg. In the current system, laptops attached to a backpack frame have been replaced with dedicated VR backpack laptops that weight about 3.5 kg. VR backpack laptops have been designed to work on battery power and allow for fast and efficient battery swapping on the run. The slowdown of the GPU memory clock that limited on-battery performance of our first setup has been eliminated.

4.7.2 Usage of ImmersiveDeck

Currently, there are two functioning setups of ImmersiveDeck: at TU Wien in Vienna and at Illusion Walk KG[5] in Berlin. Three different tracking rooms that have been used throughout the development of the setup in Vienna are illustrated in Figure 4.13. The research on ImmersiveDeck was conducted in cooperation with and funded by IllusionWalk.

[3] https://neuronmocap.com/
[4] https://www.leapmotion.com/
[5] https://www.illusion-walk.com/

Figure 4.13: Tracking setup of ImmersiveDeck at TU Wien at different stages of development, starting from the first prototype (left) to the current setup (right).

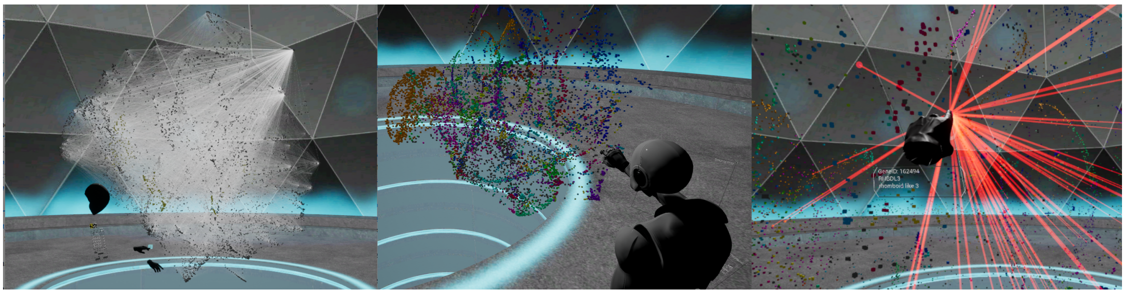

Figure 4.14: Screen-shots from BioNet, the project focusing on teh visualization of large biological networks developed on the basis of the UE4-adapted version of ImmersiveDeck. Left: a users view on the networks consisting of nodes and connecting them links and the avatar of another user. Middle: a view on the network of nodes with links made transparent and an inverse kinematics driven avatar of a co-user. Right: a node and its linked nodes are highlighted during a selection interaction.

Research platform

ImmersiveDeck has been intensively used as a research platform. The following chapters present four experiments on walking within multi-user VR in shared and non-shared scenarios that were carried out using the ImmersiveDeck setup at TU Wien. The setup in Berlin has been used to study different aspects of multi-user VR - social interdependence and cooperation [155]. In addition, single-user experiments focusing on redirected walking have been performed using the platform at TU Wien [149, 150].

The platform has also been used in two applied research projects. One of the projects uses ImmersiveDeck to train first responders. The application scenario includes group leaders training to estimate field situation at the modeled disaster site and subsequently giving commands to individual responders who carry out appropriate actions.

The domain area of the second project is scientific visualization. The project implements

Figure 4.15: Users in the tracking space of ImmersiveDeck in the setup in Berlin. Image credit: IllusionWalk KG.

Figure 4.16: A screenshot from Maratoga Mayday, an immersive four-player experience developed for the ImmersiveDeck setup in Berlin. Image credit: IllusionWalk KG.

interactive visualization of large biological networks that can be interactively explored by teams of researches. This project has been a part of our ongoing effort of porting ImmersiveDeck to UnrealEngine[6]. Figure 4.14 illustrates several views on a network and a co-user in a client application of the project.

Entertainment platform

The setup is Berlin is intended to be used as a location-based VR entertainment center. Currently, IllusionWalk are conducting beta tests of their first immersive experience. Figures 4.15 and 4.16 illustrate the tracking space of the setup and a screenshot from a beta version of the four-player team game developed for the setup.

[6]https://www.unrealengine.com/

CHAPTER 5

Experiments in Colocated Non-Shared VR

This chapter describes two experiments conducted with pairs of participants walking in the same tracking area but in individual virtual worlds. In the first experiment we investigate whether physical proximity within a shared tracking space can affect the experience of users simultaneously exploring their individual virtual spaces. The second experiment explores options of collision prevention that can be employed in various application scenarios that include colocated non-shared VR.

5.1 Mutual proximity awareness

5.1.1 Motivation

The experiment presented in this section addresses the following situation. Several users are walking around within the same shared tracking space immersed into different VEs. As they are in separate virtual worlds they do not see each other. We also assume that users do not hear each other as their VR experiences may include music or multiple environmental sounds.

Our main question is: can users notice each other being in close proximity in the physical space even if they do not see or hear each other in the virtual world? And further, if users do notice each other, what are interpersonal distances that allow this identification? What are the contributing factors?

Our research question is not simply a matter of interest but has important implications for experience of users in colocated non-shared VR. If users can really perceive the presence of someone in their proximity outside of the virtual world in which they are immersed it could lead to breaks of presence. In this case safety mechanisms that should be developed for colocated non-shared VR for collision prevention must also consider some minimal allowed interpersonal distances.

We hypothesize that three factors could allow users to perceive the presence of each other: body heat, airflow when walking users pass each other and floor vibration caused by steps. We designed an experiment to investigate these possible cues of awareness.

5.1.2 Experiment design

The above mentioned factors may have greater or lesser impact depending on relative positions of users in respect to each other, their heading directions and walking speeds. For example, if two users are heading towards each other with relatively high speeds, the airflow might be noticeable enough to reveal the presence of another person. On the contrary, when users stay close to each other during a longer time, body heat and floor vibration might have larger impact.

Interpersonal distances and walking speeds that might trigger awareness are of special interest. We conducted several pilot study sessions with experienced VR users to test the effects of various speeds and distances. During these tests, pairs of participants were asked to walk on parallel trajectories through the whole extend of our tracking space, either side by side or heading towards each other. The distance between the trajectories varied from 3m to 1m. To investigate the variation of walking speeds we asked participants to walk at comfortable speeds or as fast as they could without running. This allowed for the variation of walking speeds from 0.7 m/s to 1.5 m/s. None of the participants could notice the test partner in any of the conditions.

The pilot test resulted in several observations that were taken into account for the design of the subsequent main experiment. First of wall, walking speeds of the participants from a pair were very different. This lead to situations where one of the participants quickly walked away from their test partner in walking side by side conditions making it certainly impossible for both participants to notice each other. Secondly, the task of walking at the highest possible speed proved too involving for the participants so that they could not concentrate enough on the cues coming from the real world. Having analyzed the results of the pilot test sessions, we decided to let participants walk with low and controlled speeds on paths that were very close to each other in the main experiment.

Participants of the pilot test were informed about the purpose of the experiment. However, we initially intended to conduct the main experiment with naive users in order to avoid possible bias caused by a-priori knowledge of the research question. This made our task formulation difficult as we wanted to record precise moments when participants notice each other and not just use post-test measures. Finally, we decided to split participants into two groups, one group being naive to the research question and the other one knowing it a-priori. As a result, our experiment had a mixed design with one between-group factor, the knowledge of the research question, and one within-group factor. The within-group factor is the type of the trajectory, with four levels.

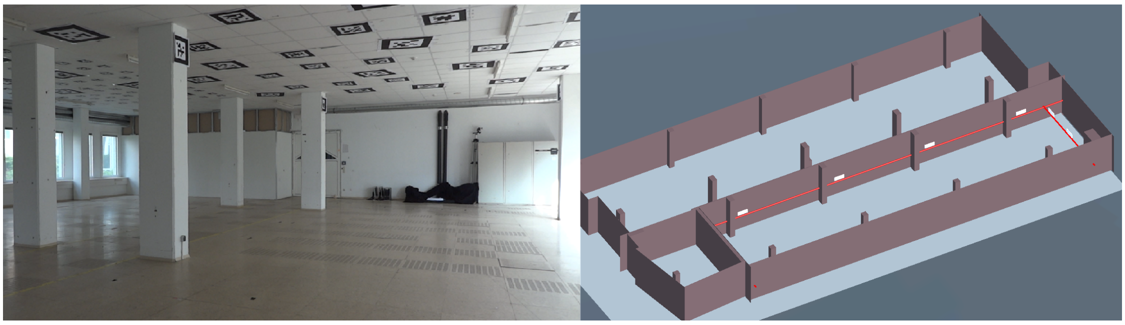

Figure 5.1: Tracking laboratory and the top view of its virtual replica used in the experiment.

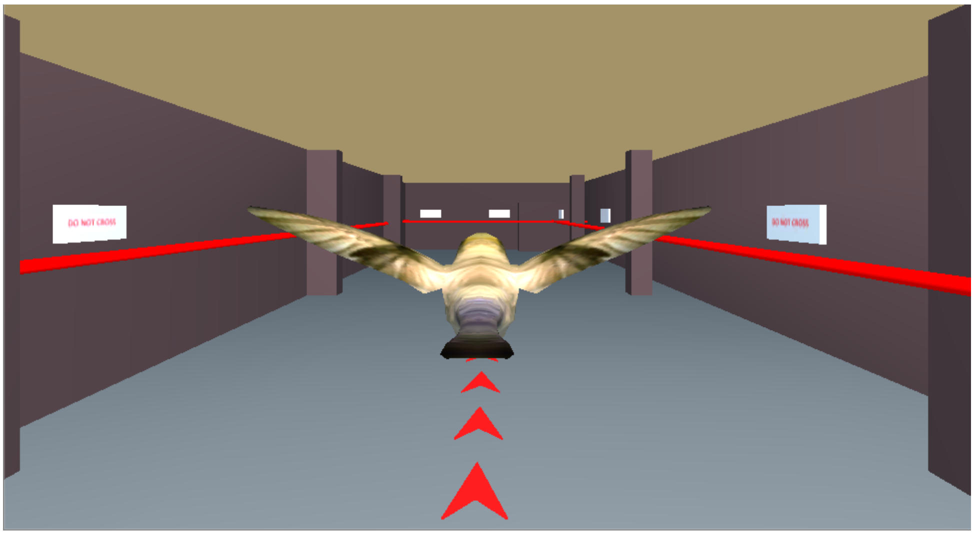

Figure 5.2: First-person view of a test sequence.

Environment and tasks

The experiment consisted of four tasks where users were asked to walk on pre-defined trajectories in a virtual room that resembled the tracking room where the experiment took place. The virtual room can be seen in Figure 5.1. Each participant from a pair could only see their own path and did not know how their test partner would walk.

In order to make users walk at about similar speeds we introduced the requirement of following a bird flying in front of a participant. Participants were instructed to always stay behind the bird and the only start walking when it starts flying. Similarly, they had to stop if the bird stopped. Figure 5.2 illustrates the view of a participant walking behind the bird.

In the first task, participants walked along the straight trajectories towards each other. Two walk sequences were performed in this task, the first at the speed of 0.3 m/s and

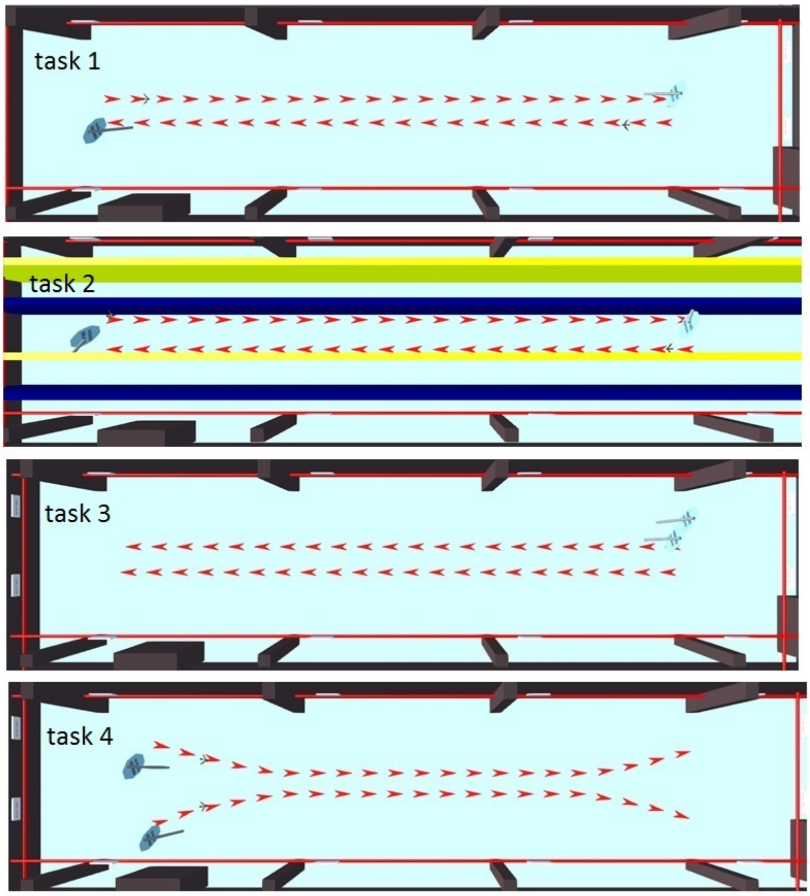

Figure 5.3: Configuration of pre-defined trajectories in the experiment tasks.

the second at the speed of 0.7 m/s. The distance between the trajectories was 0.8 m. In each walk, the participants from a pair passed each other in the middle of the room. The arm-to-arm distance between the participants at the moment of passing was about 30 cm. This task was predominantly designed to evaluate the contribution of the airflow as an awareness cue.

The configuration of the trajectories was the same in the second task. This time, however, the path of each participant was surrounded by walls producing an impression of walking in a corridor. In this task, we wanted to investigate whether the nature of a VE itself possibly contributes to the awareness of events happening outside the VE. Our assumption was that, since vision is the dominant cue for most people, users might be less aware of others when they are exploring geometrically isolated VEs like narrow streets or corridors in one walk at 0.3 m/s.

In the third task, participants were asked to walk along straight trajectories with any speed they felt comfortable without a bird to follow, the distance between the trajectories being 1m. Participants walked in the same direction. This experimental sequence was introduced as a more ecological task as walking speeds would not be restricted in an application. We also intended to collect the data about the typical walking speeds in immersive VR.

The fourth task consisted of two walk sequences side by side along the curved trajectories, again following the bird flying at 0.3 m/s and 0.7 m/s respectively. The trajectories were 3 m apart at their starting points and became progressively closer to each other towards the middle of the room reaching the shortest distance of 0.8 m. This task was designed to determine the distance at which participants would start perceiving the presence of each other, if at all. Additionally, we wanted to see if participants were more likely to notice each other in Task 4 due to the body heat or in Task 1 due to the airflow. We did not have a guess about whether floor vibration would have more contribution in the side by side walking tasks or in the passing each other tasks. The configurations of the trajectories used in all four tasks are illustrated in Figure 5.3.

We call the group of participants who knew the research question and the goal of the experiment *Informed* group. Participants from the *Informed* group were instructed to point at their test partners if they noticed them. We call the group of participants who were naive to the purpose of the experiment *Non-Informed* group. Participants from the *Non-Informed* group were asked to point if they thought there was something close to them in the real environment. We expected that participants from the *Informed* group would be able to notice their test partners more often than participants from the *Non-Informed* group.

Measures

Pointing The events of participants pointing at their test partners were recorded. The positions of both participants from a pair when one of them pointed were recorded and the distance between them calculated.

The number of times pointed (per participants and per each trajectory) and the distance between two participants from a pair are our objective metrics.

Questionnaires The pre-test questionnaire contained Simulator Sickness Questionnaire (SSQ) [72] and demographics (age, previous experience of VR and playing computer games).

Questions from the post-test questionnaire are summarized in Table 5.1. Although our main interest was in participants' awareness of each other we also included questions regarding immersion and presence in the VE using the questionnaires suggested by Witmer and Singer [158] and Slater, Usoh and Steed [130] (questions *Pr1* - *Pr5* in Table 5.1). The inclusion of presence-related questions is motivated by the fact that users may be more aware of the events occurring in the real world if the VE fails to elicit the sense of presence. Questions *Aw1*, *Aw2* and *Aw3* from Table 5.1 were used participants' mutual awareness. Questions *Aw1* and *Aw2* were asked for each of four tasks.

Table 5.1: Questions used to assess presence and awareness of the test partner during the experiment.

QUESTION CODE	QUESTION
Pr1	How aware of real events occurring in the room around you were you during the game?
Pr2	How involved were you in the experience?
Pr3	Rate your sense of being there in the virtual room.
Pr4	When you think about the virtual playground, was it rather like something that you saw or somewhere that you visited?
Pr5	There were times during the experiment where the virtual environment became reality for me.
Aw1 (T1-T4)	How aware of the presence of your test partner in the real room were you during the task?
Aw2 (T1-T4)	How safe from collisions with your test partner did you feel during the task?
Aw3	How often during the whole experiment did you notice your test partner?

Most of the questions, excluding the SSQ, were presented on the Likert scale with the range from 1 to 7, where 1 corresponds to "not at all" and 7 to "very much" in the questions asking to rate the extent of a particular phenomenon. For question *Pr5*, 1 means "strongly disagree" and 7 means "strongly agree". The answer for question *Aw3* was the number of times a participants noticed their test partner.

Participants

36 participants took part in the experiment, with nine pairs in the *Informed* and *Non-Informed* groups each. Six participants in each group were female. The age of the participants ranged between 20 and 53 years, with the median of 27 years. We did not explicitly aim to conduct the experiment with naive users only. However, only four participants had previous knowledge of VR technology, three of them having advanced knowledge in VR (in the *Informed* group) and one being an experienced user (in the *Non-Informed* group).

Apparatus

The experimental application was developed within our ImmersiveDeck platform described in Chapter 4.

Motion capture suits were not used in the experiment; participants had no self-avatars. The walkable area of the experiment was approximately 30x7 m large. The overall rendering framerate was 26 fps.

Participants wore headphones (Asus Vulcan Pro) with enabled noise cancellation to make sure that all the sounds from the real environments were blocked. Additionally, background white noise with the average loudness of 65dB was played in the headphones.

The experiment coordinator saw the positions of both participants from a pair within the VE on the server machine. Test sequences were triggered manually from the server. The events of participants pointing at their test partners were also manually registered by the experiment coordinator on the server machine. The experiment coordinator was able to trigger notifications for each of two participants in the form of signs with short instructions "Stop", "Continue walking", "Task is over", "Get ready for the next task".

Figure 5.4: Pairs of participants walking during the experiment.

The "stop" sign was used as emergency if the participants from a pair did not exactly stay on their pre-defined trajectories and were at risk of a collision.

Procedure

Participants arrived to the tracking lab in pairs. Some of them knew each other but some had never met before.

The experiment coordinator greeted the participants and explained the experimental procedure. Each participant gave an informed written consent of participation in the experiment. Then, the pre-test questionnaire comprising the SSQ and demographics was filled out. The experiment coordinator explained the task. At this moment, the general goal of the experiment was explained for participants from the *Informed* group. Additionally, participants were given hand-outs with the written explanation of the tasks.

The experiment started from Task 1. There were no immersion breaks between the tasks. After Task 2, one of the participants from a pair was instructed to cross the tracking area to reach the starting position for Task 3, while staying immersed and not seeing the other participant. The experiment completion times ranged between eight and twelve minutes. Pairs of participants walking at minimal distances between each other can be seen in Figure 5.4.

After the experiment, participants were asked to fill out the post-test questionnaire containing the repetition of the SSQ and questions concerning presence and mutual awareness. After the post-questionnaire had been filled out, a short debriefing session took place. The participants were asked if they could feel the presence of the other user next to them; if yes, during which phases of the experiment it happened and what were the main sources of their awareness. The participants were also asked if they could identify their position and their partner position in the real space during the experiment. The goal of the experiment was explained to participants from the *Non-Informed* group. Participants received sweets as a note of gratitude.

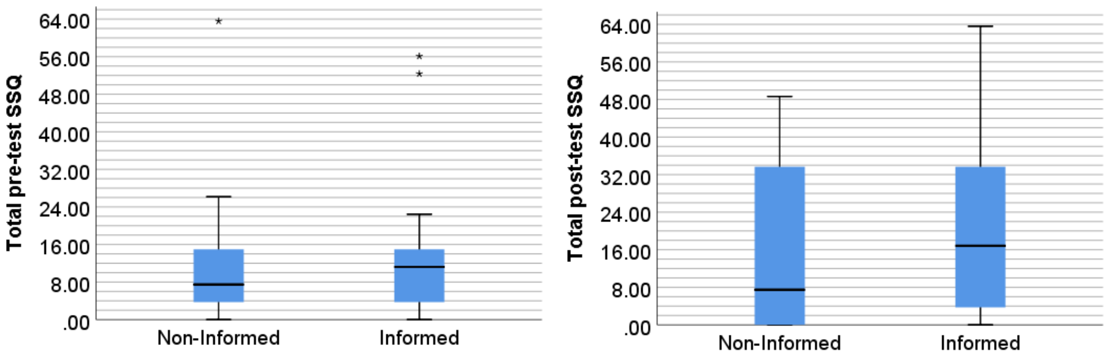

Figure 5.5: Boxplots of the pre-test and post-test total SSQ score in both groups.

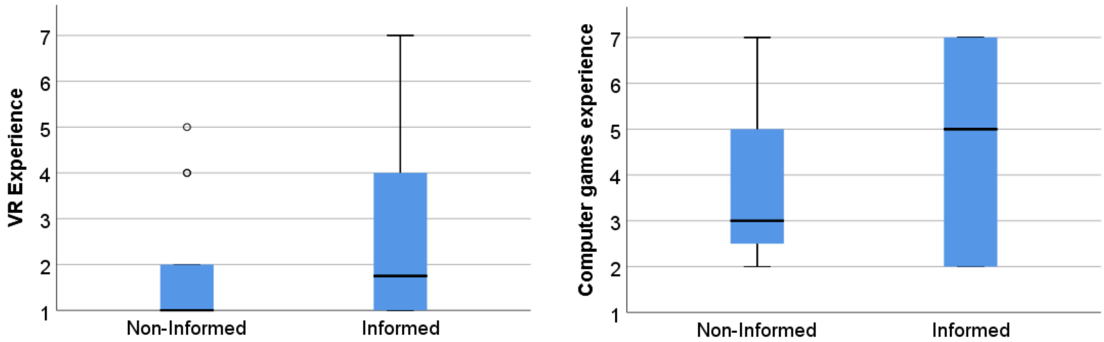

Figure 5.6: Self-report of the previous experience of VR and playing computer games.

5.1.3 Results

Data analysis

The resulting questionnaire scores were not normally distributed for any of the questions. Therefore, non-parametric tests were used for data analysis: Mann-Whitney U test to compare two independent samples, Wilcoxon signed-rank test to compare two related samples and Friedman's ANOVA to compare several related samples. Differences are reported as significant at $p < 0.5$ level.

VR experience and simulator sickness

Figure 5.5 demonstrates boxplots of the total SSQ score [72] before and after the experiment in both test groups. The difference between the pre-test and post-test total SSQ score is not significant for both groups ($T = 59.0$, $p = .682$ for the *Non-Informed* group and $T = 333.5$, $p = .126$ for the *Informed* group in the Wilcoxon signed-rank test). There is no significant difference in the post-test SSQ score between the *Non-Informed*

Table 5.2: Numbers of valid individual walks for each task and among them the number of walks where a participant pointed at their test partner.

NON-INFORMED	TASK 1	TASK 2	TASK 3	TASK 4
TOTAL WALKS	36	16	16	32
POINTED AT PARTNER	2	0	0	0
INFORMED	TASK 1	TASK 2	TASK 3	TASK 4
TOTAL WALKS	34	16	14	22
POINTED AT PARTNER	6	2	3	7

and *Informed* groups either ($U = 197.5$, $p = .265$ in Mann-Whitney U test). None of the participants reported strong simulator sickness symptoms.

Previous experience of VR and playing computer games was similar between the test groups ($U = 193.0$, $p = .339$ for VR experience, $U = 166.5$, $p = .660$ for experience of playing computer games in Mann-Whitney U test). Boxplots can be seen in Figure 5.6.

Mutual awareness

We initially intended to use the number of times when a participant pointed at their test partner as an objective metric of mutual awareness. However, in cases when pointing happened participants pointed at their test partners continuously during parts of a walking sequence. As a result, we used the number of walking sequences in which pointing happened rather than the number of discrete pointing events. We counted walks individually as often only one of the participants from a pair pointed at their test partner during a walk. This way, one sequence where two participants crossed the tracking space walking along their pre-defined paths resulted was counted as two individual walks.

18 pairs of participants performed the total of 18 x 2 x 6 = 216 individual walks. In 30 of these individual walks, the participants touched each other shoulder-to-shoulder while walking, so we discarded all the events happened in these walks. However, 1/3 of the events where participants touched each other were not noticed or remembered by the participants.

The total amount of cases where at least one of the users from a pair pointed at the other one was very low. The numbers for each task for both *Non-Informed* and *Informed* groups are given in Table 5.2. In only two out of 104 valid individual walks in the *Non-Informed* group the participants noticed their test partners. These were two different participants. In the *Informed* group, participants noticed their test partners in the total of 18 out of 86 individual walks. However, these results are produced by only six out of 18 participants in the informed group who pointed at their test partners. From six participants who noticed their test partners, only one had advanced VR knowledge. The other two participants with expert VR knowledge from the informed group never pointed at their test partners. The only participant who had extensive VR experience in the non-informed group never pointed at his test partner. In the *Informed* group, there were

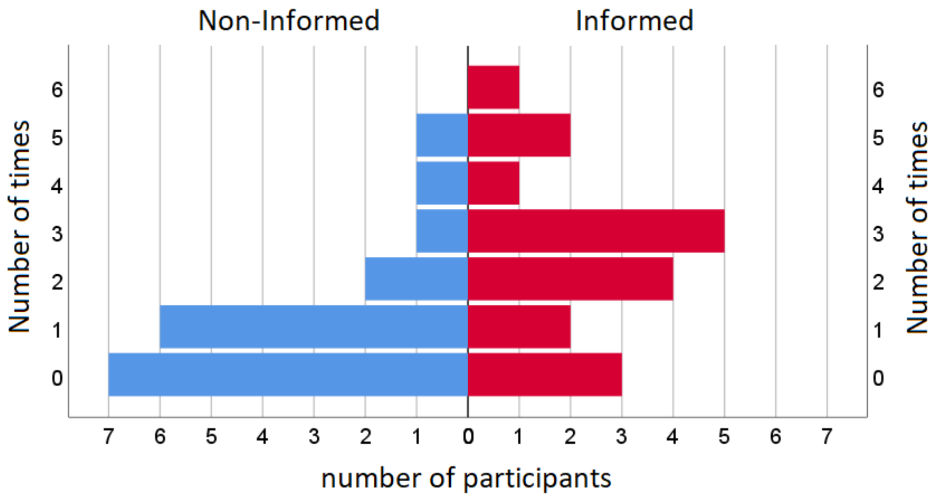

Figure 5.7: Reported number of times when a participant noticed the presence of their test partner in the same tracking space.

several occurrences of false positives, i.e. participants pointing at an empty space while they thought their test partners were there. False positives are not included in Table 5.2.

The participants who noticed their test partners pointed at them right after passing them in Task 1 and Task 2 where participants were walking towards each other. In the walking side by side tasks, the participants pointed at their test partners when they reached the head-to-head distance of 0.77 - 1.2m.

Only one of six participants who correctly pointed at their partner did it in Task 3 and Task 4 (with walking side-by-side) tasks and not in Task 1 and Task 2 (with walking towards each other). As the total number of pointing events was low there was no definite evidence of whether walking side-by-side or towards each other leads to higher awareness.

Figure 5.7 shows the distribution of answers on question Aw3 in which participants were asked to estimate how often they noticed their test partners in the course of the entire experiment. The question reflects rather the memory of a participant than the actual number of times when they may have noticed their test partner. As mentioned above, in some cases where participants touched each other shoulder-to-shoulder they did not notice or remember the event, although in others they did. Noticing the test partner reflected in the answer on question Aw3 did not necessarily lead to pointing at them. In many cases, participants reported to have noticed them in a moment but they were unable to identify where exactly the test partner was.

Box-plots of the resulting scores for questions *Aw1* and *Aw2* for each task can be seen in Figure 5.8. The median scores reflecting subjective awareness of the test partner (question *Aw1*) are higher in the *Informed* group than in the *Non-Informed* group for all

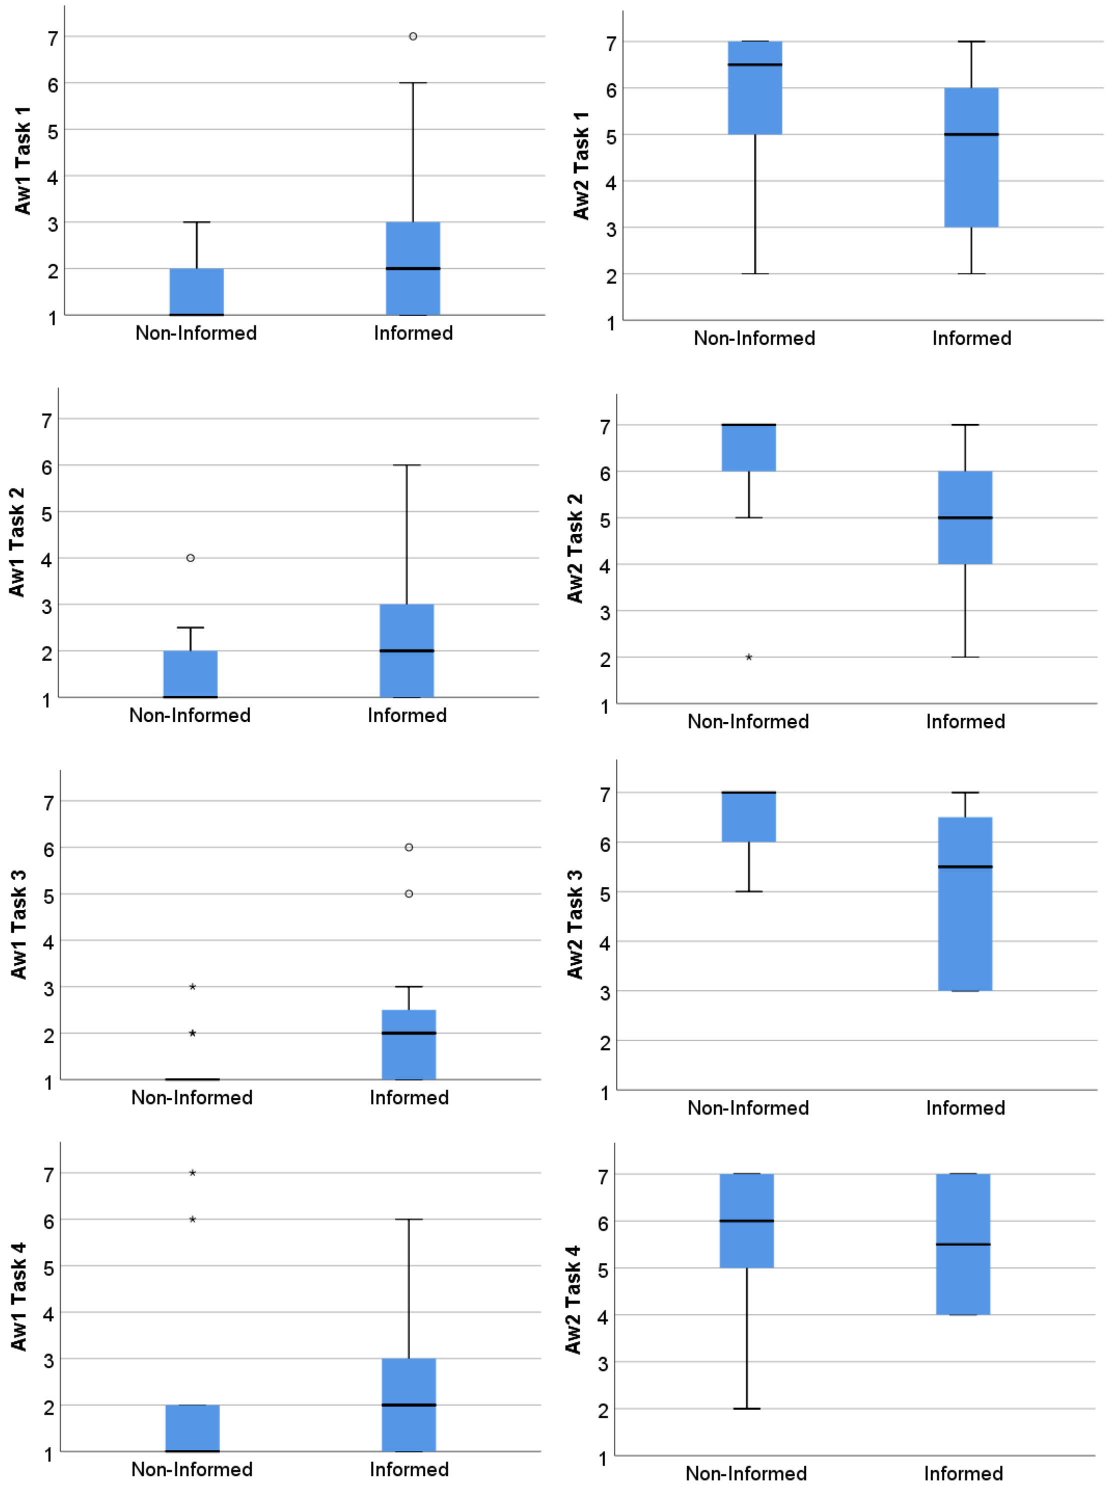

Figure 5.8: Boxplots of the resulting scores of the questions *Aw1* and *Aw2*.

Table 5.3: Results of the Mann-Whitney U test for the question *Aw1*. Test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

	U	z	p	r
TASK 1	191.5	1.772	.102	.30
TASK 2	185.5	2.358	**.029**	.32
TASK 3	157.0	1.732	.138	.31
TASK 4	49.5	1.067	.368	.24

Table 5.4: Results of the Mann-Whitney U test for the question *Aw2*. Test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

	U	z	p	r
TASK 1	94.5	-2.193	**.031**	-.37
TASK 2	70.0	-2.837	**.005**	-.48
TASK 3	89.0	-2.189	**.035**	-.37
TASK 4	77.0	-2.577	**.011**	-.44

tasks. Median scores reflecting the perceived safety from collisions with the test partner (question *Aw2*) are in contrast higher in the *Non-Informed* group. However, only some of the between-group differences in median scores are statistically significant. Tables 5.3 and 5.4 show the results of the Mann-Whitney U test for questions *Aw1* and *Aw2*. Tests with p-value shown in bold are statistically significant.

We also compared subjective awareness scores of participants within the *Informed* group where the majority of pointing event happened. We compared the scores for questions *Aw1* of those six participants who repeatedly pointed at their test partners and those who did not. The boxplots reflecting the scores can be seen in Figure 5.9. Median subjective awareness scores were indeed higher among those participants who pointed at their test partners. The difference is however only statistically significant for Task 1 (details of the statistical comparisons are presented in Table 5.5). We believe that the absence of statistical significance in Tasks 2 and 4 is due to the small sample size, especially of those participants who noticed their test partner. Task 3 can be seen as a control condition, as in the most cases participants walked at different speeds and only had a chance to notice each other in the beginning of the walking sequence. The comparison of awareness scores (question *Aw1*) between the tasks in the *Non-Informed* group did not reveal any statistically significant difference ($\chi^2(4) = 3.522$, $p = .318$ in Friedman's ANOVA).

In the de-briefing sessions after the experiment, "steps" or floor vibration caused by walking were pointed out as the main factor in noticing the test partner. Two of the participants mentioned that they had felt the increased air flow, both times in Tasks 1 and 2 where participants where walking towards each other. Body heat was never mentioned. Many of the participants who noticed their test partners when they slightly touched while walking said that they were aware of their test partner being somewhere

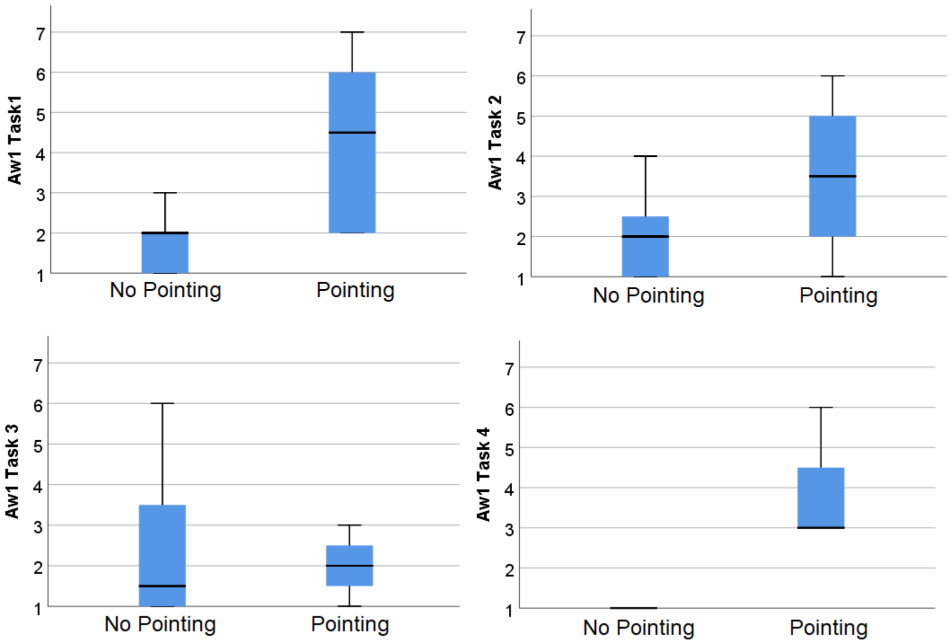

Figure 5.9: Boxplots of the resulting scores for question *Aw1* given by the participants from the *Informed* group, split by the outcome of the experiment (whether they pointed at their test partner or not).

Table 5.5: Results of the Mann-Whitney U test for the question *Aw1* answered by the participants from the *Informed* group who pointed at their test partners and those who did not. Test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

	U	z	p	r
TASK 1	59.0	2.744	**.007**	.67
TASK 2	52.5	1.606	.125	.38
TASK 3	18.0	.359	.808	.10
TASK 4	9.0	2.121	.10	.61

around them until the end of that particular walking sequence but could not any longer say where their test partner was when the next sequence started. Many participants felt disoriented and could not say where in the tracking space they had walked during the experiment, even though the test scene was a virtual replica of the real room except for one additional wall. However, those six participants from the *Informed* group who pointed at their test partners gained a clear understanding of spatial arrangements during the experiment. They reported to have quickly realized that they were walking either towards or side-by-side their test partners.

Table 5.6: Results of the Mann-Whitney U test for presence-related questions. Test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

	U	z	p	r
Pr1	155.0	-.231	.839	-.04
Pr2	132.0	-.714	.503	-.12
Pr3	141.0	-.708	.521	-.12
Pr4	137.5	-.787	.443	-.13
Pr5	158.5	-.113	.913	-.02

Table 5.7: Results of the Mann-Whitney U test for presence-related questions in the *Informed* group. Median values for those participants who did not notice their test partner during the test (*No Pointing*) and those who did (*Pointing*), test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

	Median *No Pointing*	Median *Pointing*	U	z	p	r
Pr1	2	2.5	43.5	.730	.494	.17
Pr2	4.5	6	46.5	1.024	.335	.24
Pr3	5	5	44.5	.877	.437	.21
Pr4	4.5	5	30.5	-.522	.616	-.12
Pr5	4	5	38.0	.190	.892	.04

Presence

Immersion and presence-related questions were used in the post-test questionnaire to track possible differences that could lead to different levels of mutual awareness among the participants.

The resulting scores show that participants from both *Non-Informed* and *Informed* groups judged their experience as very immersive. Scores of any of the questions are not different between the groups. The details of the corresponding Mann-Whitney U tests used for the analysis can be found in Table 5.6. Boxplots of the resulting scores of questions *Pr1* - *Pr5* are shown in Figure 5.10.

As with the subjective mutual awareness scores, we compared the scores for presence-related questions of the participants who pointed at their test partners during the test and those who did not as their higher awareness of the test partner could be explained by lower immersion. However, the scores of questions *Pr1* - *Pr5* were not different between those participants who pointed at their test partners and those who did not (the details of the statistical comparisons are summarized in Table 5.7).

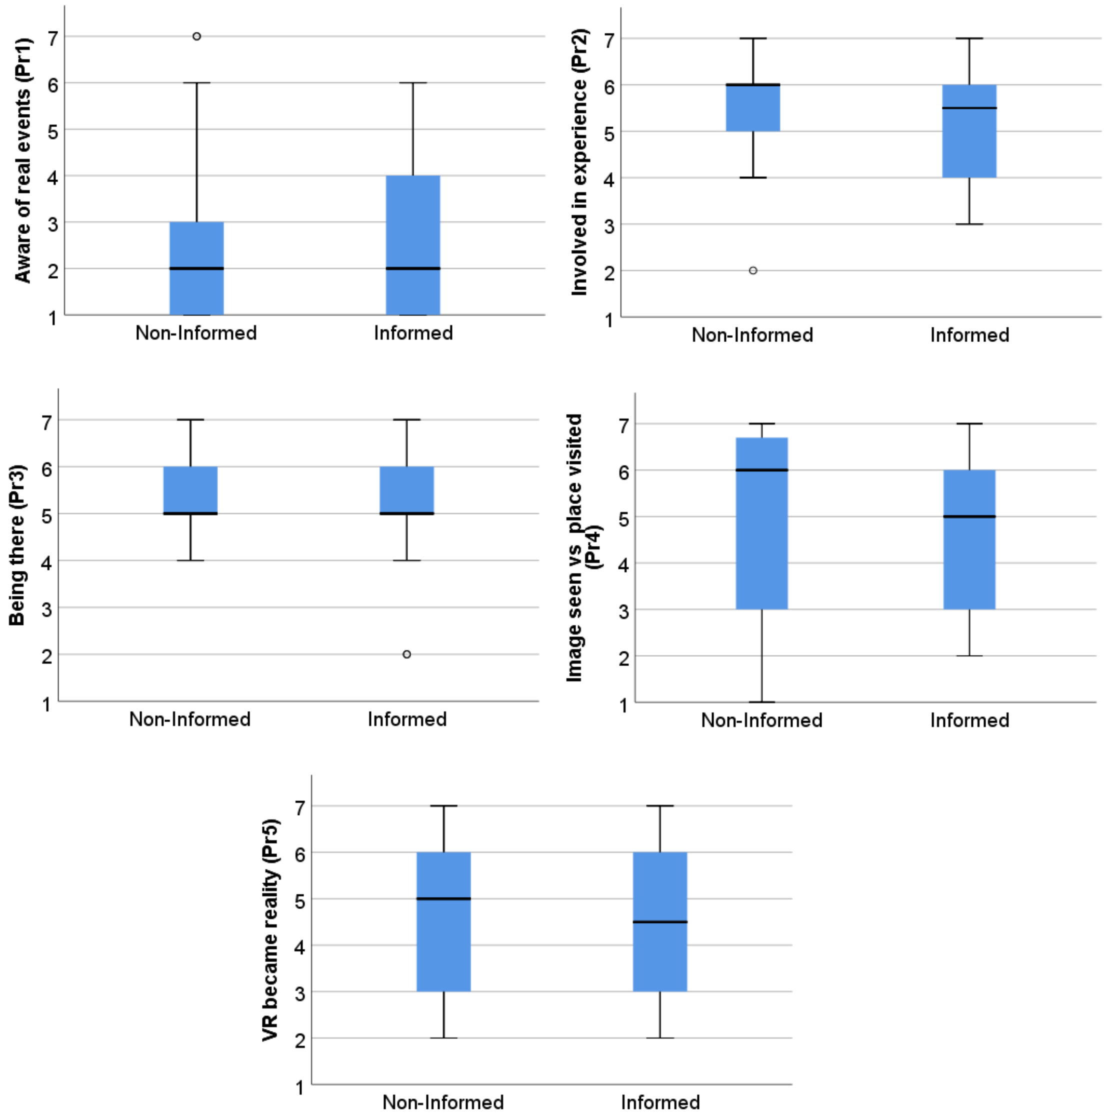

Figure 5.10: Boxplots of the resulting scores of the self-report of presence.

5.1.4 Discussion

The main question of our experiment could be answered with high degree of certainty: in the majority of cases, colocated users do not notice each other in the shared tracking space if they do not see or hear each other in VR, even at very small interpersonal distances. The ability to notice the test partner seems to be dependent on personal attention. Those six participants from the *Informed* group who repeatedly pointed at their test partners did not report lower immersion as the rest of the participants, so it can be excluded as a possible explanation of their increased awareness. The same conclusion can be made about familiarity with VR technology as advanced VR users did not perform better at recognizing the presence of others than naive users. The only factor that distinguished those especially perceptive participants from other participants was that they maintained a good understanding of their movement within the real space and could imagine how their test partners were moving during the experiment. The knowledge of the experimental task however played the decisive role in the outcome as participants from the *Non-Informed* group did not identify their test partners, except for in two walking sequences out of 104. In fact, a number of the participants from the *Non-Informed* group reported to have thought that they would be walking within a separate physical space each.

The configuration of the trajectories did not produce any effect. The effect of steps was indicated as the primary source of awareness. It is possible that body heat was masked by heat generated by the laptop worn on a user's back. The impact of airflow was only reported by two users. However, we chose low walking speeds to assure user comfort. It is possible that provided with a more light-weight equipment and an application that naturally requires fast motion users would be able to perceive the airflow caused by motion of their nearby co-players.

It is a curious result that participants from the *Informed* group were slightly more concerned about their safety in terms of colliding with each other, as demonstrated by significant differences in the scores for question *Aw2* for each task, although absolute perceived safety scores were high for both groups. Participants from the *Non-Informed* group were not given any information about their test partners during the experiment, except for that they would not see them. Possibly, the assumption of separate walking spaces that some of the participants maintained contributed to their perceived safety from collisions. High involvement in the experimental task might have furthermore made them forget about their test partners. The attention of the participants from the *Informed* group was, in contrast, brought to their test partners by the very formulation of the experimental task.

The result of very low mutual awareness is positive in terms of immersion: outside factors like proximity of others are unlikely to introduce breaks of presence. It is reasonable to assume that visually more impressive scenes and scenarios where users interact with the VE would lead to even higher involvement and ever decreased attention paid on real surroundings. As a conclusion, collision prevention techniques appear to be especially important for immersive VR setups with non-shared VEs to guarantee the safety of all users.

5.2 Towards collision prevention

5.2.1 Motivation

As the results of our experiment on mutual awareness in colocated non-shared VR showed, users walking in individual VEs within a shared tracking space do not notice each other in the majority of cases, even if they are side-by-side to each other. Collision prevention methods need to be provided to ensure safe exploration of virtual worlds in colocated non-shared VR setups.

An ideal collision prevention method would be seamless. It would predict all future collisions of walking users and apply an actual prevention strategy once a collision is predicted. For example, RW techniques could steer users away from each other either by adding gain to their movement or by modifying the VE on the fly and thus making them to choose a virtual path that leads away from the future collision, all of it without a user noticing the manipulation. However, even though first attempts at multi-user RW exist such an ideal collision prevention method has not been presented yet. Moreover, such a collision prevention algorithm is unlikely to be generic and suitable to all types of VEs and any number of users. The task of the application also needs to be taken into account as it would influence the distribution of interest points in the VE that would usually serve as attractors in the part of collision prevention that deals with the prediction of future user paths. A number of users is also likely to influence the performance of collision prevention as it imposes constraints on available safe redirection trajectories. In short, the creation of a robust and broadly applicable seamless collision prevention method is a very difficult task. However, the absence of such algorithm should not be an impediment for the development and use of colocated non-shared VR experiences.

In this section, we investigate which measures could be taken if a seamless collision prevention algorithm is not available or in cases when it fails. Especially in cases when a seamless collision prevention method fails, time constraints for preventing a collision become stricter. As a result, a robust method for averting an approaching collision becomes necessary.

The state of the art literature suggests to stop users in dangerous situations where a collision with either a wall in the real environment or another user is imminent [69], accompanied with the suggestion of blanking the HMD of a user in a collision danger. In fact, the number of times when users need to be stopped and the VE reset is used as a measure of the success for RW techniques. However, there are several unsolved questions regarding this approach of stopping users and blanking their HMDs. First of all, it is mostly a theoretical suggestion. No research was conducted on stopping users while they are immersed into a virtual world. It is reasonable to expect that a user would not stop at once while exploring a VE, therefore a safety margin needs to be introduced. Secondly,

Figure 5.11: Visual signals used in the experiment for stopping walking users.

it has not been studied which signals are better to use to stop users. Blanking the HMD hardly seems to be an ideal way to stop a user as it would make them disoriented and would certainly introduce breaks in presence.

In the experiment presented in this section, we first investigated four different signals used to stop a walking user without blanking their HMD and, secondly, explored an alternative strategy of warning a walking user about a possible collision based on our first experiments with colocated non-shared VR during the development of ImmersiveDeck.

5.2.2 Experiment design

Part 1: Stopping users

The goal of this experiment was to compare the effectiveness and plausibility of different types of signals that could be used to stop a user walking in a VR environment. This part of the experiment had a 4 x 1 within-subject design. The stopping signal was the within-group factor with four levels.

Stopping Signals Two visual and two auditory signals were tested. The visual signals were a human figure (*VirtualHuman*) and a stop sign (*StopSign*) appearing in front of a walking user. The auditory signals were a short "stop" command (*StopSound*) and a short clap sound (*ClapSound*). The visual signals as seen by a user are presented in Figure 5.11.

The motivation behind using a stop sign and a stop sound was to give a short and simple command known to everybody. In everyday situations, it is common to stop a walking person by talking to them. However, the visual system is likely to dominate other senses during a VR experience potentially making visual signals more effective. We tested a visual and an auditory stop command separately to investigate this reasoning.

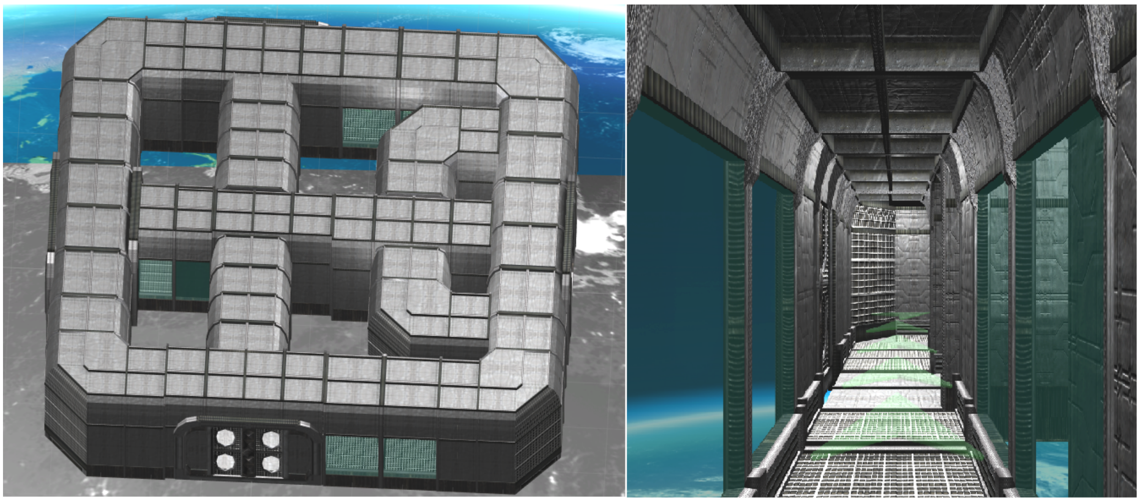

Figure 5.12: The VE used in the stopping test.

The figure of a virtual human was used as the second visual signal. The virtual human was used in connection with our collision warning strategy that was explored in the second part of the experiment. The assumption behind the choice of a virtual human was that a walking person would naturally stop if they saw a human appearing in front of them. In a real application, a virtual human would appear at the position of a user the collision with whom needed to be averted thus identifying the location of a possible collision.

Finally, the clap sound was used as a comparison for the auditory signal containing the stop command semantics (stop sound).

The position of the virtual human was fixed in the VE whereas the stop sign was always shown at 1.5m distance in front of a user. This way, it was clearly visible even when the user kept walking for some time after receiving the stopping signal. Both sound signals were non-spatialized.

Environment and experimental task The main task for a participant was to walk in a virtual corridor and stop immediately when they saw or heard a stopping signal. Each participant was asked to walk at a comfortable pace and to not slow down in the anticipation of a signal. The VE where the task took place is shown in Figure 5.12. The VE was designed to look appealing enough to not disappoint participants but at the same time to be simple and non-distracting. The size of the VE is 8x11 m, the width of the corridor is 1.2m. We chose a space-themed environment to avoid possible comparisons with the real world.

The signals were invoked by invisible triggers placed at different locations in the middle of the virtual corridor. When a participant reached a pre-defined distance towards each trigger, a stopping signal was given. Thus, the triggers represented a second user the collision with whom would need to be prevented by stopping the walking participant.

Each participant performed two walking sequences with all four stopping signals being given in each of them. The order of the signals was different between the sequences. The distance at which the triggers invoked the stopping signals was 2 m in the first sequence and 1 m in the second sequence. Effectively, this difference only affected the *VirtualHuman* signal that appeared closer to a participant in the second walking sequence.

Measures The main quantitative measure was the distance that a user walked before the full stop after a stopping signal had been issued. We chose to use distance over time that it takes a participant to stop completely as it is a metric that can be directly used as a reference for the design of multi-user VR experiences.

The stopping distance was measured as a difference between the tracked position of a user when a stopping signal was given and the position of the user registered when the user fully stopped.

Additionally, subjective data on the signal perception was collected in a post-test questionnaire. Participants were asked to rate how difficult it was to stop following each signal on a Likert scale ranging from 1 to 7 where 1 meant "very easy" and 7 meant "very difficult". Participants also indicated their preferred signal.

Part 2: Displaying notification avatars

In the second part of the experiment we explored our alternative strategy of collision prevention. This strategy expand the simple "ghost avatar" approach that we used in the demonstration of non-shared virtual spaces during the development of ImmersiveDeck.

The ghost avatar used in ImmersiveDeck was a simple semi-transparent box of a user height that was constantly displayed at the positions of users who were on a different level within the virtual scene and therefore not visible as standard avatars. This way, users where always aware of the positions of others within the tracking space.

However, it may be not necessary to notify walking users of the positions of others all the time. Constantly seeing ghost avatars can be distracting for users. If users are dispersed within the tracking space the danger of collisions is minimized and there is no need to display any notifications. Therefore, we decided to modify our strategy by only notifying the user of a potential collision if another user is sufficiently close. A notification avatar is displayed for each user during the time when users are within a threshold distance from each other but not when the distance between them is larger. This way, each user can decide on their own how to avoid a potential collision, by either stopping or just slowing down or modifying their trajectory.

The first question of our strategy is which avatars to use for notification.

Secondly, we hypothesized that different scenarios might require different types of notification avatars. It was decided to evaluate the strategy in a more ecological situation where users had a specific task in a VR environment.

It was important to conduct the experiment with two users at a time so that a collision avoidance task could be real.

This experiment had a 4×2 mixed design, with the within factor being the notification avatar and the between factor the type of the VE.

Avatars Avoiding collisions with others while walking is a natural task daily performed by everybody. A human figure displayed at the position of an approaching colocated user would likely allow to avoid collisions in a most intuitive way. However, a human avatar is likely to produce an illusion of another person being in the same VE whereas its intended use is notification only. A different solution would be to use a very generic non-human shape that would only identify a currently occupied volume in the VE. Such an avatar is unlikely to induce a feeling of the VE being shared with somebody else, although it would be less familiar to users and might cause confusion. To explore the range between these two border cases, we introduced four levels of anthropomorphism and realism of a notification avatar, ranging from a photo-realistic human to a completely non-human shape.

We used two human avatars: a photo-realistic avatar from the first part of the experiment (*Human*) and a detailed but cartoon-styled and less realistic avatar (*Spaceman*). A lower level of avatar anthropomorphism was represented by a robot (*Robot*), obviously non-human, faceless and featuring simplified geometry. Finally, we used a semi-transparent cylindrical shape (*Shape*) that did not resemble a human at all. All avatars are shown in Figure 5.13.

Avatars were displayed at respective participants' positions when the distance between two participants reached the minimum threshold set to 2 m. Several pilot tests were conducted prior to the main experiment to determine this minimum threshold. During the pilot tests, different display modes of the notification avatars were tested as well. In particular, we tried using semi-transparent anthropomorphic avatars as well as a fade-in mode in which avatars first appeared as semi-transparent and progressively changed their rendering mode to fully opaque over a short time. However, semi-transparency and fading in were negatively met by pilot test users; fully opaque rendering for anthropomorphic avatars was chosen for the main experiment.

The avatar's position and rotation was set according to the position and head rotation of the participant whose proximity the avatar notified. Avatars' limbs were not animated.

Walk group The main task for participants from the *Walk* group was to walk following pre-defined paths inside the corridors of the same VE that was used in the first part of the experiment. The paths were displayed with semi-transparent arrows in the middle of the corridor.

Figure 5.13: Notification avatars used to test the collision notification strategy.

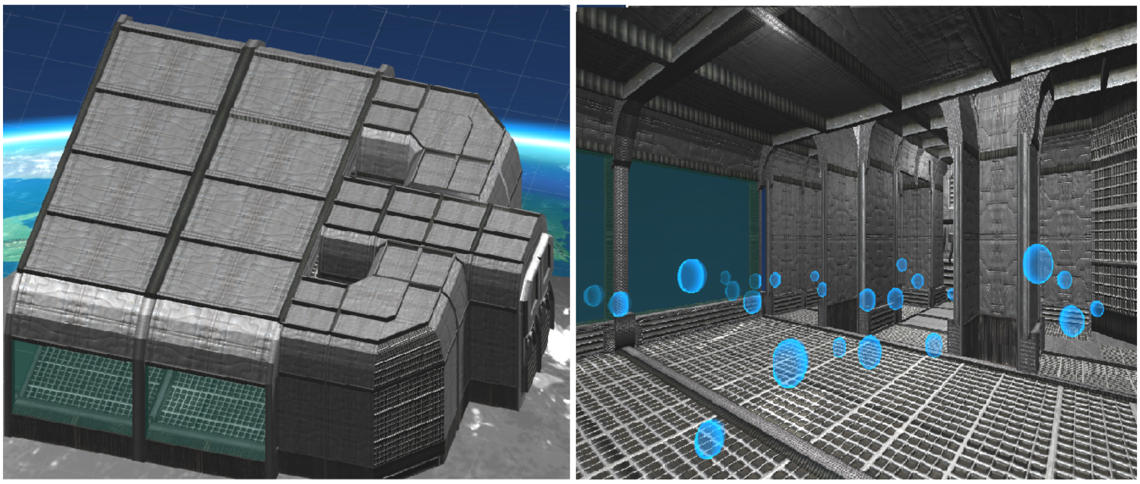

Figure 5.14: The VE used for participants from the *Game* group.

The paths for both participants from a pair were arranged symmetrically and had several common segments. This way, the participants were guaranteed to approach each other sufficiently to see a notification avatar several times during a walking sequence. Each pair of participants performed four walking sequences, each time with a different notification avatar.

Game group The game was set in the VE shown in Figure 5.14, 8 x 8 m large. The virtual space was filled with blue spheres, except for one colored in pink. Participants were instructed to collect spheres by walking into them, but only the pink sphere could be collected each time. When a participant walked into the pink sphere, it disappeared and a new sphere turned pink. The game was finished when all spheres were collected.

The spheres were not shared between the players. Each participant saw the same initial set of spheres, and the same sphere was pink in the beginning. The next sphere turning pink for each participant was chosen in the following way. A position was calculated that

Table 5.8: Questions used to assess presence and experience of collision notification strategy.

QUESTION CODE	QUESTION
Pr1	How aware of real events occurring in the room around you were you during the game?
Pr2	Rate your sense of being there in the virtual room.
Pr3	When you think about the virtual playground, was it rather like something that you saw or somewhere that you visited?
Pr4	There were times during the experiment where the virtual environment became reality for me.
Surprised	To what extent were you surprised when you saw the notification avatar?
Scared	To what extent were you scared when you saw the notification avatar?
Disturbed	To what extent were you disturbed when you saw the notification avatar?
Associate	To what extent could you associate the notification avatar with your test partner?

was equidistant from both players. Then, all remaining spheres were found that were located within 2 m radius from this equidistant point. One of these spheres was then chosen to change its color into pink. If there were no spheres found in this area, one of the remaining spheres was turned pink at random. With this approach, participants were drawn to walk towards each other when trying to collect the next sphere but did not stay too close to each other during the whole game. As in the *Walk* group, players saw notification avatars at each others' respective positions when the distance between them was lower than 2 m.

Each pair of participants from the *Game* group performed one game sequence in which all four notification avatars were used. The notification avatar was chosen in a pseudo-random manner each time when the participants from a pair came closer than a threshold distance to each other.

Measures Subjective questionnaires were used to assess participants' experience with the notification strategy.

As in the experiment on mutual awareness, presence-related questions were used to control for possible between-group differences in the experiment outcome (questions *Pr1* - *Pr4* in Table 5.8).

In a post-test questionnaire, participants were asked to rate how surprised, scared and disturbed they felt when they saw each avatar, as well as how well they could associate each avatar with their test partner (precise formulations can be further seen in Table 5.8). The aim of questions *Surprised*, *Scared* and *Disturbed* was to assess how disruptive the appearance of each avatar was for user experience. With question *Associate*, we intended to determine whether participants were fully aware that a notification avatar represented their test partner or, in contrast, whether they saw a notification avatar rather as just a graphical entity within the VE. All four questions were asked about each of notification avatars. The answers were given on a Likert scale from 1 to 7. SSQ was administered only after the experiment.

In addition, an in-depth discussion with each pair of participants took place focusing on their experience of the notification strategy and preferences for notification avatars.

Apparatus

The ImmersiveDeck platform upgraded to Unity3D version 5.3.2 was used to develop the experimental scenarios. An updated version of the tracking algorithm allowed an update rate of 45 fps. The size of the walkable area was 14 x 9 m.

Motion capture suits were not used and participants did not have self-avatars. Participants did not wear headphones, so they could hear each other's and the test coordinator's steps and eventually voices. We were aware that such a setup would possibly lower participants' immersion. Nevertheless, we decided that participants would feel safer and more confident if they could easily communicate with the experiment coordinator, and possibly also hear steps, especially in the second part of the experiment.

Participants

29 volunteers (10 female and 19 male) in the age from 21 to 48 years (median 25) took part in the study. 14 participants had no previous experience of VR. The remaining 15 had either tried Oculus Rift before or taken part in a different user study in VR, including the ones with real walking. All 29 participants accomplished the first part of the experiment, the stopping test. 17 participants accomplished the second part of the experiment in the *Walk* group and 12 in the *Game* group. Figure 5.15 shows a pair of participants from the *Walk* group and a pair of participants from the *Game* group during the second part of the experiment.

Procedure

Participants came to the laboratory in pairs. First, the general course of the experiment was explained to participants. The setup was shown to participants and each participant from a pair was given time to freely walk in the environment with corridors used in the first part of the experiment. Then, instructions for the first part of the experiments were explained to the participants.

The stopping test was conducted with one user at a time. Stopping signals were demonstrated to each participant before the test. The second participant from a pair to do the stopping test waited in an adjacent room while the first participant was doing the test. After finishing the stopping test, each participant filled out a short questionnaire where they assessed the difficulty of each stopping signal, followed by a short discussion of the stopping test with both participants.

The general purpose of the experiment and the notification strategy used in the second part of the test were explained to the participants afterwards. The participants received instructions for the second part of the experiment, either *Walk* or *Game* scenario. Then the second part of the experiment took place, followed by the post-test questionnaire. The experiment finished with a discussion of the participants' experience and their impressions of the collision notification strategy.

Figure 5.15: Left: a pair of participants from the *Walk* group at a moment where they see notification avatars. Right: a pair of participants from the *Game* group during the experiment.

Table 5.9: Results of pairwise comparisons with *VirtualHuman* displayed at 1 m distance. Test statistic T, standardized statistic z, p-value and effect size r are shown.

	T	z	p	r
VirtualHuman 1m - VirtualHuman 2m	-1.750	-3.834	**.001**	-.55
VirtualHuman 1m - StopSign	-1.208	-2.647	.081	-.38
VirtualHuman 1m - StopSound	-2.167	-4.747	**<.001**	-.68
VirtualHuman 1m - ClapSound	-1.229	-2.693	.071	-.39

5.2.3 Results

Data analysis

The resulting data was not normally distributed. Non-parametric tests were used for data analysis: Mann-Whitney U test to compare two independent samples, Wilcoxon signed-rank test to compare two related samples and Friedman's ANOVA to compare several related samples. Differences are reported as significant at $p < 0.5$ level.

Part 1

The results were obtained from the data of 24 participants. Data of the remaining 5 participants was excluded due to missing values. The resulting stopping distance was averaged from the data of two walking sequences for *StopSign*, *StopSound* and *ClapSound* as the conditions for these signals were the same between two trials. The stopping

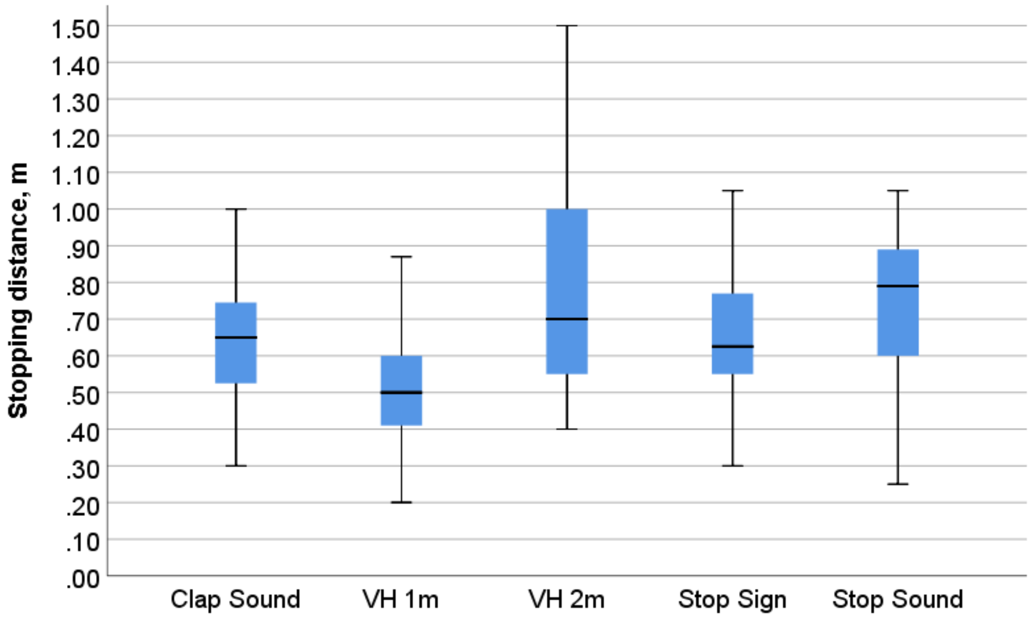

Figure 5.16: Boxplots of the resulting stopping distance for each of the tested signals.

distance was analyzed separately for *VirtualHuman* shown at 2 m (in the first walking sequence) and at 1 m (in the second waking sequence) in front of a participant. The resulting boxplots can be seen in Figure 5.16.

Friedman's ANOVA showed statistical significance of signal type for the stopping distance ($\chi^2(5) = 26.690$, $p < .001$). The follow-up pairwise comparisons with Wilcoxon matched-pair signed-rank test showed that *VirtualHuman* displayed at 1 m distance had the largest contribution to this result, leading to a significantly shorter stopping distance than two other signals. The details of these pairwise comparisons are presented in Table 5.9. Significance values were adjusted by the Bonferroni correction for multiple tests. Statistically significant comparisons are indicated with p-values displayed in bold.

The boxplots of the resulting scores of the subjective difficulty of each signal can be seen in Figure 5.17. The type of the signal was statistically significant for the perceived stopping difficulty in Friedman's ANOVA ($\chi^2(4) = 24.165$, $p < .001$). The follow-up pairwise comparisons showed that *ClapSound* was perceived as significantly more difficult than *StopSign* (T = -1.479, z = -3.969, $p < .001$, effect size r = -.21) and *VirtualHuman* (T = -1.021, z = -2.739, p = -.037, r = .15).

13 participants indicated *StopSign* and eight *StopSound* as the preferred stopping signal. Three participants found *VirtulHuman* to be the easiest signal. However, participants reported to have often walked into the figure shown in front of them (at 1 m distance)

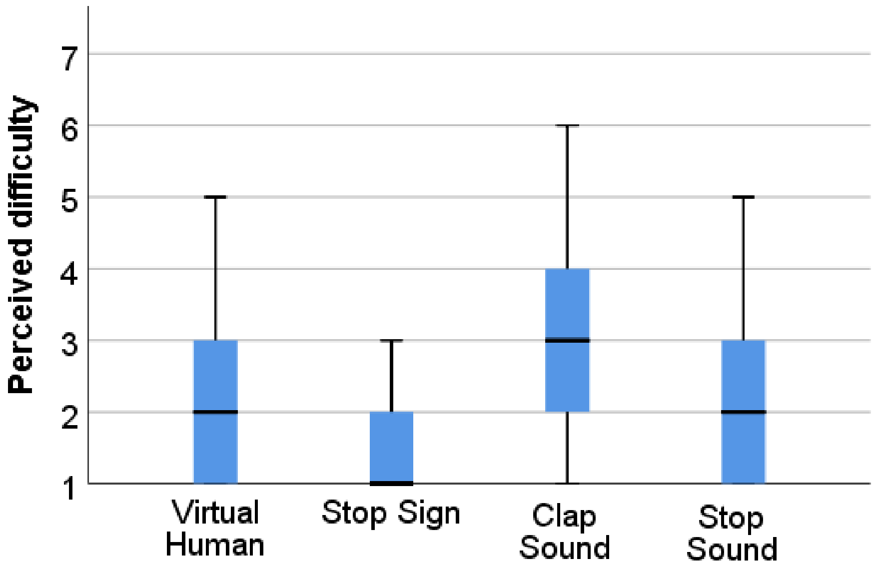

Figure 5.17: Boxplots of the resulting scores of the perceived difficulty of each signal.

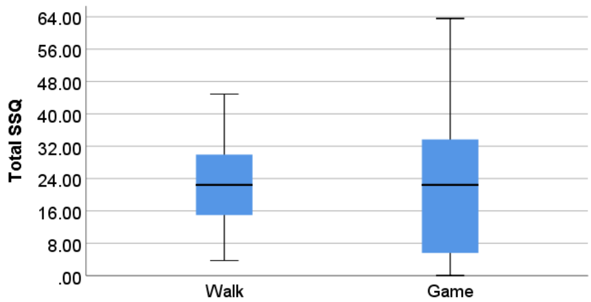

Figure 5.18: Boxplots of the post-test total SSQ score in the *Walk* and *Game* groups.

in the second walking sequence and judged it as being too close. *ClapSound* was not preferred by any of the participants, a result that well coincides with it being found significantly more difficult than the visual signals. Seven participants indicated that they found the visual signals generally easier, five said the same about the auditory signals.

Part 2

VR experience and simulator sickness None of the participants reported strong symptoms of simulator sickness. The total post-experiment SSQ score was not different between the *Walk* and *Game* groups (Mann-Whitney U test, $U = 93.5$, $p = .711$). The

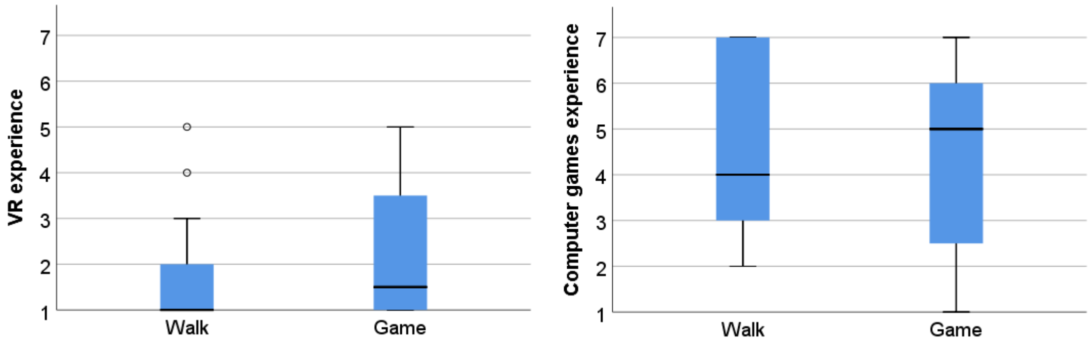

Figure 5.19: Previous experience of VR and playing computer games reported by the participants.

corresponding boxplots are presented in Figure 5.18.

The reported experience of VR and playing computer games was not different between the groups in Mann-Whitney U test (U = 121.0, p = .419 for VR experience, U = 97, p = .845 for experience of playing computer games). The corresponding boxplots are displayed in Figure 5.19.

Notification avatars Figure 5.20 presents the boxplots of of the resulting scores for questions *Surprised*, *Scared*, *Disturbed* and *Associate*.

In the *Walk* group, Friedman's ANOVA indicated statistical significance of the avatar type for questions *Surprised* (($\chi^2(4)$) = 8.286, p = .04), *Disturbed* (($\chi^2(4)$) = 10.180, p = .003) and *Associate* (($\chi^2(4)$) = 12.625, p = .006). The follow-up pairwise comparisons performed with the Wilcoxon signed-rank test did not show significant difference for all pairs. However, the effect sizes calculated based on the pairwise comparisons ranged from medium to large in many cases. For the questions *Surprised* and *Disturbed*, the largest effects were found in comparisons with *Shape*. For *Associate*, effect sizes grow with the increase of the difference in the degree of anthropomorphism of the compared avatars: from small for *Human* - *Spaceman*, *Spaceman* - *Robot*, *Robot* - *Shape*, to median for *Human* - *Robot*, *Spaceman* - *Shape*, to large for *Human* - *Shape*. The details of the pairwise comparisons can be seen in Table 5.10.

In the *Game* group, Friedman's ANOVA showed significant differences in the scores for *Scared* (($\chi^2(4)$) = 8.556, p = .036). There were no statistically significant differences in the follow-up pairwise comparisons.

Presence Participants from the *Walk* group were more aware of real events (question *Pr1*) happening in the tracking laboratory than the participants from the *Game* group, according to the Mann-Whitney U test performed on data without outliers (U = 21.5, z

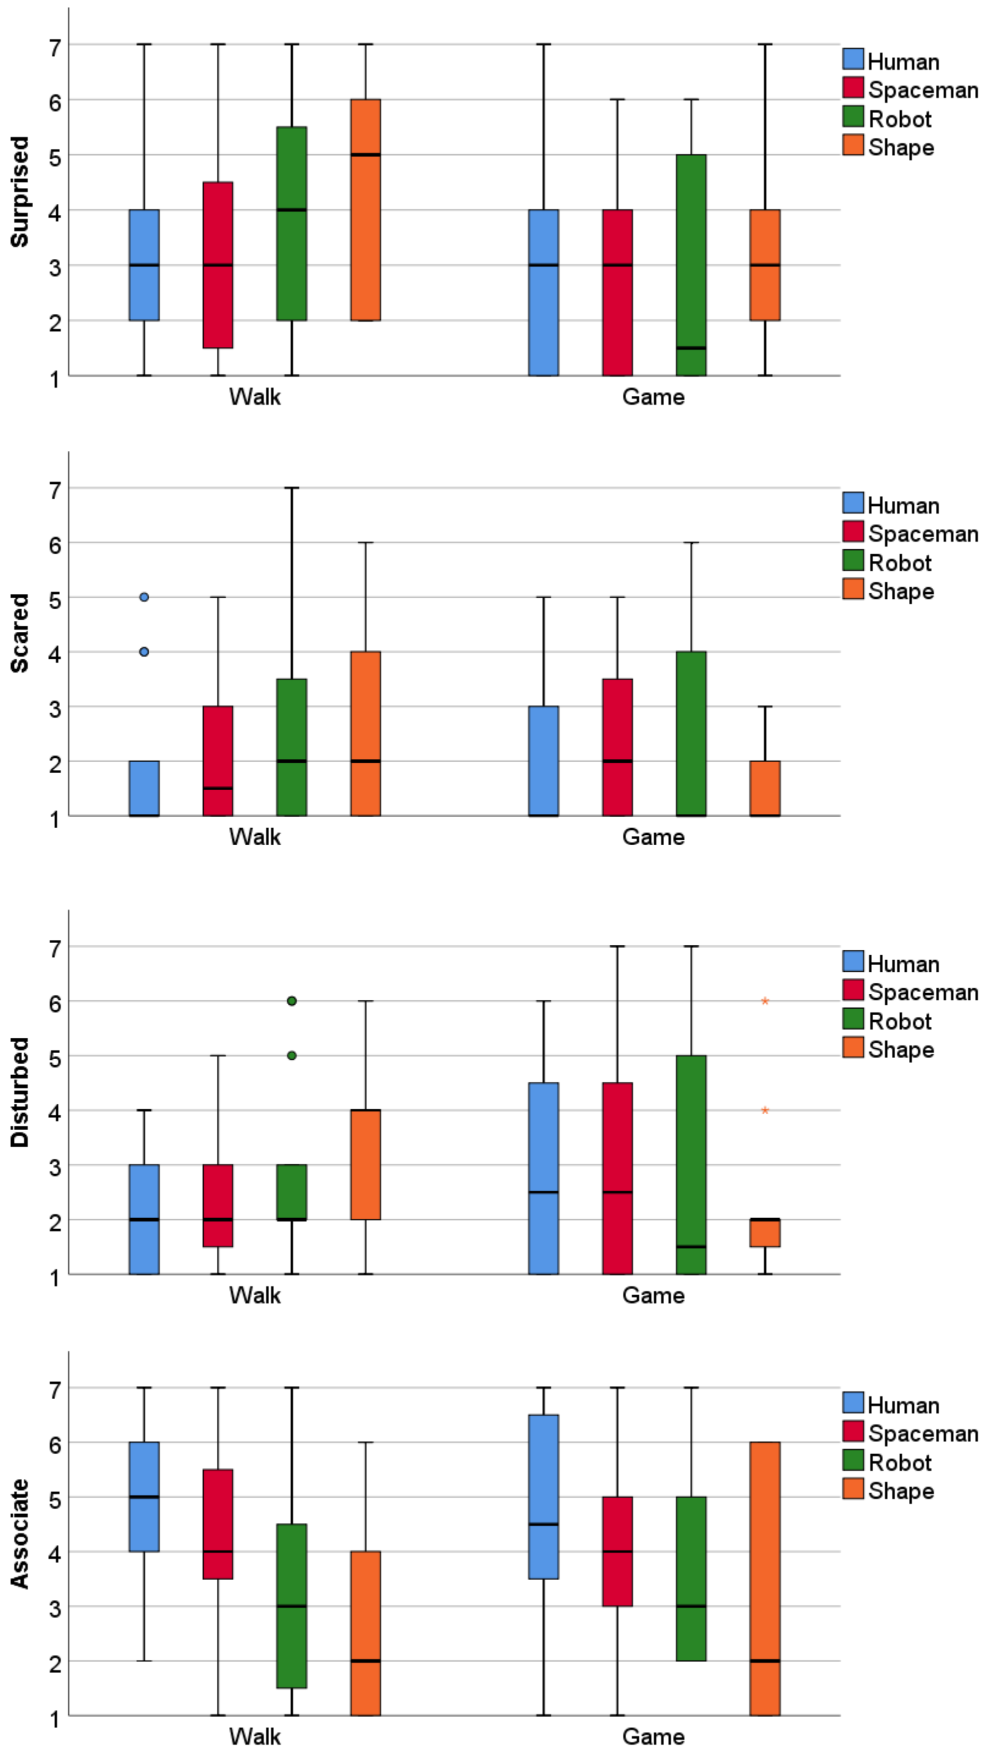

Figure 5.20: Resulting scores of four questions used to explore the notification strategy with each avatar type. 1 corresponds to "not at all" and 7 to "very much" on a Likert scale.

Table 5.10: Results of the post-Friedman's ANOVA pairwise comparisons in the *Walk* group.

Surprised	*T*	*z*	*p*	*r*
Human - Spaceman	-.167	-.316	1.0	-.06
Human - Robot	-.333	-.632	1.0	-.13
Human - Shape	-1.167	-2.214	.161	-.46
Spaceman - Robot	-.167	-.316	1.0	-.06
Spaceman - Shape	-1.0	-1.897	.347	-.39
Robot - Shape	-.833	-1.581	.683	-.32
Disturbed	*T*	*z*	*p*	*r*
Human - Spaceman	-.250	-.474	1.0	-.09
Human - Robot	-.292	-.553	1.0	-.11
Human - Shape	-1.125	-2.135	.197	-.43
Spaceman - Robot	-.042	-.079	1.0	-.02
Spaceman - Shape	-.875	-1.660	.581	-.33
Robot - Shape	-.833	-1.581	.683	-.50
Associate	*T*	*z*	*p*	*r*
Human - Spaceman	.414	.791	1.0	.16
Human - Robot	.875	1.660	.581	.34
Human - Shape	1.375	2.609	.055	.54
Spaceman - Robot	.458	.870	1.0	.18
Spaceman - Shape	.958	1.818	.414	.37
Robot - Shape	.5	.949	1.0	.19

= -2.258, p = .013, effect size r = -.52). This difference was possibly due to the more involving nature of the experimental task in the *Game* group. Many participants reported that they would have rated their awareness of real events lower if they had not heard steps during the experiment. However, they reported to have stopped paying attention to steps after some short adjustment time.

Scores of the remaining presence-related questions were not different between the *Walk* and *Game* groups in Mann-Whitney U test. The corresponding boxplots can be seen in Figure 5.21.

Comments of participants In the *Walk* group, eight out of 17 participants strongly preferred to see *Human* as a notification of the proximity of their test partner. Another four participants preferred any human-shaped avatar. Only one participant from this group preferred to see *Shape* as a notification, whereas five participants strongly disliked it.

In the *Game* group, nine out of 12 participants preferred to be notified of the proximity of their test partner by *Shape*. However, none of them disliked any other avatar.

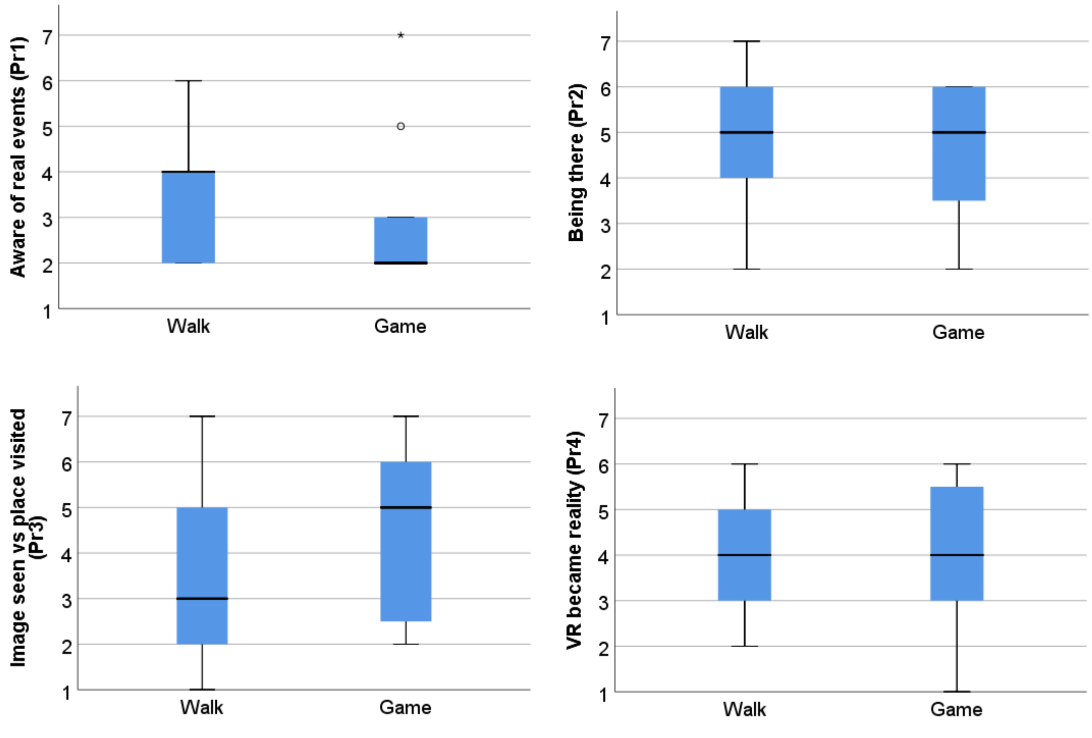

Figure 5.21: Boxplots of the resulting scores of the self-report of presence.

Those participants who preferred *Human* and human-like avatars said that human figures looked natural and expected. They also helped to avoid collisions with a test partner as they indicated the direction of their movement. The non-human shape was described as "unnatural" by those participants who strongly preferred *Human* as a notification. Three participants reported that they could get used to every avatar except *Shape*.

Those participants who preferred *Shape* avatar felt less distracted from fulfilling their task when seeing *Shape* as compared to other notification avatars. They reported that they did not feel like they had to look at *Shape* for too long while the visuals of human-like avatars attracted their attention making them to pause the main task. Furthermore, the semi-transparent material of *Shape* allowed to see what was going on in the game in the area behind it.

5.2.4 Discussion

Stopping users

The results of the first part of the experiment indicated that only *VirtualHuman* displayed at 1 m distance in front of a walking user was more efficient for stopping than other signals. However, 1 m distance resulted in the signal being reported as uncomfortably close. Two visual signals displayed at 1.5 m and 2 m distance in front of a user and two auditory signals were equally effective in stopping participants. On average, participants walked between 0.5 m and 0.8 m before the full stop after a stopping signal had been given, although the longest distance a participant walked after a signal had been given was 1.5 m (with *VirtualHuman* displayed at 2 m distance). Thus, in a situation where users are walking without seeing each other on paths that potentially lead to a collision, they should be stopped no later than when the distance between them reaches 2 m. Our test was conducted in a fully immersive but not interactive VE where the only task was to walk and stop when a signal is given. Therefore, our results are likely to reflect users' fastest reaction times.

Even though the tested signals were similarly effective, they were differently perceived by users. A sound signal not carrying any command (*ClapSound*) was found to be more difficult than visuals. The absolute majority of participants preferred a stop command to be used, although their preferences were divided between its visual and auditory version. In general, we observed a strong split of preferences for either visual or auditory signals.

In a virtual world filled with dynamic content and various game sounds, the best practice for quickly stopping users might be a combination of a visual and an auditory signal. However, a follow-up study in such a sensory-intense scenarios might be required to confirm this conclusion.

Notification strategy

The notification strategy used in the second part of the described experiment proved successful for imminent collision prevention: all participants could avoid their test partners when they saw notification avatars for short periods of time. Users' preference for a notification avatar was strongly affected by the type of the task performed in the virtual world. A human or at least an anthropomorphic notification avatar was found to suit better an exploration task whereas a featureless non-anthropomorphic shape was positively met in a scenario with a specific goal and interactive content.

The core idea of our notification strategy was to provide a non-distracting yet effective visual guidance for collision avoidance. The ideal collision avoidance experience that we intended for a user would be to see a notification avatar, recognize that it was a notification about someone else close in the same physical space, walk past and immediately forget about the encounter. The results suggest that the strategy might be better suited for scenarios similar to our *Game* task, where users' attention is concentrated on tasks but not the VE itself.

While notification avatars were effective for collision prevention in the exploration scenario performed by the participants from the *Walk* group they inevitably attracted participants' attention. This fact is reflected in the significant differences found by Friedman's ANOVA for avatar comparisons in the *Walk* group. In the *Walk* group, the scores for *Surprised*, *Scared* and *Associate* questions reflected the decrease in the level of realism and anthropomorphism from *Human* to *Shape*, with the most abstract avatar, *Shape*, provoking the strongest reactions. While the boxplots of the scores (Figure 5.20) and the effect sizes (Table 5.10) indicate a correspondence between the level of the anthropomorphism and the perception of each tested avatar by the participants from the *Walk* group the post-Friedman's ANOVA pairwise comparisons were not statistically significant. We believe that this fact reflects the limitation of the sample size in our experiment. We believe that more data would allow to find statistical significance in post-ANOVA pairwise comparisons that produced effects of medium to large size. The impact of the degree of avatar anthropomorphism could then be established in a more detailed manner.

However our experiment lead to a following main conclusion. In situations where users' attention is attracted to a notification avatar they prefer it to have a familiar human appearance. In this case, even small differences in the avatar appearance are likely to be noticed by users and may provoke strong emotional responses. In scenarios with other goals than exploration, on the contrary, the notification avatar is required to have as few visual details as possible. In this case, the non-distracting nature of the visual feedback is more important than familiarity.

The strategy of notifying users about possible collisions with colocated others was revisited in two experiments published later. A pilot study of Scaravelli and Teather [118] compared an anthropomorphic (although not realistic) avatar, a wireframe box (comparable to our *Shape*) and camera overlay of HTC Vive in the task of collision avoidance of a simulated colocated user. The experimental task consisted of walking between pre-defined positions in a simple non-interactive virtual environment. The anthropomorphic avatar and camera overlay were strongly preferred compared to the bounding box that provoked negative reactions among participants. This result is close to the result of our *Walk* task in which the human notification avatar was preferred. In the experiment of Langbehn and Steinicke [81], a semi-transparent humanoid (although moderately realistic) shadow avatar was used to visualize a simulated colocated user. Two conditions were compared: in the first one, the shadow avatar was visible all the time; in the second one, it appeared at 1.25 m distance from the user and became progressively more opaque as the user approached closer. As in the experiment of Scaravelli and Teather, the main experimental task was to walk towards a goal within the VE, the task that is similar to our *Walk* condition. The constantly visible avatar was strongly preferred by participants of the user study and produced fewer collisions. The walkable space used in the experiment was 4 x 4 m, several times smaller than the 14 x 9 m tracking area used in our experiment. Therefore, even when the shadow avatar was displayed constantly it was never very far away from a user. This result might indicate the existence of a trade-off between navigation safety and introducing possible disturbances in the user experience: while room-sized environments require constant visual feedback to prevent potential collisions between users, such feedback could be limited to a closer range around the user in large-scale setups. Determining such a range in which comfortable collision avoidance through visual feedback can be ensured is a task for future research.

5.3 Summary

In this chapter, we presented two experiments aimed at exploring the use of colocated non-shared VR. In the first one, we investigated mutual awareness of users colocated in the same tracking space but being immersed into individual VEs without seeing or hearing each other. In the second experiment, we explored strategies of preventing user-on-user collisions in colocated non-shared VR. Both presented experiments were conducted before the release of HTC Vive, an affordable HMD that is accompanied by Lighthouse tracking technology allowing an easy setup of a room-scale walkable VR experience. At the time, even single-user walkable VR setups were rare, whereas large-scale multi-user VR systems only existed at few select laboratories. Therefore, the presented experiments were among the first published explorations of colocated non-shared VR. The following section summarizes the findings of both experiments.

- **Users of colocated non-shared VR are not aware of each other's presence in the shared physical space even in the immediate proximity**. Only a third of all participants who were specifically instructed to try and locate their test partners in the shared tracking space could successfully do. High mutual awareness of these select participants might be connected to the fact that they could trace their own movement in the tracking space and make guesses about the movement of their test partners, whereas the remaining participants reported to have lost track of their location in the real environment. However, those participants who did not receive the specific instruction to locate their test partners were mostly absolutely unaware of their respective location during the experiment, even when both participants from a pair were right next to each other.

- **If users need to be urgently stopped, visual and auditory signals are equally efficient although participants' preferences are divided between these two categories**. Although two visual and two auditory signals were equally efficient in terms of the measured distance that participants walked before the full stop, the auditory signal not containing the "stop" command received higher scores for the perceived difficulty than other signals. The absolute majority of participants preferred a "stop" command to be contained in the signal regardless of whether it was visual or auditory. A combination of a visual and an auditory signal might be appropriate to provide a preferable cue to all users.

- **Notification avatars may be used as visual feedback for immediate collision prevention**. We suggested a collision prevention method based on visual feedback of a colocated user. In this method, a notification avatar is displayed at the position of a colocated user if the distance between two users is below a collision-critical value. As opposed to stopping both involved users at a danger of an immediate collision, this method provides users with information about the location of the potential collision and is less likely to introduce breaks in presence.

- **The suitability of a particular type of the notification avatar is determined by the application and user tasks in it.** In an exploration scenario without any specific task participants preferred to see a human notification avatar whereas an abstract avatar provoked strong negative reaction. Although the goal of the notification strategy was to be as least distracting as possible notification avatars inevitably attracted participants' attention during the exploration task. In the scenario in which a specific sufficiently engaging task was provided, on the contrary, the abstract notification avatar with few prominent visual features was favored on the account of being less distractive than anthropomorphic avatars.

CHAPTER 6

Shared VR in Colocated, Distributed and Mixed Colocation Setups

This chapter is dedicated to the investigation of shared walkable VR scenarios. Two experiments on shared VR are described. In the first experiment we investigate proxemics and locomotion patterns of pairs of users walking within a shared VE. Proxemics behavior demonstrated in shared VR in a colocated and in a distributed setup is compared to that of a real environment. In the second experiment we study proxemics patterns and copresence experienced by triples of users in a mixed colocation setup where users are split between two tracking areas.

6.1 Proxemics and locomotion in shared VR

6.1.1 Motivation

The process of walking together in either real or virtual environments includes the task of avoiding collisions with colocated walkers. In the real world, this task is so natural and well-trained for any individual that it is normally not perceived as a difficulty. We easily navigate around obstacles and avoide collisions with other persons during everyday tasks by adapting our movements based on mutli-sensory feedback from our visual, vestibular and proprioceptive systems. In VR, however, visual cues produced by a rendering system differ from those that we normally receive from our visual system, leading to differences in perception of virtual space. For example, multiple studies have demonstrated that distances are underestimated in immersive virtual environments [85, 143].

Since the perception of space in VR differs from the perception of the real space there is a reason to expect that proxemics and collision avoidance behavior in VR might be different from the real world as well. The purpose of the experiment presented in this section was to investigate whether such differences in locomotory behavior really exist.

What is the practical importance of this research question for the development of shared VR experiences? First of all, some multi-user VR applications where users move quickly and a lot may require the exact coordination of users' movements. Such applications include team rescue training scenarios, sports simulation and play-throughs, performing arts rehersals. Furthermore, the designers of large-scale multi-user VR experiences would need the information about collision avoidance behaviour in virtual environments to estimate how many users can be accommodated simultaneously in a restricted physical tracking space.

While previous research addressed the topic of collision avoidance behaviour with static obstacles in walkable VEs (see Chapter 3), there is a gap in the understanding of mutual collision avoidance of users simultaneously walking in VR. We intended to bridge this gap.

6.1.2 Experiment design

Our goal was to compare locomotor trajectories of users simultaneously walking in the matched real and virtual environments. Since we were to produce a comparison with real-world behavior we intended to achieve virtual experience of high fidelity. Collision avoidance in the real world is based on body-to-body interaction where physical bodies of interacting walkers provide a rich continuum of cues for the interpretation of the future intent of the other. Thus, high fidelity primarily meant providing users with self-avatars that would precisely repeat the movements of their real bodies.

Shared VR can be experienced in colocated or distributed physical setups. Both scenarios were investigated in the experiment. Our general hypothesis was that users would display more cautious proxemics behavior in colocated shared VR compared to the real setting, based on previous research describing collision avoidance of static obstacles in VR. However, we did not have a hypothesis for user behavior in the distributed setup.

In the real world, collision avoidance is never the main aim of a walker but rather a minor task that they have to perform on the way to their chosen destination. With the aim to preserve ecological validity of the test conditions we did not explicitly ask participants to avoid their test partners during the experiment instead providing them with a goal within the environment.

The experiment had a mixed design, with two within-subject independent variables and one between-subject variable. Within-subject variables were the type of trajectory and the type of physical colocation. We tested two collision trajectories, *Frontal* and *Crossing* and three colocation conditions, *Real*, *VirtualColocated* and *VirtualDistributed*. The between-subject variable was the order in which *VirtualColocated* and *VirtualDistributed* conditions were tested.

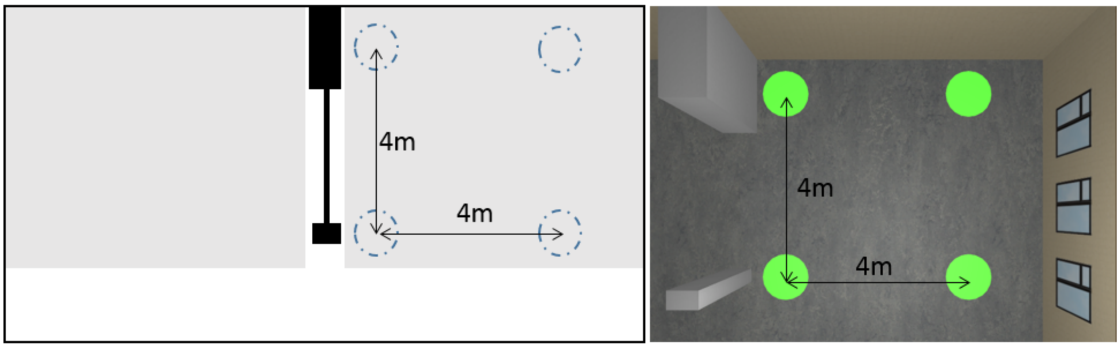

Figure 6.1: Scheme of the laboratory with two tracking areas depicted in grey color(left) and the top view of the matching VE.

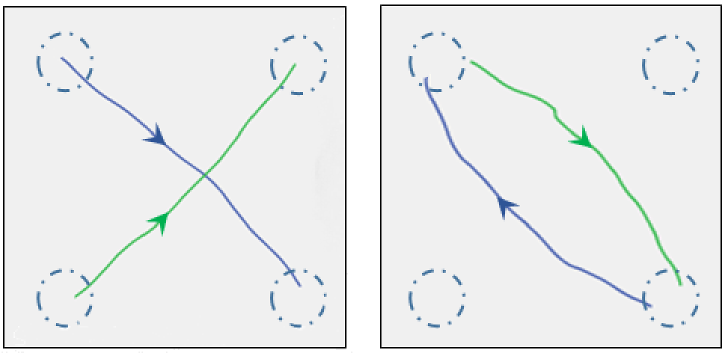

Figure 6.2: *Crossing* (left) and *Frontal* (right) walking trajectories.

Environment and tasks

The experiment took place in a laboratory with two symmetric tracking areas separated by a wall section and a curtain. Figure 6.1 shows a schematic view of the lab.

In the *Real* condition, each pair of participants walked between markings on the floor of one of the tracking areas forming a square with the 4 m long side. The task for each participant from a pair was to walk to the floor marking diagonal from their starting position. In each sequence of the *Real* condition, both participants started walking simultaneously at the signal of the experiment coordinator. The experiment coordinator also instructed them to change starting positions when necessary. Starting positions of the participants from a pair where either on one side of the square formed by the floor markings leading to a *Crossing* collision course for the participants or on the opposite sides of a diagonal leading to a *Frontal* collision course. Both collision courses are illustrated in Figure 6.2. A pair of participants walking on *Frontal* and *Crossing* trajectories in *Real* condition can be seen in Figure 6.3.

The shared VE used in the *VirtualColocated* and *VirtualDistributed* conditions matched

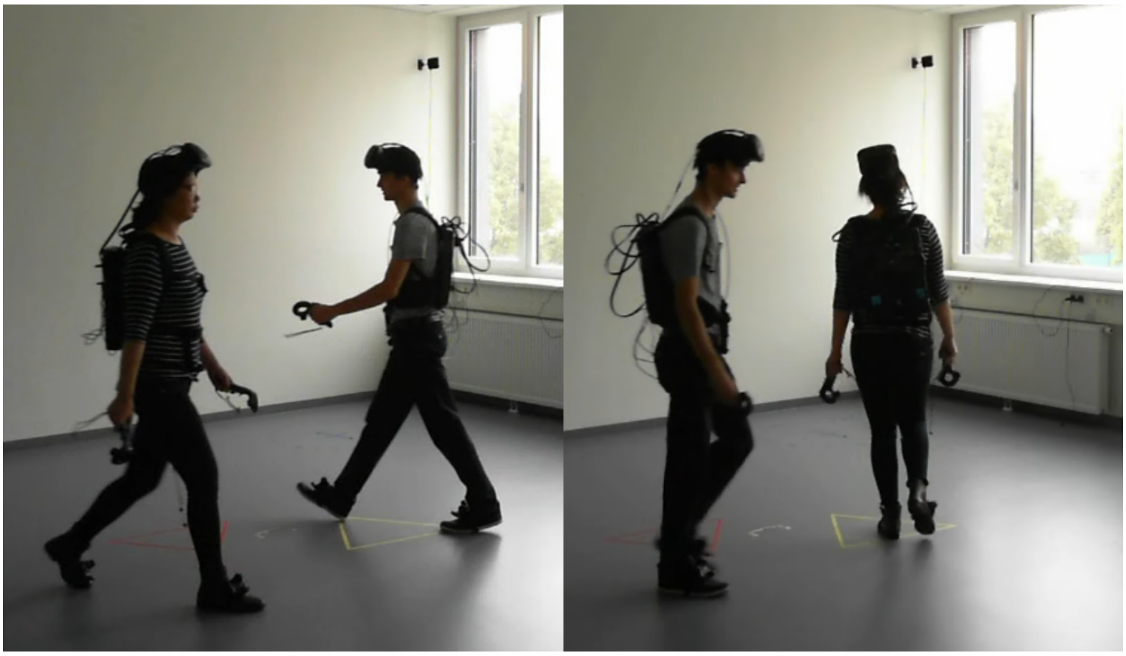

Figure 6.3: A pair of users in the *Real* condition of the experiment.

the real laboratory. In the *VirtualColocated* condition, both participants from a pair shared the tracking area with the floor markings. They walked between virtual markings (Figure 6.1) the positions of which were exactly mapped to the positions of the floor markings in the tracking area. In the *VirtualDistributed* condition, participants were performing the experimental task in the same shared VE as in the *VirtualColocated* condition while being located in the separate tracking areas. In this condition, there was no possibility of a physical collision.

In both virtual conditions, the task was to collect items appearing in the virtual scene. In the beginning of the test, each participant saw a green dot on the floor corresponding to the position of one of the four floor markings. Participants were instructed to take their starting position on this dot. When both participants from a pair where at their starting positions a collectible item appeared in each participant's view at the position of a diagonal marking. Then, a sign saying "WALK" appeared simultaneously for each participant and they started walking across the tracking space to collect the item. The trajectories of the participants were either *Frontal* or *Crossing* depending on the relative arrangement of starting points. After the item was collected a new starting point appeared for each participant. Each participant always saw only their individual starting point and collectible item. Figure 6.3 shows a pair of users walking in *VirtualDistributed* and *VirtualColocated* conditions.

Participants had full-body human avatars of the matching gender and height in the virtual conditions. The avatars were animated with a motion capture software that

Figure 6.4: A pair of participants walking in the *VirtualDistributed* (left) and the *VirtualColocated* (right) colocation conditions.

Figure 6.5: First-person view of a test partner during walking on a *Frontal* (left) and *Crossing* (right) collision course.

produced accurate real-time poses. Full-body motion capture allowed each user from a pair to see their own movements and movements of their test partner as they would have been during walking in the real environment. Figure 6.5 illustrates a participant's avatar as seen by their test partner in the virtual conditions. In the virtual conditions, sounds coming from the real world were blocked by earphones in which rain-like white noise was playing. Sound blocking was introduced to mimic the conditions of a likely multi-user application where in-game sounds would prevent users from hearing noise from the real surroundings.

Participants performed five *Frontal* and five *Crossing* walks in each colocation condition. *Frontal* and *Crossing* collision courses were randomised within each colocation condition.

The between-subject variable was the order in which participants from a pair performed the experimental task in *VirtualColocated* and *VirtualDistributed* colocation conditions, splitting all participants into the *Co-First* and *D-First* groups respectively. The *Real* condition was tested either before both virtual conditions or between them or in the end, all three variants being counter-balanced.

Measures

Trajectory metrics Head tracking data of participants was recorded when they performed *Crossing* and *Frontal* collision courses. The following characteristics of participants' trajectories were used to describe their locomotor behavior.

Clearance is the shortest distance between two participants from a pair during a walk from the starting position towards a goal.

Radius of curvature at each point of the trajectory is the radius of a circle tangent to the trajectory in this point. It is obtained with the following formula

$$r = \frac{(\dot{x}^2 + \dot{z}^2)^{1.5}}{\dot{x}\ddot{z} - \ddot{x}\dot{z}} \tag{6.1}$$

where x and z are Unity3D world coordinates in the horizontal plane.

To characterize the overall curvature of a trajectory we calculated the median value of the radius of curvature. The median curvatures of the trajectories performed by two participants from a pair are likely to be interdependent since avoiding a collision is a collaborative effort. For example, if one participant walks straight to their goal without modifying their path at all their test partner is forced to take a larger detour to avoid a collision. Therefore, we used the sum of median radii of curvature as a cumulative measure for both participants from a pair and denoted this measure by *SumOfMedianCurvatures*.

Just as path curvatures, walking speeds of two participants from a pair are likely to be correlated, especially on *Crossing* trajectories. For example, if one of the participants from a pair starts walking faster to pass first, the other one has to slow down to avoid a collision. Therefore, we calculated the mean walking speed of each participant from a pair during one test walk and considered the sum of their mean speeds, denoted *SumOfSpeeds*, as a cumulative measure in the evaluation.

Questionnaires The pre-test questionnaire contained questions on demographics, previous experience of VR and playing computer games.

Table 6.1 summarizes the post-experiment questionnaire.

Table 6.1: Questions from the post-experiment questionnaire.

QUESTION CODE	QUESTION
Pr1	How aware of real events occurring in the room around you were you during the game?
Pr2	Rate your sense of being there in the virtual room.
Pr3	When you think about the virtual playground, was it rather like something that you saw or somewhere that you visited?
Pr4	There were times during the experiment where the virtual environment became reality for me.
CoPr1 (Colocated and Distributed)	To what extent did it feel like your test partner was in the same virtual environment?
CoPr2 (Colocated and Distributed)	To what extent were you worried that you would collide with your test partner?
Atn (Colocated and Distributed)	Were you rather focused on the task or on avoiding your test partner?

Presence-related questions *Pr1 - Pr4* were used to assess the overall experience of users as well as to track to verify whether participants from *Co-First* and *D-First* had comparable experiences.

Questions *CoPr1* and *CoPr2* were used as an exploratory measure of copresence. These questions were asked separately for the *VirtualColocated* and *VirtualDistributed* conditions. The answers were recorded on the Likert 1 to 7 scale. Question *Atn* was used to assess the attentional allocation of participants during the experiment and was also asked separately the *VirtualColocated* and *VirtualDistributed* conditions. The possible answers for this questions were "task", "avoiding collisions" or "both"; participants were encouraged to add their detailed comments.

Apparatus

The experimental application was developed within our ImmersiveDeck platform used with an HTC Vive HMD and Lighthouse tracking technology. Unity3D version 5.6.3 was used.

Each participant was equipped with an HTC Vive HMD, two Vive controllers held in each hand and three Vive trackers placed on each foot and the hip for motion capture, a VR backpack and a pair of earphones for the virtual conditions. The used VR backpacks were an MSI VR One and an XMG Walker, each containing NVIDIA GeForce GTX 1070 graphics card and Intel Core i7 quad-core processor.

Motion capture was achieved with the use of IKinema Orion[1] inverse kinematics solution that reconstructed the pose of a participant's avatar based on six tracked points.

The experiment coordinator observed the progress of the experiment on the server machine. "WALK" signs for each walking sequence in the virtual conditions were triggered manually from the server application. Participants' trajectories were recorded in the server application as well.

Participants

19 pairs of participants took part in the experiment, ten pairs in the *Co-First* group and nine pairs in the *D-First* group. Age of the participants ranged from 22 to 60 years, with the median of 30 years. 14 participants were female and 24 male. All participants were

[1] https://ikinema.com/orion

naive to the purpose of the study. None of the participants withdrew from the study or reported symptoms of simulator sickness.

Procedure

Participants came to the laboratory in pairs. Participants from each pair knew each other; usually, they were friends or partners. They were greeted by the experiment coordinator and asked to read and sign a consent form. The experiment coordinator then explained the procedure to follow. The exact instructions for each of the colocation conditions were given before the start of that condition. Before the start of the *Real* condition the experiment coordinator briefly demostrated walking between the markings on the floor.

A short training phase without earphones took place before the start of the first virtual condition. First, both users from a pair were given time to freely walk around the tracking space to get familiar with locomotion and avatars. After the free exploration phase, the experimental task was introduced to the participants in a trial run. During this trial, the experiment coordinator gave the participants from a pair step-by-step instructions until both participants understood the task.

Three colocation conditions were separated by short breaks during which the experiment coordinator explained the next condition and recalibrated the tracking equipment of one of the participants if the change of the tracking area took place. After the last of three sessions the participants filled out the questionnaire. The experiment coordinator then explained the purpose of the study and answered questions. The whole procedure took around 45 mins.

6.1.3 Results

Observations

The difference in collision avoidance behavior in the *VirtualDistributed* condition between the participants from the *Co-First* and *D-First* groups was visible already during the experiment. The majority of the participants from the *Co-First* group avoided each other's virtual avatars in the *VirtualDistributed* condition. Only two pairs tried walking through each other's avatars, and only in one walking sequence each.

Participants from the *D-First* group, on the opposite, walked through each other's avatars in the most cases, often already during the first walking sequence or even during the instruction period. Only one pair from the *D-First* group did not walk through each other's avatars in the *VirtualDistributed* condition, four pairs did it one or two times, and another four pairs let their avatars collide in five to eight walking sequences out of ten.

When asked about why they decided to walk through each other's avatars, participants from the *D-First* group said that they were either curious about what would happen if they collide with the avatars or they did not want to deviate from the straight path towards the goal. Participants from the *Co-First* group said that it did not occur to them that they could walk through each other's avatars. Furthermore, they felt "respect" towards the human figure and did not want to violate its space even though they knew it would not lead to an actual collision.

Data analysis

Mixed design ANOVA (trajectory type × colocation type × order of virtual conditions) was used for the evaluation of trajectory metrics. Data resulting from five walks in each combination of collision trajectory (*Frontal* or *Crossing*) and colocation type (*Real*, *VirtualColocated* and *VirtualDistributed*) was collapsed to produce one average value for a pair of participants for each of three metrics. For the analysis of the questionnaire responses, non-parametric tests were used as the data was not normally distributed in many cases. All results are reported as statistically significant at p <.05 level.

Trajectory metrics

Figure 6.6 illustrates the trajectories walked by the participants in the *Real* and *VirtualColocated* conditions. It can be seen that the participants normally took a larger detour around each other in the *VirtualColocated* condition then in the *Real* condition. This difference is confirmed by the results of mixed ANOVA for *Clearance*. For the *VirtualDistrbuted* condition, the difference is in the proxemics behavior of the participants from the *Co-First* and *D-First* groups is also reflected in the resulting trajectories which are displayed in Figure 6.7. The differencesare especially noticeable in the *Frontal* collision courses.

There was no significant main between-subject effect of the order of virtual conditions for any of the metrics.

For *Clearance*, there was a significant main effect of the collision trajectory (F = 29.012, p < .001), colocation type (F = 44.967, p < .001) as well as a significant interaction between the colocation type and the order of virtual sequences (F = 11.441, p < 0.001).

Contrasts analysis showed that *Clearance* was larger for *Crossing* (Mean = 0.989 m) collisions than for *Frontal* (Mean = 0.775 m) ones (F = 29.012, p < .001).

Contrasts for the colocation type revealed that *Clearance* in the *VirtualColocated* condition (Mean = 1.068 m) is significantly larger than in the *Real* condition (Mean = 0.784 m) (F = 91.933, p < .001).

The difference between *VirtualDistributed* and *Real* conditions was however not statistically significant (F = .066, p = .801). This fact can be explained by the significant interaction between the colocation type and the order of virtual conditions. Contrasts for this interaction effect showed difference in the comparisons of *Clearance* in the *VirtualDistributed* and the *Real* colocation conditions between the *Co-First* and *D-First* groups (F = *6.778*, p = .019). In the *Co-First* group, *Clearance* was larger in the *VirtualDistributed* condition (Mean = 0.936 m) than in the *Real* (Mean = 0.827 m) condition. In the *D-First* group, *Clearance* was smaller in the *VirtualDistributed* condition (Mean = 0.651 m) than in the *Real* (Mean = 0.741 m) condition. The plots of *Clearance* illustrating differences

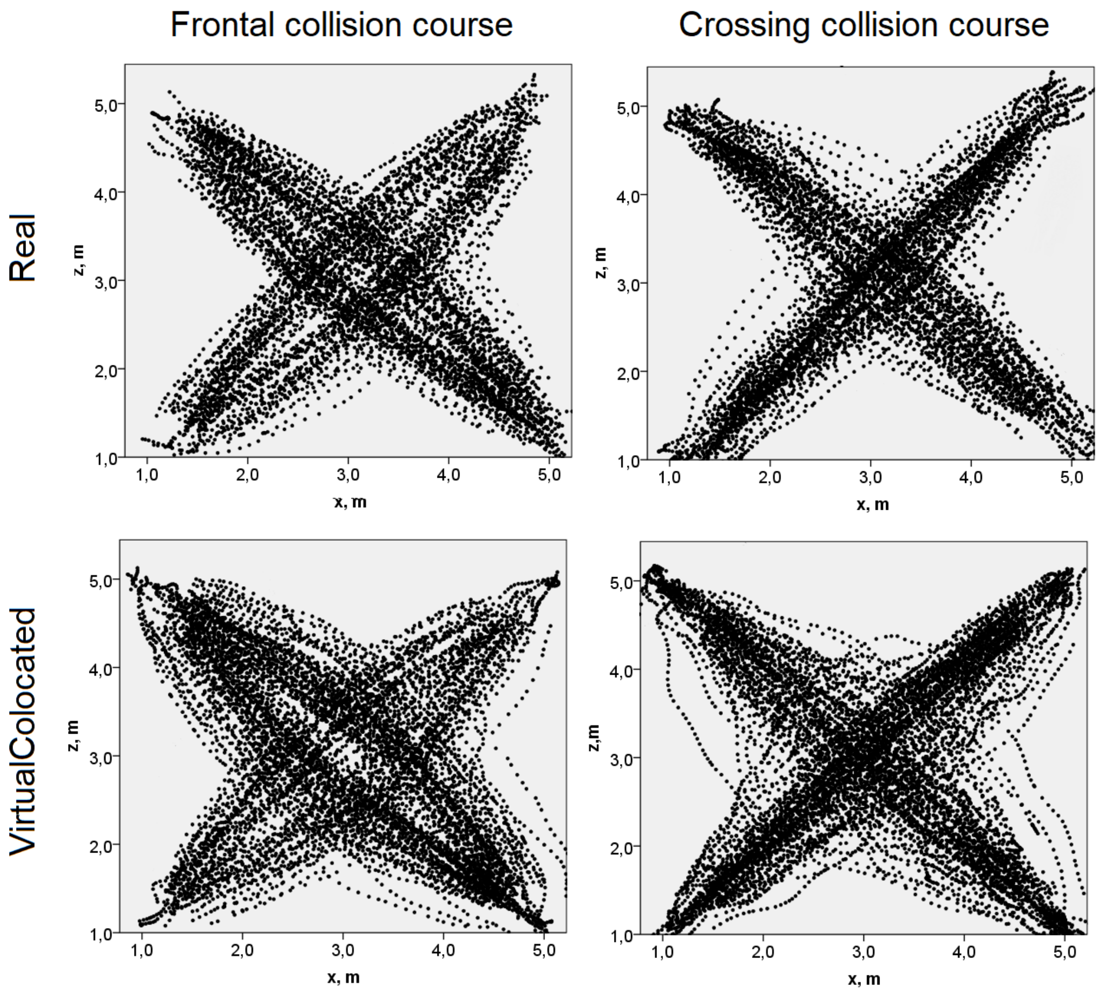

Figure 6.6: All trajectories walked by the participants in the *Real* and *VirtualColocated* conditions.

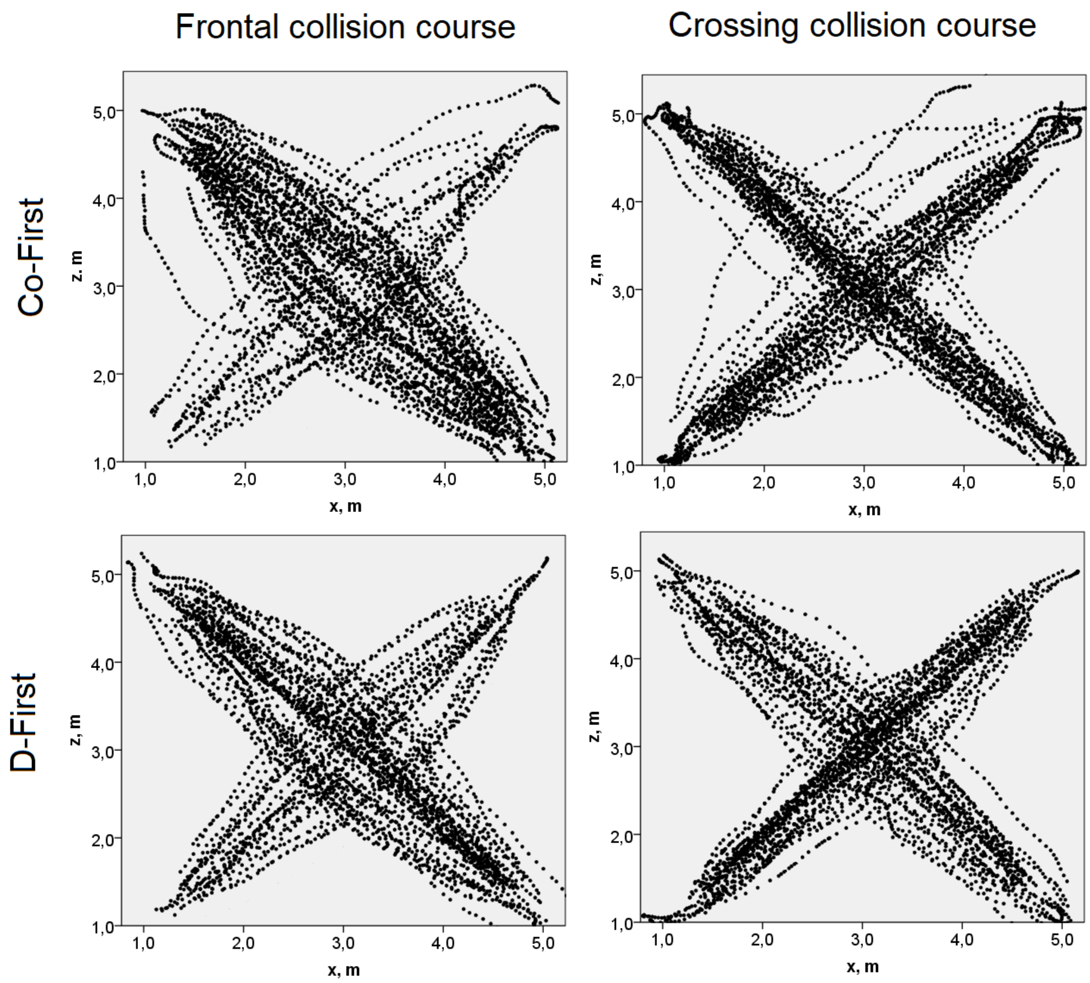

Figure 6.7: All trajectories walked by the participants from the *Co-First* and *D-First* groups in the *VirtualDistributed* condition.

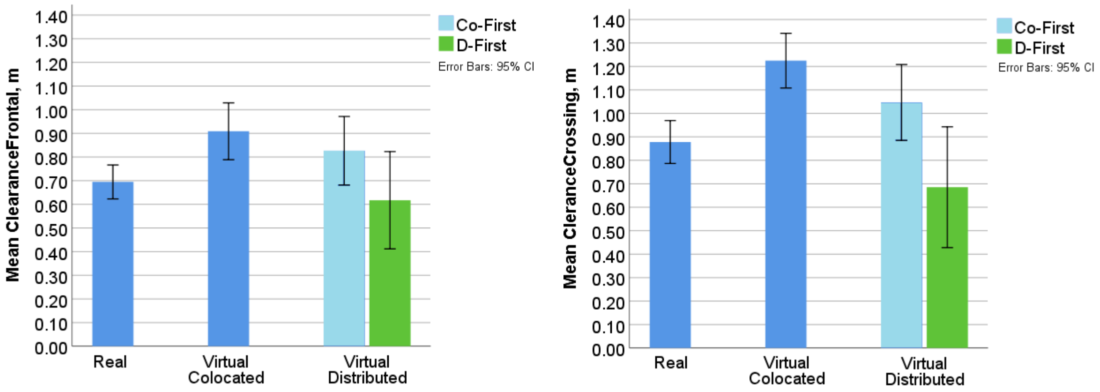

Figure 6.8: Mean *Clearance* for *Frontal* and *Crossing* trajectories in the *Real*, *Virtual-Colocated* and *VirtualDistributed* conditions.

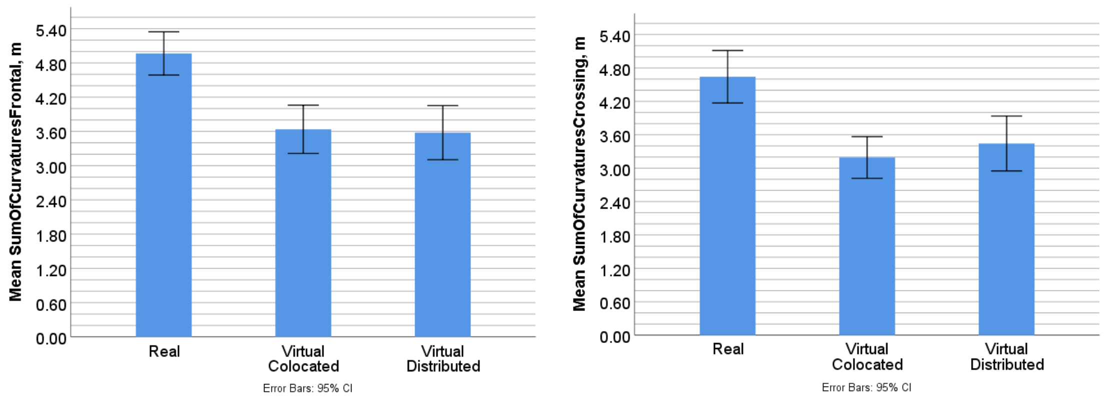

Figure 6.9: Mean *SumOfMedianCurvatures* for *Frontal* and *Crossing* trajectories in the *Real*, *VirtualColocated* and *VrtualDistributed* conditions.

between the *Real*, *VirtualColocated* and *VirtualDistributed* (separately for *Co-First* and *D-First*) conditions are presented in Figure 6.8.

For *SumOfMedianCurvatures*, there was a significant main effect of the collision trajectory ($F = 10.605$, $p = .005$) and of the colocation type ($F = 54.714$, $p < .001$). There were no significant interactions. Contrasts revealed that *SumOfMedianCurvatures* was larger for *Frontal* collision trajectories (Mean = 4.072 m) than for *Crossing* collision trajectories (Mean = 3.770 m) ($F = 10.605$, $p = .005$). Furthermore, *SumOfMedianCurvatures* was smaller in the *VirtualColocated* condition (Mean = 3.419 m) than in the *Real* condition (Mean = 4.821 m) ($F = 37.258$, $p < .001$), and smaller in the *VirtualDistributed* (3.523 m) than in the *Real* condition ($F = 31.947$, $p < .001$). *SumOfMedianCurvatures* in the *Real*, *VirtualColocated* and *VirtualDistributed* conditions is illustrated in the plots in Figure 6.9.

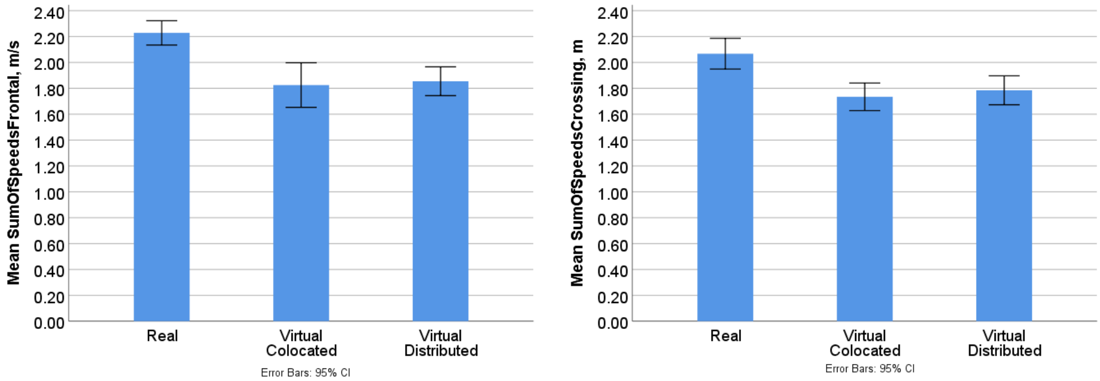

Figure 6.10: Mean *SumOfSpeeds* for *Frontal* and *Crossing* trajectories in the *Real*, *VirtualColocated* and *VrtualDistributed* conditions.

Table 6.2: Deviations of *Clearance* from the *Real* condition.

Clearance		
	Frontal	*Crossing*
VirtualColocated	Δ = **+0.22** m	Δ = **+0.35** m
VirtualDistributed Co-First	Δ = **+0.10** m	Δ = **+0.10** m
VirtualDistributed D-First	Δ =-0.05 m	Δ = -0.13 m

Table 6.3: Deviations of *SumOfSpeeds* from the *Real* condition.

SumOfSpeeds		
	Frontal	*Crossing*
VirtualColocated	Δ = **-0.4** m/s	Δ = **-0.34** m/s
VirtualDistributed	Δ = **-0.38** m/s	Δ = **-0.29** m/s

For *SumOfSpeeds*, there was a significant main effect of the collision trajectory (F = 14.041, p = .002) and of the colocation type (F = 37.935, p < .001). There were no significant interactions. Contrasts showed that *SumOfSpeeds* was larger for *Frontal* collision trajectories (Mean = 1.97 m/sec.) than for *Crossing* collision trajectories (Mean = 1.86 m/s) (F = 14.041, p = .002). *SumOfSpeeds* was smaller in the *VirtualColocated* condition (Mean = 1.78 m/s) than in the *Real* condition (Mean = 2.15 m/s) (F = 49.852, p < .001), and smaller in the *VirtualDistributed* (Mean = 1.82 m/s) than in the *Real* condition (F = 37.321, p < .001). *SumOfSpeeds* in the *Real*, *VirtualColocated* and *VirtualDistributed* conditions is illustrated in the plots in Figure 6.10.

Tables 6.2, 6.3 and 6.4 summarize deviations from the *Real* condition for each metric. Values that reflect statistically significant deviations from the metric of the *Real* condition are shown in bold.

Table 6.4: Deviations of *SumOfMedianCurvatures* from the *Real* condition.

SumOfMedianCurvatures		
	Frontal	*Crossing*
VirtualColocated	Δ = **-1.33** m	Δ = **-1.47** m
VirtualDistributed	Δ = **-1.39** m	Δ = **-1.21** m

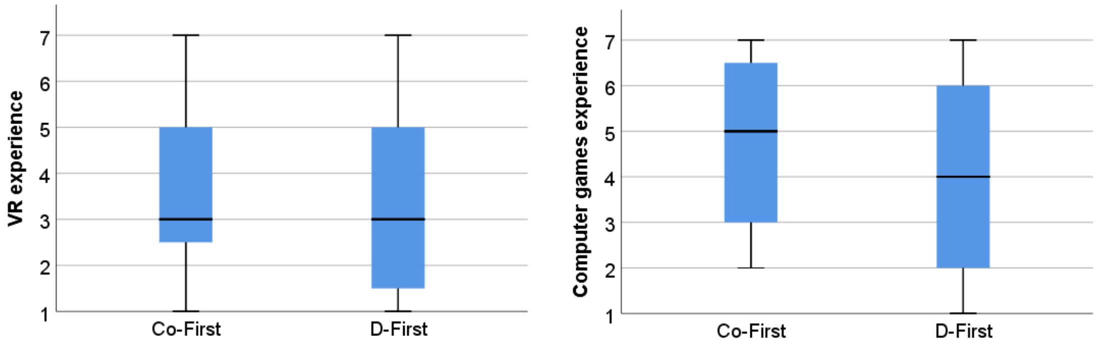

Figure 6.11: Resulting scores of the self-report of previous VR experience and the experience of playing computer games in the *Co-First* and *D-First* group.

Questionnaire data

VR experience and presence Figure 6.11 illustrates the resulting scores for previous VR experience and the experience of playing computer games given by the participants from the *Co-First* and *D-First* groups. The distributions of the scores were not different between the groups.

However, there was a statistically significant difference in the reported familiarity with VR and computer games between male and female participants. The corresponding boxplots can be seen in Figure 6.12 and the details of the Mann-Whitney U tests are summarized in Table 6.5.

The resulting scores for questions *Pr1* - *Pr4* were high for both groups. The corresponding boxplots are presented in Figure 6.13. There were no statistically significant differences between the *Co-First* and *D-First* groups in the Mann-Whitney U test. There were no between-gender differences in the self-report of presence.

Copresence The resulting scores of perceived copresence (question *CoPr1*) were higher in the *VirtualColocated* condition than in the *VirtualDistributed* condition. The difference was statistically significant in the Wilcoxon signed-rank test (T = 12.0, z = -3.186, $p <$ .001, r = -.52). There was no statistically significant effect of the group for *CoPr1* in neither *VirtualColocated* no *VirtualDistribted* condition.

However, female participants reported higher copresence scores in the *VirtualDistributed*

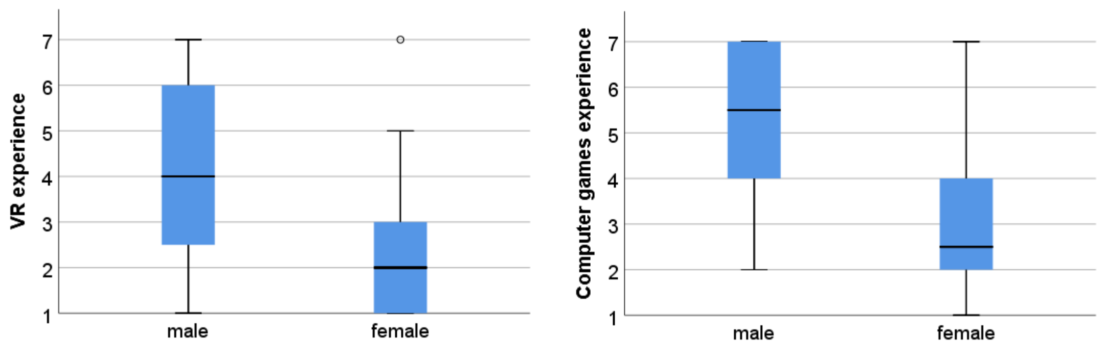

Figure 6.12: Boxplots of the resulting scores of the self-report of previous VR experience and the experience of playing computer games among male and female participants.

Table 6.5: Results of the Mann-Whitney U test on the self-report of VR and computer games experience. Test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

Co-First vs *D-First*				
	U	z	p	r
VR experience	147.0	-.991	.339	-.16
Computer games experience	136.5	-1.308	.201	-.21
male vs female				
	U	z	p	r
VR experience	92.5	-2.315	**.021**	-.38
Computer games experience	60.5	-3.311	**.001**	-.54

condition than male participants. The difference is borderline not significant (in the Mann-Whitney U test, U = 232.0 , z = 1.977, p = .54, r = .32). Boxplots reflecting reflecting the differences in the scores of the question *CoPr1* can be seen in Figure 6.14.

Participants reported to have been more worried about colliding in the *VirtualColocated* condition than in the *VirtualDistributed* condition (in the Wilcoxon signed rank test, T = 22.0, z = -4.150, $p < .001$, r = -.66). The scores were not different between the *Co-First* and *D-First* groups, according to the Mann-Whitney U test. Boxplots illustrating the resulting scores of the question *CoPr2* are shown in Figure 6.15.

Tables 6.6 and 6.7 summarize the answers about participants' main focus during the test in the *VirtualColocated* and *VirtualDistributed* conditions. It can be seen that the majority of participants from both groups mostly focused on avoiding collisions with their test partners during the *VirtualColocated* test. During the *VirtualDistributed* test, on the contrary, most participants from both groups focused on the experimental task (collecting items), although the answers of the participants from the *Co-First* group were somewhat more spread out.

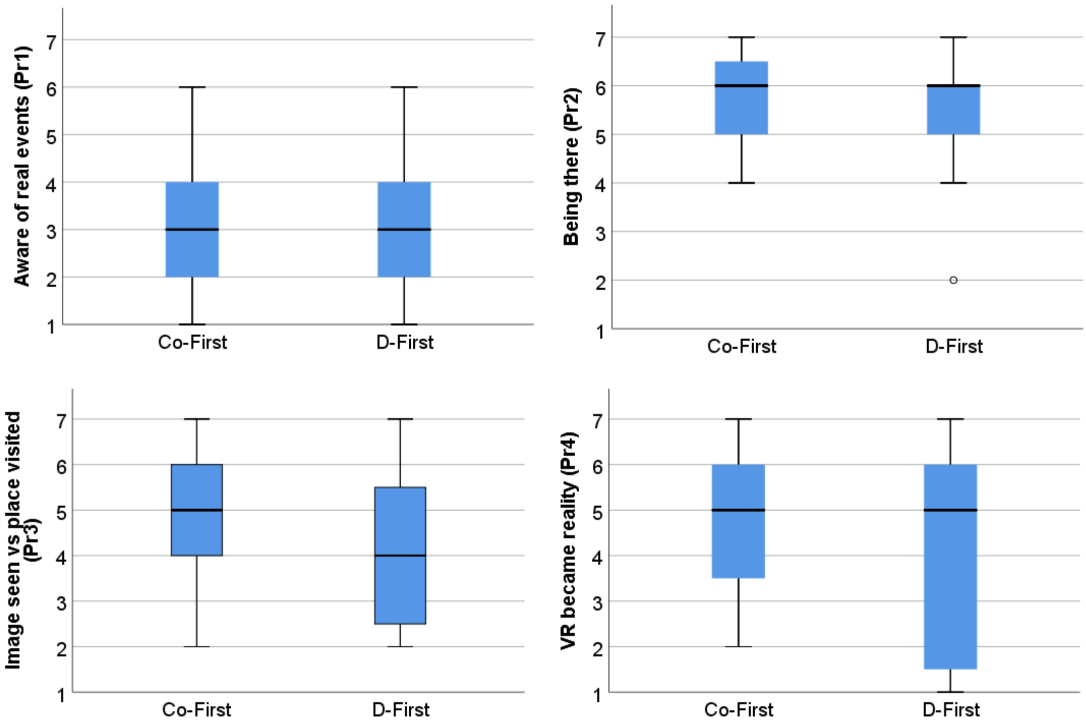

Figure 6.13: Scores of presence-related questions resulting from the answers of the participants from the *Co-First* and *D-First* groups.

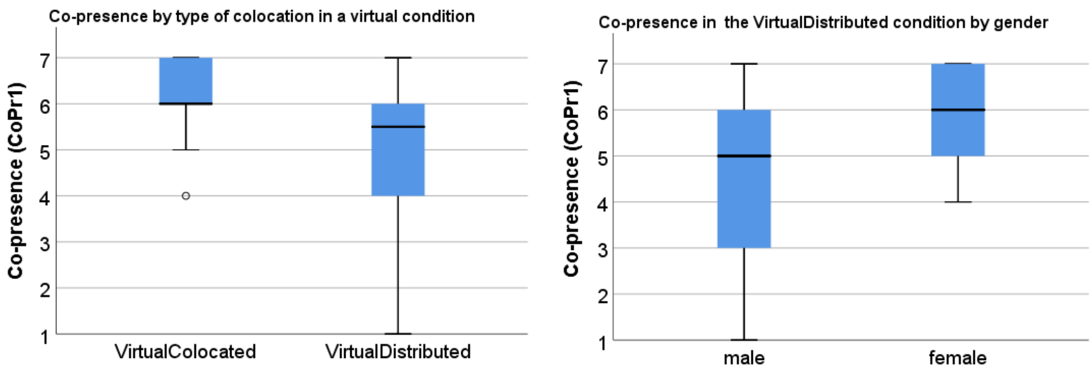

Figure 6.14: Scores of the question *CoPr1* in the *VirtualColocated* and *VirtualDistributed* conditions.

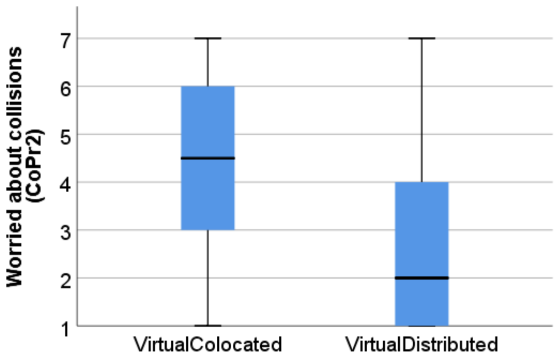

Figure 6.15: Scores of the question *CoPr2* in the *VirtualColocated* and *VirtualDistributed* conditions.

Table 6.6: Self-reported of the main focus of attention during the *VirtualColocated* test. Each cell contains the number of participants from the respective group who chose the answer in bold.

Group	**Task**	**Avoidance**	**Both**	**Cannot tell**
Co-First	2	12	3	2
D-First	3	14	0	2

Table 6.7: Self-reported of the main focus of attention during the *VirtualDistributed* test. Each cell contains the number of participants from the respective group who chose the answer in bold.

Group	**Task**	**Avoidance**	**Both**	**Cannot tell**
Co-First	10	4	2	3
D-First	14	4	0	1

6.1.4 Discussion

Clearance distances were larger for *Crossing* collision trajectories compared to *Frontal* trajectories for all colocation conditions. This result is consistent with the observations of the elliptic shape of personal space [57]. Slightly slower speeds for *Crossing* collisions in all three colocation conditions are likely to be due to the fact that locomotor adaptations in case of crossing trajectories involve speed as well as heading adjustment. It is however hard to explain that trajectories of the *Crossing* collision type were slightly more curved than those of the *Frontal* type. Overall, the absence of interactions between the collision type and any of two other independent variables indicates similarity of general mutual collision avoidance principles in real and virtual environments.

In previous experiments on collision avoidance, subjects did not know in advance where the obstacle, whether moving or static, would appear [44, 108, 107]. Such design was used in order to make participants achieve their natural walking speeds before they had to adjust their locomotory trajectory to avoid an obstacle. To ensure high-quality motion capture we had to limit our experiment to a smaller tracking area not allowing for additional walking space. Our participants knew the location of their test partners at all times. This fact might have made collision avoidance easier, however equally in real and virtual conditions.

As expected, participants displayed more cautious behavior in the *VirtualColocated* condition compared to the *Real* condition. They kept larger distance to each other (larger *Clearance*), walked slower (smaller *SumOfSpeeds*) and their paths were less straight (smaller *SumOfMedianCurvatures*) than in the *Real* condition. Slower walking speed and larger clearance in VR were demonstrated in previous experiments with static obstacles [44, 3]. Speed differences in our experiment ranged from $\Delta = 0.15$ m/s to $\Delta = 0.2$ m/s if to assume equal speeds for two participants from a pair, comparable to $\Delta = 0.13$ m/s found by Fink et al. [44].

The influence of the order of virtual conditions on the collision avoidance behaviour in *VirtualDistributed* colocation scenario is an interesting result. The *VirtualDistributed* test was similar to previous experiments on locomotion in the presence of obstacles in VR in a sense that the obstacles (test partners in this case) in this condition were purely virtual. In contrast to the previous studies, we did not instruct participants to avoid obstacles. Our results indicate that the choice of whether to avoid them or not was influenced by the order of the virtual conditions. The *VirtualColocated* test session clearly had a priming effect on participants' avoidance behaviour during the subsequent *VirtualDistributed* session. Priming has been previously observed and studied in VR research. For example, in a recent paper on social presence with virtual agents [37] priming exposure to a social situation including a virtual agent was used in an attempt to induce higher sense of social presence towards the social agent in the subsequent interaction. Social priming was reported to result in a significant increase of social presence assessed with self-report measures. Although unintended, the priming effect of our *VirtualColocated* condition was strong enough to influence not only self-report but also the actual proxemics behavior of participants in the *VirtualDistributed* condition. This result demonstrates once again that a careful use of priming could help achieving the desired experience for users of a VR scenario.

It is rather surprising that users did not try avoiding collisions in many cases during the *VirtualDistributed* test when it was the first virtual experience. Such behavior could be interpreted by the fact that the elicitation of presence or copresence did not happen for these participants. However, self-report of presence was equally high among the participants from the *Co-First* and *D-First* groups. Participants reported slightly lower copresence in *VirtualDistributed* condition than in the *VirtualColocated* condition, however copresence scores in the *VirtualDistributed* condition were not different between the *Co-First* and *D-First* groups. The discrepancy between the behavioral characteristics and the self-report might reflect the weakness of the particular used questionnaire or of the self-report-based assessment approach in general.

Interestingly, female participants reported higher copresence in the *VirtualDistributed* condition than male participants. The only other difference in the self-report between female and male participants was the experience of VR and playing computer games. It has been previously shown that presence and copresence in VR may be influenced by the reported computer usage [53]. However, we did not have sufficient amount of data to determine whether higher copresence in the *VirtualDistributed* condition reported by female participants was connected to their lower exposure to VR and computer games.

Overall, when participants chose to avoid colliding with each other in VR they did it with more caution than in the real space, in both *VirtualColocated* and *VirtualDistributed* scenarios. Although participants were less worried about colliding with their test partners and generally remembered that they were safe from actual physical collisions during the *VirtualDistributed* test their walking speeds and the curvatures of their locomotor trajectories were similar to those in the *VirtualColocated* condition. The clearance distances were however smaller in the scenario with separate tracking spaces, even when users avoided colliding with each other (the *Co-First* group). Moreover, the majority of participants focused on collision avoidance and not on the task itself when walking in the shared tracking space.

6.2 Copresence, social presence and proxemics in a mixed setup

6.2.1 Motivation

In shared VR with mixed physical colocation a physical tracking space may be used by one person or may be shared by a pair or a small group of users. The virtual world is however shared by all connected users. In this case, each user within a shared tracking space has co-players who are colocated both physically and virtually (like in a colocated shared VR setup) and co-players who are colocated only virtually while physically being in a different geographical location (like in a distributed shared VR setup). For those users who have a separate tracking space of their own the setup is essentially the same as in distributed shared VR. However, the asymmetric setup may influence the experience of all connected users. The experiment presented in this section was designed to investigate such possible influence.

The focus of the described experiment was on copresence, both self-reported and interpreted from the behavioral metrics. Shared VR scenarios may elicit the sense of copresence and social presence. Whereas not all shared VR scenarios require the elicitation of strong sense of social presence the copresence is always the goal of creating a shared environment. For example, in rescue scenarios each training user may need to exercise the ability to react quickly and move precisely in the presence of others according to their specific role but without the emphasis on social interaction.

The first goal of our experiment was, in a setup where three players are immersed within two separate tracking spaces, to investigate whether there is a difference in how users perceive their colocated and distributed co-players. Secondly, we wanted to see if there is a difference in how users perceive their co-players depending of whether they are alone in a tracking space themselves or share it with others.

Prompted by the results of the previously described experiment in which we observed participants walk through each other's avatars we also intended to see whether users display different proxemics behavior depending on whether they are alone in a tracking space or share it with a co-player.

Finally, we wanted to examine whether the nature of the task accomplished in the VE could influence proxemics behavior and self-report of copresence and social presence.

6.2.2 Experiment design

Three participants at a time took part in the experiment. Two of them shared a tracking area whereas the third one was in a separate tracking area. We will refer to participants who shared the tracking area as "co-placed" and to the ones who were alone in the tracking area as "separated". From a user perspective, each co-placed participant had a colocated and a distributed co-player, whereas each separated participant had two distributed co-players. This way, "colocated" and "distributed" will be used to refer to participants from the point of their test partners.

A simple game where participants had to build a tower from building blocks scattered in the VE was chosen as the main experimental task. This task prompted participants to move across the whole tracking area while avoiding colliding with their test partners. Two versions of the task were tested: a competitive scenario where each of three participants had to build a separate tower and a collaborative scenario where all three participants built one tower.

Prior to the main experiment, we conducted a pilot test with two teams of participants. The participants from the pilot teams were not restricted in the way they could interact with each other. In both teams, separated participants walked through avatars of their co-players, in both competitive and collaborative scenarios. Some of the co-placed participants did not like this behavior in the beginning of the test but soon learned that it was "the rule" of the virtual world and started walking through the avatars of distributed co-players as well. In the de-briefing sessions after the test, both separated and co-placed participants reported very low degree of copresence towards their distributed co-players. Following these results, we decided to introduce limitations on the allowed interactions. Whereas such limitations may not be natural for a real-life scenario, they allowed us to have a deeper insight into the illusion of copresence and social presence in shared VR. All teams were split into two groups. In the first group, only separated participants received instructions limiting their proxemics behavior. In the second group, all three participants from a team received such instructions.

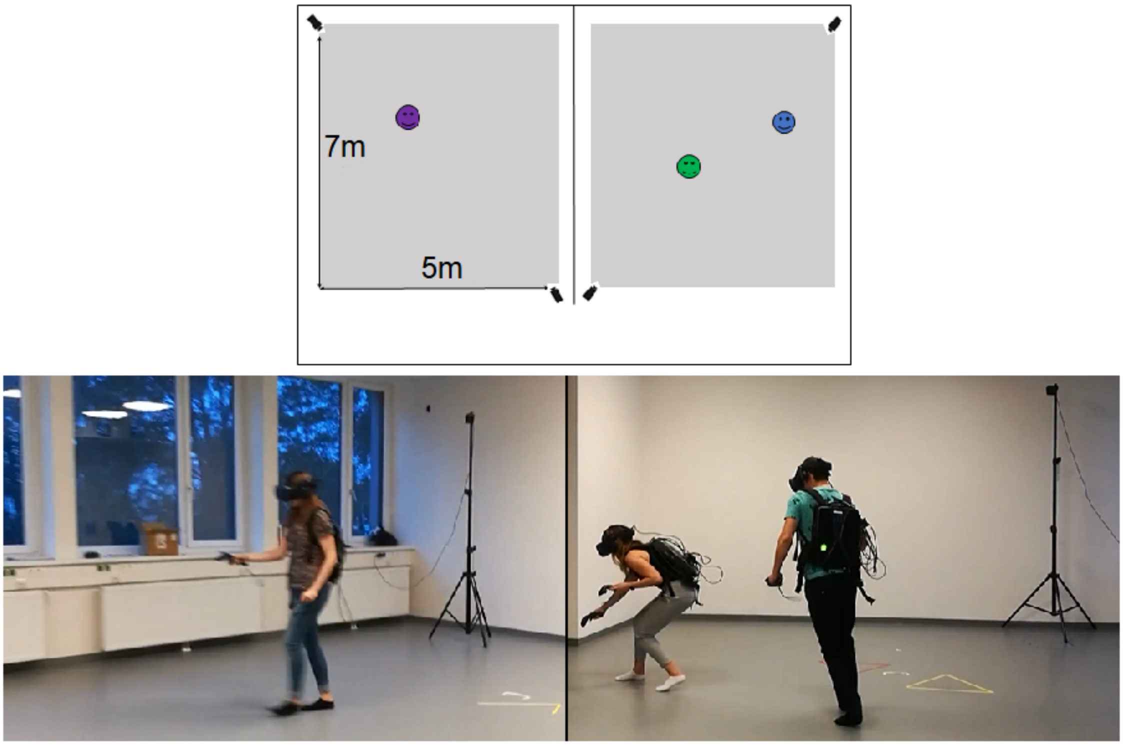

Figure 6.16: The schematic placement of participants within two tracking areas (top) and a separated (bottom left) and two co-placed (bottom right) participants during the experiment.

Following the described goals, our experiment had a 2 (placement in a tracking space) x 2 (type of instructions) x 2 (type of scenario) mixed design. The type of placement and the type of instructions were between-subject factors and the type of scenario (competitive vs collaborative) was a within-subject factor. The gender of a participant was added as another between-subject factor during the data analysis.

Environment

The experiment took place in the laboratory with two tracking areas separated by a wall segment and a thick curtain, each 7 × 5 m large. Figure 6.16 illustrates the placement of participants in the tracking areas.

The VE used in both scenarios consisted of a large terrain with a sandy texture and a skybox. The walking area was marked with red border. Each participant had a designated starting position in the VE that depended on the color of their avatar's suit.

Competitive scenario

In the competitive scenario, each participant had to build a tower out of ten cubes scattered in the playground in a specified location within the game area. Each player had their own set of cubes for construction that color-matched the suit of the player's avatar. The cubes were positioned in such a way that each participant had to walk across the playground to collect a cube and bring it to the building location. To collect a cube, a player had to touch it with one of their Vive controllers and press the trigger button on the controller. The cube was held while the trigger button was pressed and could be released when the player released the trigger button. The virtual models of the Vive controllers held by a player were only shown in their local application and not replicated to the server and other players.

Participants were told that the player who builds their tower first would win the game. When this happened, a particle system effect began playing around the winner so that all three participants could see who won the game. It was possible to pick up only one cube a time. This way, participants were forced to walk across the tracking area 20 times to finish the construction of their towers. Participants could not interact with cubes of their co-players in any way thus not being able to cheat by destroying other towers. The competitive scenario is illustrated in Figure 6.17 where all three players can be seen in the server view not long after the start of the game.

Collaborative scenario

In the collaborative scenario, all three players from a team constructed one big tower in the middle of the playground. Just as in the competitive scenario, each player could only pick up cubes of their matching color. The amount and the initial positions of the cubes were the same as in the competitive scenario. The task was completed when all the cubes were put into the central tower.

To ensure that trajectories walked by the participants were somewhat similar to those in the competitive task, the players were made to "charge" their cubes by shortly holding them inside colored spheres positioned on the poles at the locations of the towers from the competitive task. The server view of the collaborative scenario is shown in Figure 6.17.

Participants were asked not to talk to each other during both experimental scenarios to create similar audio conditions for separated and co-placed participants. Additionally, rain-like noise recording was playing in the HMDs' headphones during both scenarios.

CollisionIgnorant group

In this group of teams, only separated participants received instructions restricting their proxemics behavior. They were secretly instructed to not walk through the avatars of their test partners. If the co-placed participants from a team, in their turn, tried to walk

Figure 6.17: A team of participants in the competitive (left) and collaborative (right) scenarios.

through the avatar of the separated participant the latter was instructed not to try to avoid a virtual collision.

In the *CollisionIgnorant* group, each of three participants from a team was assigned a fixed color, blue, green or purple. The suit of their avatar had this color in both games and all participants know each other's colors in advance.

CollisionInformed group

In this group, all three participants from a team were instructed to not touch or walk through each other's avatars. Participants did not know in advance which color they would have in the game. The colors were assigned differently in each scenario. With these restrictions, we intended to see if co-placed participants could easily identify their distributed and colocated co-players.

Avatars

Participants had gender-matching human avatars created with Autodesk Character Generator[2]. All avatars had identical suits only differing in colors that can be seen in Figure 6.18. The avatars were animated with the use of inverse kinematics. The avatar's head followed the position and rotation of the HMD and its hands followed the positions and rotations of the HTC Vive controllers. The positions of the avatar's elbows and pelvis followed the rules of inverse kinematics. The range of movements of an avatar was however limited due to the low amount of tracking points. The avatars had a locomotion component that allowed them to make steps when the user's HMD was moving horizontally. These steps did not necessarily coincide with the participant's own steps however looking like a walk to the co-players.

[2]https://charactergenerator.autodesk.com/

Figure 6.18: *Blue*, *Purple* and *Green* avatars in female and male version used in the experiment.

Measures

Trajectory metrics The movement trajectory of each participant was recorded in both scenarios. We then calculated the average walking speed (*MeanSpeed*) of each participant and the average clearance between each two participants from a team (*MeanClearance*). The clearance was obtained by averaging several minimal distances between each pair of participants when they passed each other during the game. The closest distance between each pair of participants from a team achieved during a game was analyzed as well (*MinClearance*).

Questionnaires The pre-test questionnaire included demographics as well as questions about the previous experience of VR and playing computer games.

The main objective of the post-questionnaire was to assess the self-report of copresence and social presence experienced by participants. The illusions of copresence and social presence are however complex phenomena the elicitation of which is conditioned by the experience of place presence and the sense of embodiment, according to the theory of

Table 6.8: Questionnaire including the following dimensions: Presence (*Pr*), Enjoyment (*Enj*), Body Ownership (*BdOwn*) and Agency (*Ag*), CoPresence (*CoPr*), Mutual Awareness (*Aw*), Attentional Allocation (*Atn*) and Behavioral Engagement (*BhvEng*).

QUESTION CODE	QUESTION
Pr1	How aware of real events occurring in the room around you were you during the game?
Pr2	How involved were you in playing the game?
Pr3	Rate your sense of being there in the virtual playground
Pr4	When you think about the virtual playground, was it rather like something that you saw or somewhere that you visited?
Enj	How much did you enjoy playing the game?
BdOwn1	I felt as if the virtual body was my own body
BdOwn2	It felt as if the virtual body I saw at the place of my own body was someone else
BdOwn3	It seemed as if I might have more than one body
Ag1	It felt like I could control the virtual body like if it was my own
Ag2	The movements of the virtual body were caused by my own movements
Ag3	I felt as if the movements of the virtual body were influencing my own movements
Ag4	I felt as if the virtual body was moving by itself
CoPr1	To what extent did you have the feeling of the COLOR player being together with you in the virtual playground?
CoPr2	Rate how closely the sense of being together with COLOR resembles the sense of being with others in the real world
CoPr3	To what extent were you worried that you would collide with the COLOR player?
Aw1	The COLOR player caught my attention
Aw2	I hardly noticed the COLOR player
Aw3	The COLOR player did not notice me in the virtual playground
Atn1	I paid close attention to the COLOR player
Atn2	The COLOR player paid close attention to me
Atn3	I tended to ignore the COLOR player
Atn4	The COLOR player tended to ignore me
BhvEng1	My actions depended on the COLORs actions
BhvEng2	The COLORs actions depended on my actions

Embodied Social Presence [93]. Therefore, questions intended to assess presence and embodiment were also included in the post-questionnaire.

The questions from the post-test questionnaire are summarized in Table 6.8. The assessment of the experienced presence was done with the same questions that we used in the previous experiments, adapted from the questionnaires of Witmer and Singer [158] and Slater, Usoh and Steed [130]. A question about the enjoyment of the game was added. To assess embodiment, we used questions related to body ownership and agency adapted from Gonzalez et al. [59].

According to Biocca et al. [22], social presence is achieved by evoking copresence, psychological involvement and behavioral engagement. These three phenomena are in their turn affected by mutual awareness (for copresence), mutual attention, mutual understanding and empathy (for psychological involvement) and behavioral interaction, mutual assistance and dependent action (for behavioral engagement). We used questions suggested by Biocca et al. [22, 21] to address these groups of phenomena as well as questions adapted from an experiment with three users of Slater et al. [126].

All answers to the questions were given on a 7-point Likert scale. For questions formulated as statements, 1 meant strongly disagree and 7 - strongly agree. The questions concerning co-players were asked twice, each time with appropriate color (instead of COLOR in Table 6.8) for a co-player.

Apparatus

As in our previous experiment on shared VR, the experimental application was developed within the ImmersiveDeck platform used with an HTC Vive HMD (Pro version) and Lighthouse tracking technology. Unity3D version 2017.0.3 was used.

Each participant was equipped with an HTC Vive Pro, two Vive controllers held in each hand a VR backpack. The rainy noise used to block real sounds was played in the in-built headphones of the Vive Pro. The backpacks were XMG Walker, each with a NVIDIA GeForce GTX 1070 graphics card and Intel Core i7 quad-core processor. Final IK plugin for Unity3D implementing an inverse kinematics solution was used to animate participants' avatars. The recording of participants' trajectories was done on the server machine.

Participants

17 teams of participants took part in the main experiment (51 participants in total, 18 female, 33 male), ten of them in the *CollisionIgnorant* group and seven in the *CollisionInformed* group. In the *CollisionIgnorant* group, two teams were female, five male, two teams with two female and one male participant and one with one female and two male participants. In the *CollisionInformed* group, two teams were female, four male and one team with one female and two male participants. The age range of participants was 20 to 49 years, with the median of 27 years.

Procedure

Participants arrived at the lab in teams. All three participants from a team knew each other. The participants from a team were given a general overview of the experiment, gave their consent to participate and filled out the demographics questionnaire. The experiment coordinator then explained the scenarios and tasks and showed how to use the controllers to pick up cubes. A short training session took place before each scenario. During this session, participants had time to accommodate to wearing an HMD and walking within the borders of the playground in the virtual space. They also trained to pick up and carry the cubes and to fulfill the scenario task.

In the *CollisionIgnorant* group, the secret instruction for the separated participant was given written on a sheet of paper when the experiment coordinator was helping them with the hardware and two co-placed participants already had their HMDs on in the different tracking space.

Participants took down the hardware and filled out the post-test questionnaire after each scenario. The experiment was conducted in the following order: training for the competitive game, the competitive game, break and questionnaire, training for the collaborative game, the collaborative game, final questionnaire and de-briefing. The time spent by each team in the lab was about 40 mins whereas the duration of each game was under 10 mins, with the competitive game being quicker.

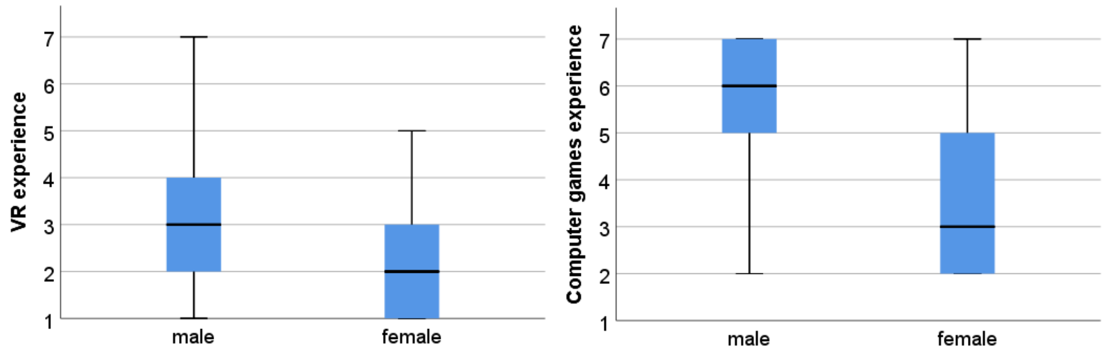

Figure 6.19: Reported previous experience of VR and playing computer games.

6.2.3 Results

Data analysis

Mixed design ANOVA (scenario × placement × group × gender) was used for the evaluation of trajectory metrics.

For the analysis of the questionnaire responses, non-parametric tests were used as the data was not normally distributed in many cases. All results are reported as statistically significant at $p < .05$ level.

VR experience

Previous experience of VR and playing computer games was not different between the participants from the *CollisionIgnorant* and *CollisionInformed* groups, or between the separated and co-placed participants. However, female participants reported to have less experience of both VR and playing computer games than male participants. The difference in the resulting scores was statistically significant. Boxplots illustrating the responses of female and male participants can be seen in Figure 6.19. The results of statistical comparisons are summarized in Table 6.9.

Observations and comments

In both groups, female participants were generally careful with movements. Female participants never walked through avatars, many of them reported that they were scared to collide with their distributed co-players even though there was no danger of a real collision. Many of the female participants also walked around obstacles in the playground trying not to step on other players' cubes or the central building spot in the collaborative scenario. This clearly more careful proxemics behavior displayed by female participants prompted us to introduce gender as another between-group factor during the analysis of the results.

Table 6.9: Results of the Mann-Whitney U test on the self-report of VR and computer games experience. Test statistic U, standardized test statistic z, corresponding p-value and effect size r are shown.

CollisionIgnorant vs *CollisionInformed*				
	U	z	p	r
VR experience	331.5	.324	.746	.05
Computer games experience	330.0	.292	.770	.04
separated vs co-placed				
	U	z	p	r
VR experience	377.0	1.803	.071	.25
Computer games experience	308.0	.387	.699	.05
male vs female				
	U	z	p	r
VR experience	182.5	-4.415	**.015**	-.34
Computer games experience	81.5	-3.311	**<.001**	-.62

In the *CollisionIgnorant* group, co-placed participants from two male teams did not walk through the avatar of the distributed co-player. In all other male or mixed teams from the group, co-placed male participants walked through the avatars of the distributed co-players at least once during the test. They explained this behavior of walking through avatars by the fact that it would not lead to real collisions.

In the *CollisionInformed* group, all co-placed participants could identify the colors of their colocated co-players after a period of time, mostly due to fast motions in the near proximity of each other. In two teams, the co-placed participants said to have forgotten that the separated participant was in the separate tracking space. The co-placed participants often commented that they had to focus on thinking about which co-player was in the same tracking space to be sure to avoid collisions.

Finally, many participants said to have liked the collaborative scenario more because they did not need to hurry and could work together. Participants expressed regret that they could not talk to each other during the collaborative game. However some participants rather preferred competing.

Trajectory metrics

Figure 6.22 demonstrates an example of the recorded trajectories of a team of participants in the competitive and in the collaborative scenario. It illustrates that participants walked on mostly straight trajectories between their building spots and cubes in the playground in the competitive scenario. In the collaborative scenario, these trajectories curved around the central building spot.

For *MeanSpeed*, the only within-subject factor in the mixed ANOVA analysis was the type of scenario (competitive and collaborative). The between-subject factors were the

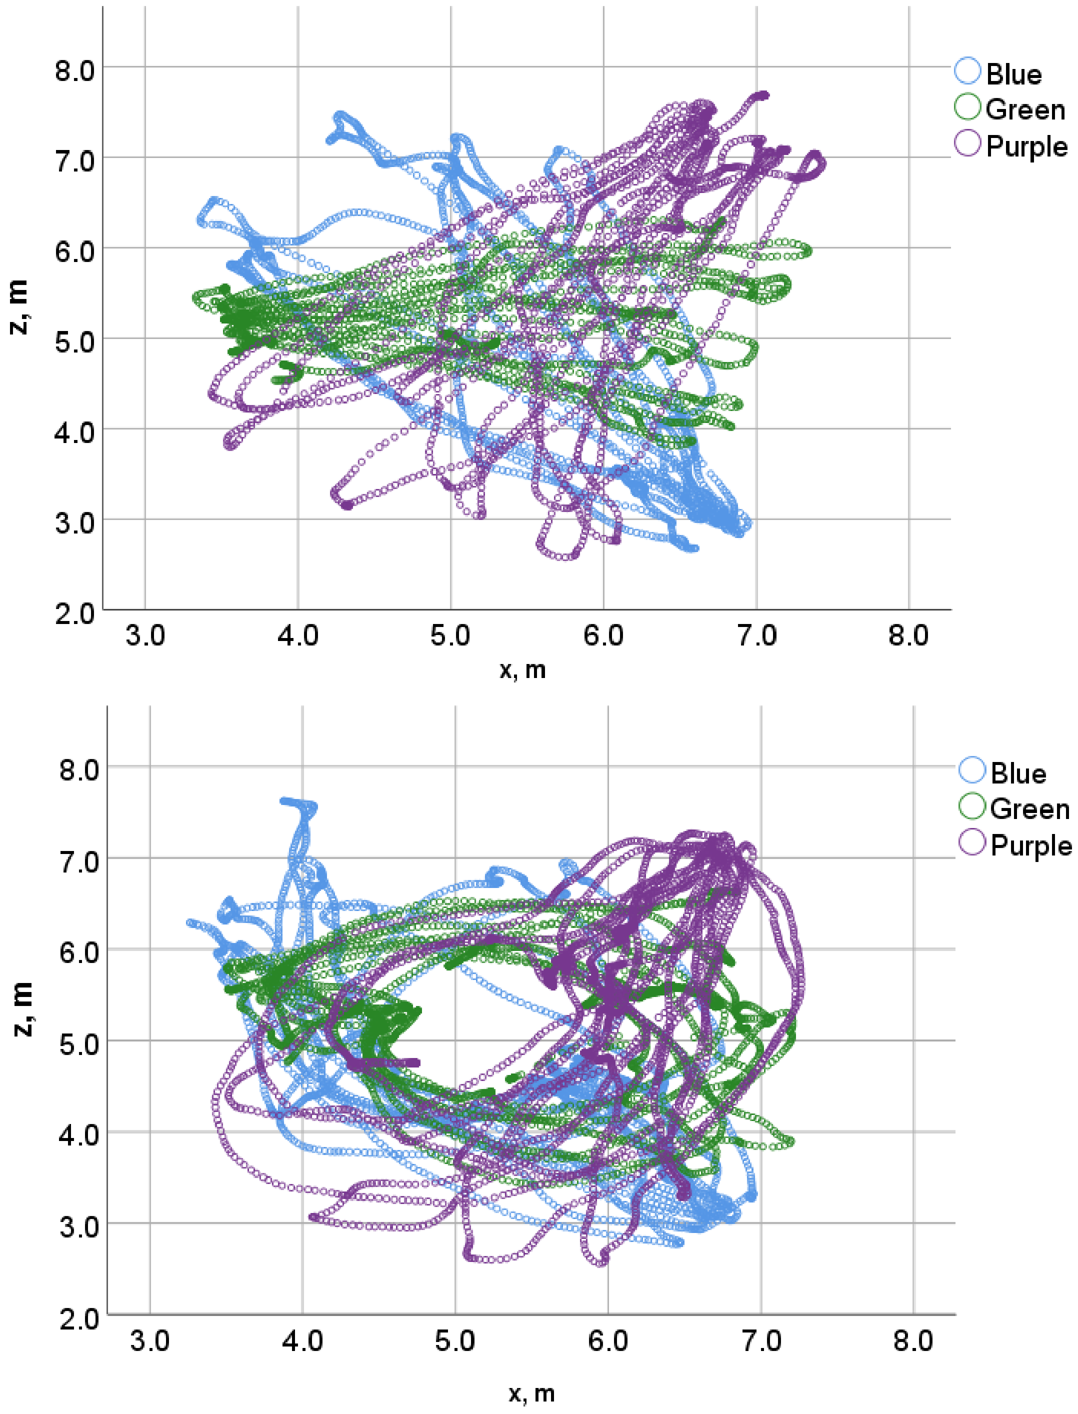

Figure 6.20: Trajectories of one team of participants in the competitive (top) and collaborative (bottom) scenarios. The coordinates are given in the horizontal plane of the Unity3D coordinate system.

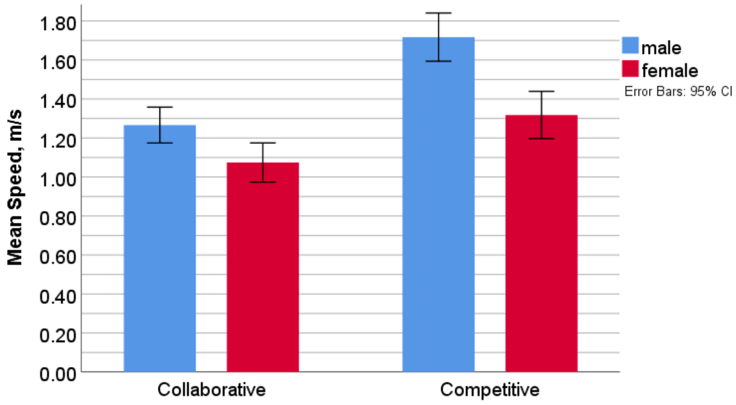

Figure 6.21: *MeanSpeed* of female and male participants in the competitive and collaborative scenarios.

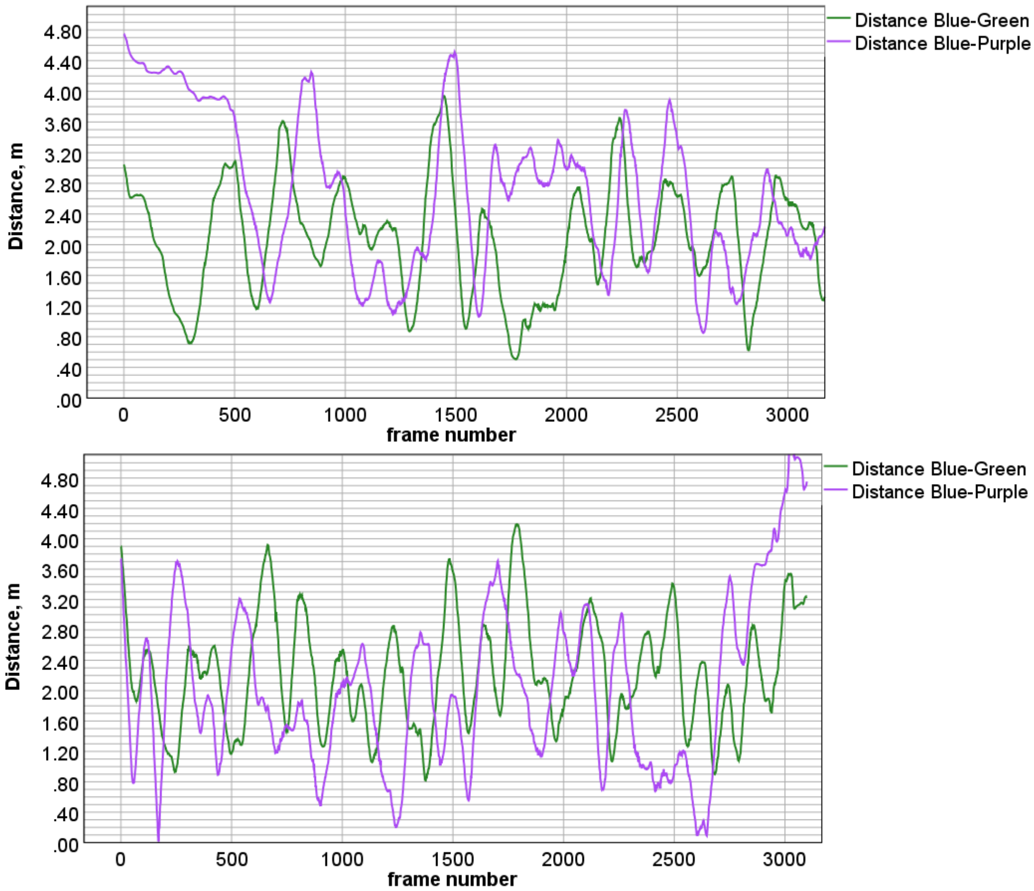

Figure 6.22: Distance between two colocated (green line) and two distributed (purple line) co-players during the test on the example of a team from the *CollisionInformed* (top) and the *CollisionIgnorant* (bottom) group.

type of placement (co-placed vs separated), the instruction group *CollisionIgnorant* vs *CollisionInformed*) and gender.

The main effect of the group was not statistically significant ($F(1,40) = 0.003$, $p = .956$) and neither any of its interactions. The main effect ($F(1,40) = 0.950$, $p = .335$) and interactions of placement were not significant either.

The main effect of gender was statistically significant ($F(1,40) = 13.54$, $p < .001$): male participants walked faster than female ones.

The main effect of scenario is significant as well ($F(1,40) = 59.03$, $p < .001$): participants walked faster in the competitive than in the collaborative scenario. Plots of *MeanSpeed* reflecting gender and scenario differences can be seen in Figure 6.21. The interaction between gender and scenario was also significant ($F(1,40) = 8.94$, $p = .005$), showing that the difference in average walking speeds between the competitive and collaborative scenarios was smaller for female than for male participants.

Figure 6.22 illustrates an example of the distance between two pairs formed by three participants from a team during the test. The top graph in the figure corresponds to a team from the *CollisionInformed* group whereas the bottom graph is produced from the trajectory data of a team from the *CollisionIgnorant* group. In both teams whose data was used for the illustration, the separated participant was the purple player during the game and the co-placed participants were the blue and green players. The local minima on each distance graph correspond to moments where two players passed each other while walking between their points of interest in the virtual playground. For each distance graph, such local minima were averaged to produce the mean clearance distance between a pair of players during the game. In the example in Figure 6.22, the blue player from the *CollisionInformed* grouped demonstrated similar proxemics behavior towards the colocated (green) and the distributed (purple) co-player, whereas the blue player from the *CollisionIgnorant* grouped walked through the avatar of the distributed (purple) co-player several times.

Each game played by a team of participants resulted in three averaged distance metrics : one average clearance between colocated co-players (between two co-placed participants) and two average clearance distances between distributed co-players (between the separated participant and each of two co-placed participants). To be able to compare these metrics, we calculated an average of two clearance distances between distributed co-players. This way, two *MeanClearance* distances were taken for the analysis of each game, one between the colocated participants and one between the distributed participants. The calculation of *MinClearance* followed the same procedure.

In the mixed ANOVA for *MeanClearance*, the within-subject factors were the colocation of players(colocated vs distributed) and the type of scenario (collaborative vs competitive). The between-subject factors were the instruction group (*CollisionIgnorant* vs *CollisionInformed*) and the gender.

Colocation and scenario type produced statistically significant main effects ($F(1,11) =$

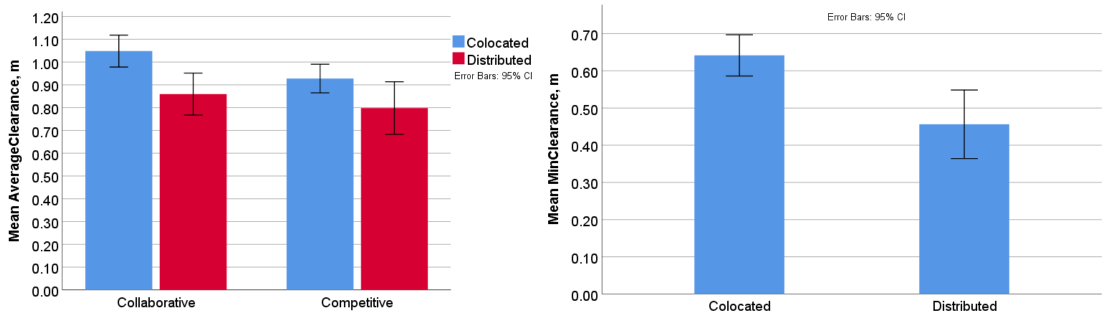

Figure 6.23: Left: *MeanClearance* between colocated and distributed participants in the collaborative and competitive scenarios. Right: *MinClearance* between the colocated and distributed participants.

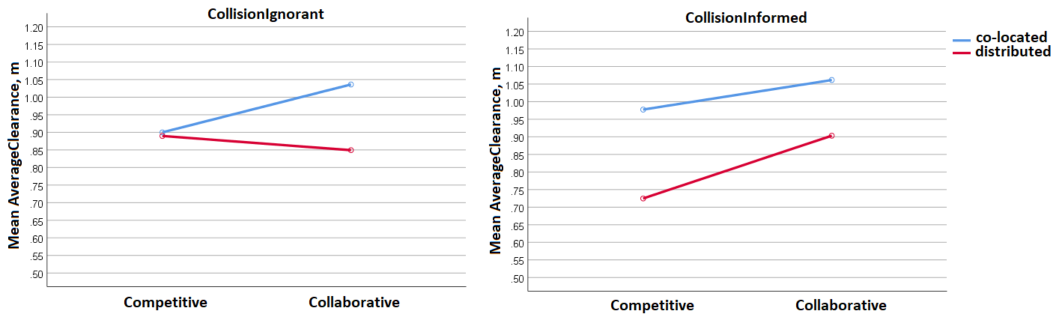

Figure 6.24: ANOVA plots of the group-scenario-colocation interaction for *MeanClearance*.

9.05, p = .012 for colocation, $F(1,11)$ = 7.34, p < .001 for scenario). Average clearance was larger between colocated co-players than between distributed co-players, and was larger in the collaborative scenario compared to the competitive one. These results are illustrated in Figure 6.23.

Neither of the between-subject factors had a significant main effect ($F(1,11)$ < .001 , p = .973 for group, $F(1,11)$ = 3.57, p = .085 for gender). The interaction between group, scenario and colocation was however significant ($F(1,11)$ = 8.35 , p = .015). The interaction effect is illustrated in Figure 6.24. While in the *CollisionInformed* group average clearance increased in the collaborative scenario compared to the competitive one independently of whether co-players were colocated or distributed, in the *Collision-Ignorant* group only the average clearance between colocated co-players increased in the collaborative scenario.

Another mixed ANOVA test with the same parameters was conducted for *MinClearance*. For *MinClearance*, only the main effect of colocation was significant ($F(1,11)$ = 10.75 , p = .007): the minimal clearance between distributed co-players was smaller than between colocated co-players. This result is illustrated in Figure 6.23.

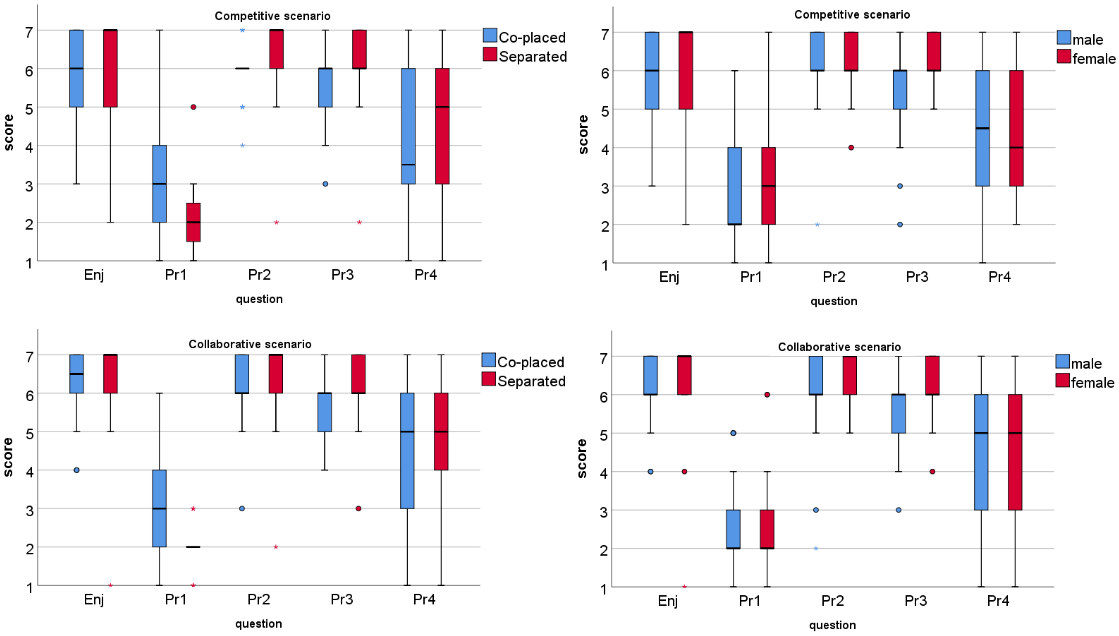

Figure 6.25: Scores of the presence-related questions in the competitive and collaborative scenarios.

Presence and embodiment

The distributions of scores of the questions *Pr1 - Pr4* were not different between the *CollisionIgnorant* and *CollisionInformed* groups in both competitive and collaborative scenarios, compared with the Mann-Whitney U tests.

Statistically significant differences between genders were found for questions *Pr2* and *Pr3* in the collaborative game, indicating that female participants were slightly more involved in the experience and rated their sense of being in the VE higher. Co-placed participants reported to have been more aware of real events (*Pr1*) and rated their sense if being in the VE (*Pr3*) slightly lower than separated participants in both scenarios.

Finally, participants reported to have enjoyed the collaborative game more than the competitive game (in the Wilcoxon signed rank test, T = 59.0, z = -2.528, p = .011, r = -.37), in accordance with many comments in the de-briefing phase.

Overall, presence scores very high in both competitive and collaborative scenarios. Corresponding boxplots can be seen in Figure 6.25. The details of the statistically significant Mann-Whitney U tests are summarized in Table 6.10.

In the self-report of embodiment, there were no differences between the *CollisionIgnorant* and *CollisionInformed groups* or between the separated and colocated participants (in the Mann-Whitney U tests). Female participants reported higher body ownership (*BdOwn1*) in both scenarios, as well as higher sense of control over their virtual bodies (*Ag1*) in the

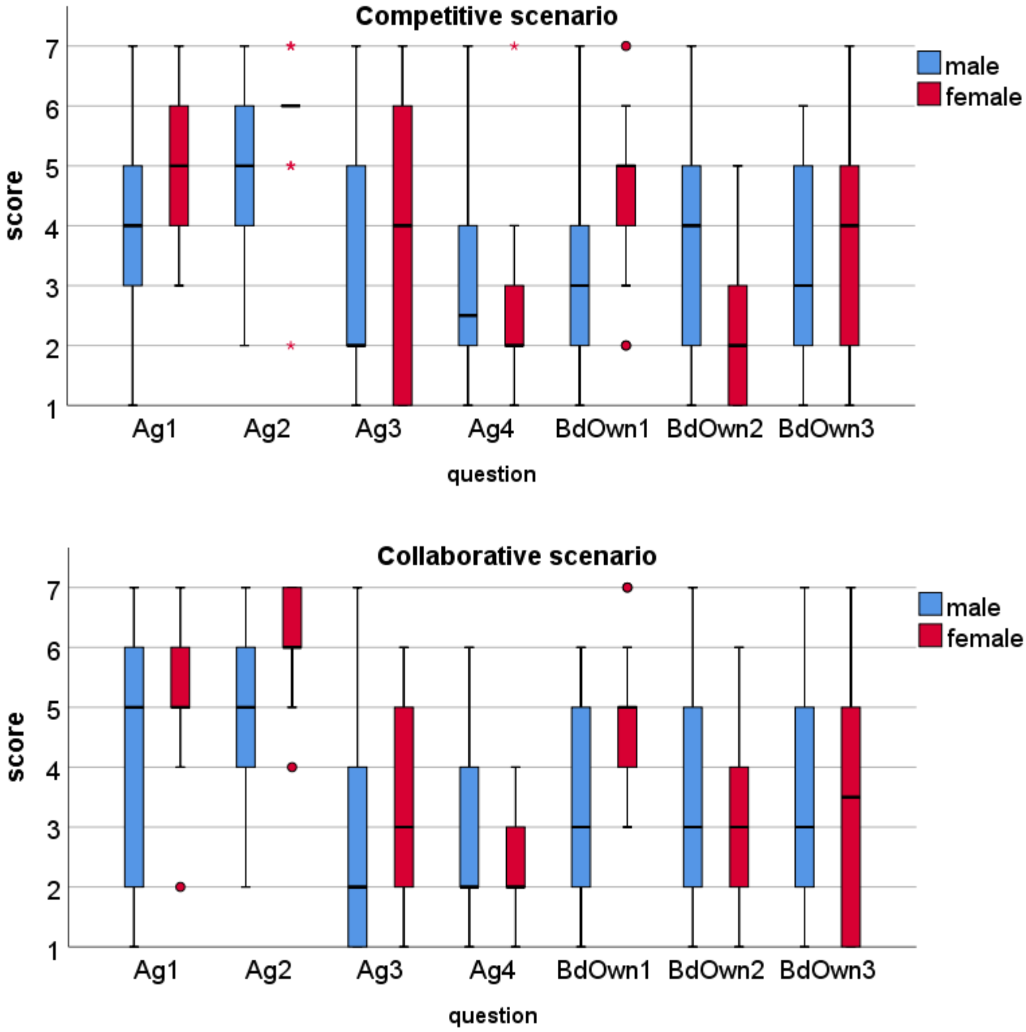

Figure 6.26: Scores of the embodiment-related questions in the competitive and collaborative scenarios.

Table 6.10: Results of the statistically significant comparisons of the presence scores in the Mann-Whitney U tests. Test statistics, p-values and the effect sizes r are shown.

question	U	z	p	r
male vs female				
Pr2 collaborative	410.0	2.234	.025	.31
Pr3 collaborative	397.0	1.965	.049	.28
co-placed vs separated				
Pr1 competitive	145.5	-2.240	.025	-.31
Pr3 competitive	324.0	2.095	.036	.29
Pr1 collaborative	149.0	-2.966	.003	-.42
Pr3 collaborative	546.5	2.265	.024	.32

Table 6.11: Results of the statistically significant comparisons of the embodiment scores of male and female participants in the Mann-Whitney U tests. Test statistics, p-values and the effect sizes r are shown.

question	U	z	p	r
BdOwn1 competitive	550.5	2.645	.008	.39
BdOwn1 collaborative	413.5	2.179	.029	.31
BdOwn2 competitive	155.5	-2.239	.025	-.33
Ag1 competitive	513.5	2.387	.017	.35
Ag2 collaborative	415.5	2.433	.015	.34

competitive scenario and higher agency of movements (*Ag2*) in the collaborative scenario. In addition, male participants reported higher sense of the virtual body belonging to somebody else (*BdOwn2*) in the competitive scenario. Corresponding boxplots of the scores are shown in Figure 6.26. The details of statistical comparisons are summarized in Table 6.11.

Copresence and social presence

Each co-placed participant had one colocated and one distributed co-player while each separated participant has two distributed co-players. Separated participants often gave different answers concerning their two distributed co-players. To estimate this difference in answers, we computed a metric which was the sum of difference points in answers on all questions about two distributed co-players, however separate for the competitive and collaborative scenarios. This metric was bigger for the competitive scenario (median M = 25 points) than for the collaborative scenario (median M = 9 points) in the Wilcoxon matched-pair signed-rank test, p = 0.001. To be able to compare the perception of distributed co-players by separated and co-placed participants, we calculated an average value for each separated player if the difference between their answers about both distributed co-players was not greater than 3 points. This limitation was introduced to

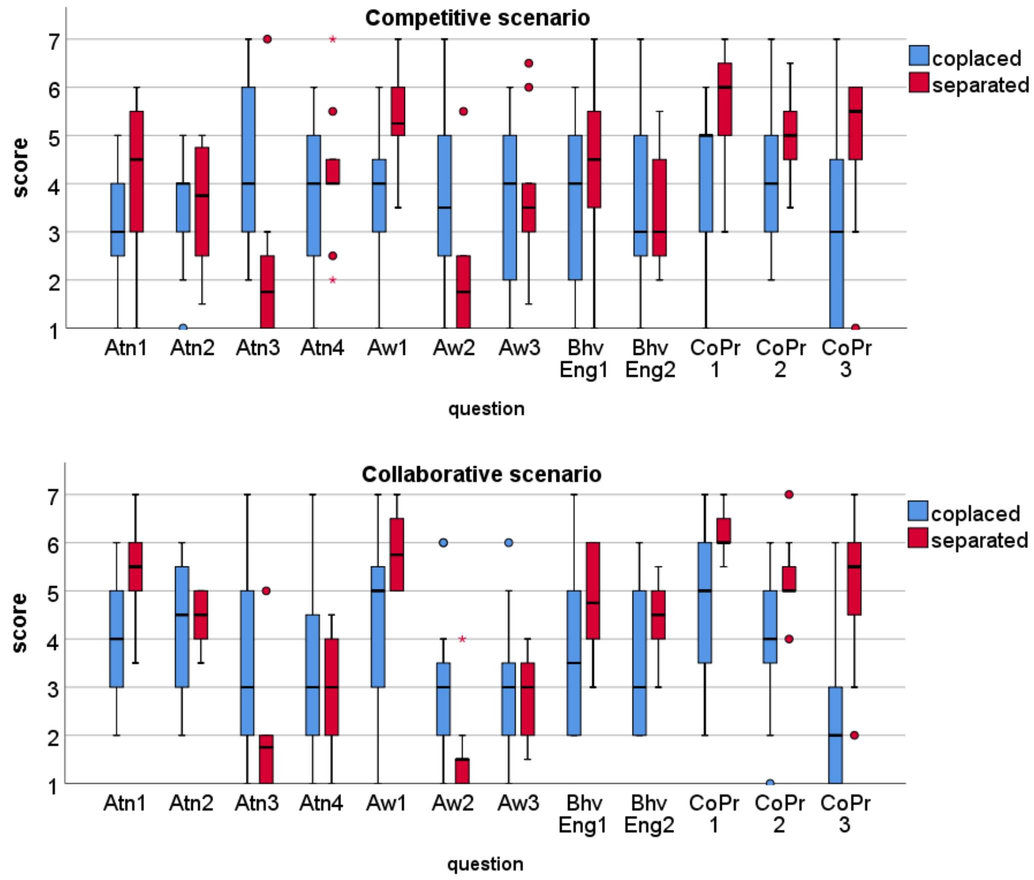

Figure 6.27: Scores of the perceived copresence of distributed co-players resulting from the answers of co-placed and separated participants from the *CollisionIgnorant* group.

filter out rare cases in which the difference between responses was large.

In the *CollisionIgnorant* group, distributed co-players were perceived as less co-present by co-placed participants (who also had colocated co-players) than by separated participants (who only had distributed co-players) in the resulting scores of the direct copresence questions (*CoPr1 - CoPr3*) as well as in the scores of questions awareness and attentional allocation, in both competitive and collaborative scenarios. Corresponding boxplots of the resulting scores are shown in Figure 6.27. The details of statistically significant comparisons are presented in Table 6.12.

Co-placed participants from the *CollisionIgnorant* group judged their colocated co-players being significantly more copresent than their distributed co-players in all questions except one (*Aw3* in the collaborative game) in both scenarios. Boxplots of the resulting scores are shown in Figure 6.28. The details of the statistical comparisons are summarized in Table 6.13.

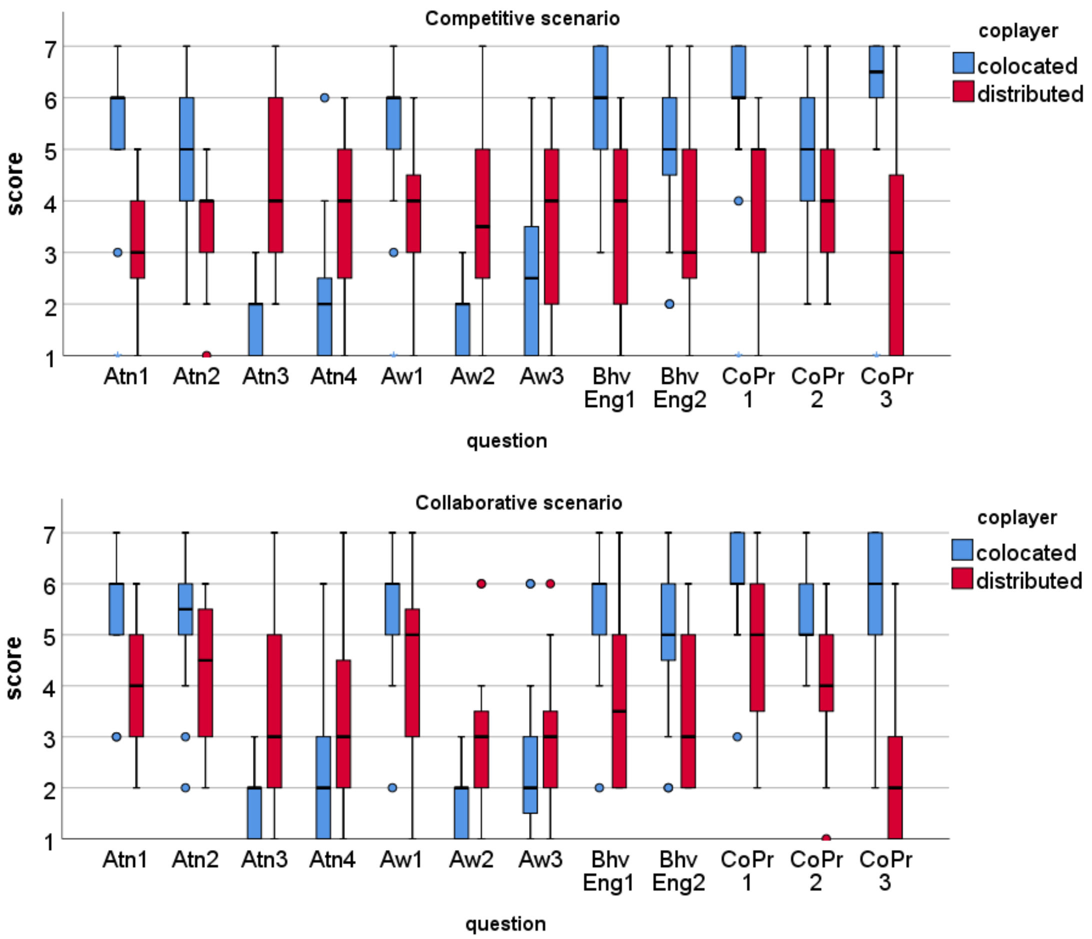

Figure 6.28: Scores of the perceived copresence of colocated and distributed co-players resulting from the answers of co-placed participants from the *CollisionIgnorant* group.

In the *CollisionInformed* group, there were no statistically significant differences in the self-report of copresence of distributed co-players reported by separated and co-placed participants. Co-placed participants from the *CollisionInformed* group judged their colocated co-players as more copresent than their distributed co-players only in five questions after the competitive game and in three questions after the collaborative game. Plots of the copresence scores in the *CollisionInformed* group can be seen in Figure 6.29. The details of the statistically significant comparisons of the resulting scores in regard of colocated and distributed co-players are summarized in Table 6.14.

For separated participants, belonging to the *CollisionIgnorant* or *CollisionInformed* group produced no significant differences in answers about co-players. This result was expected as the instructions were in fact the same for separated participants in both

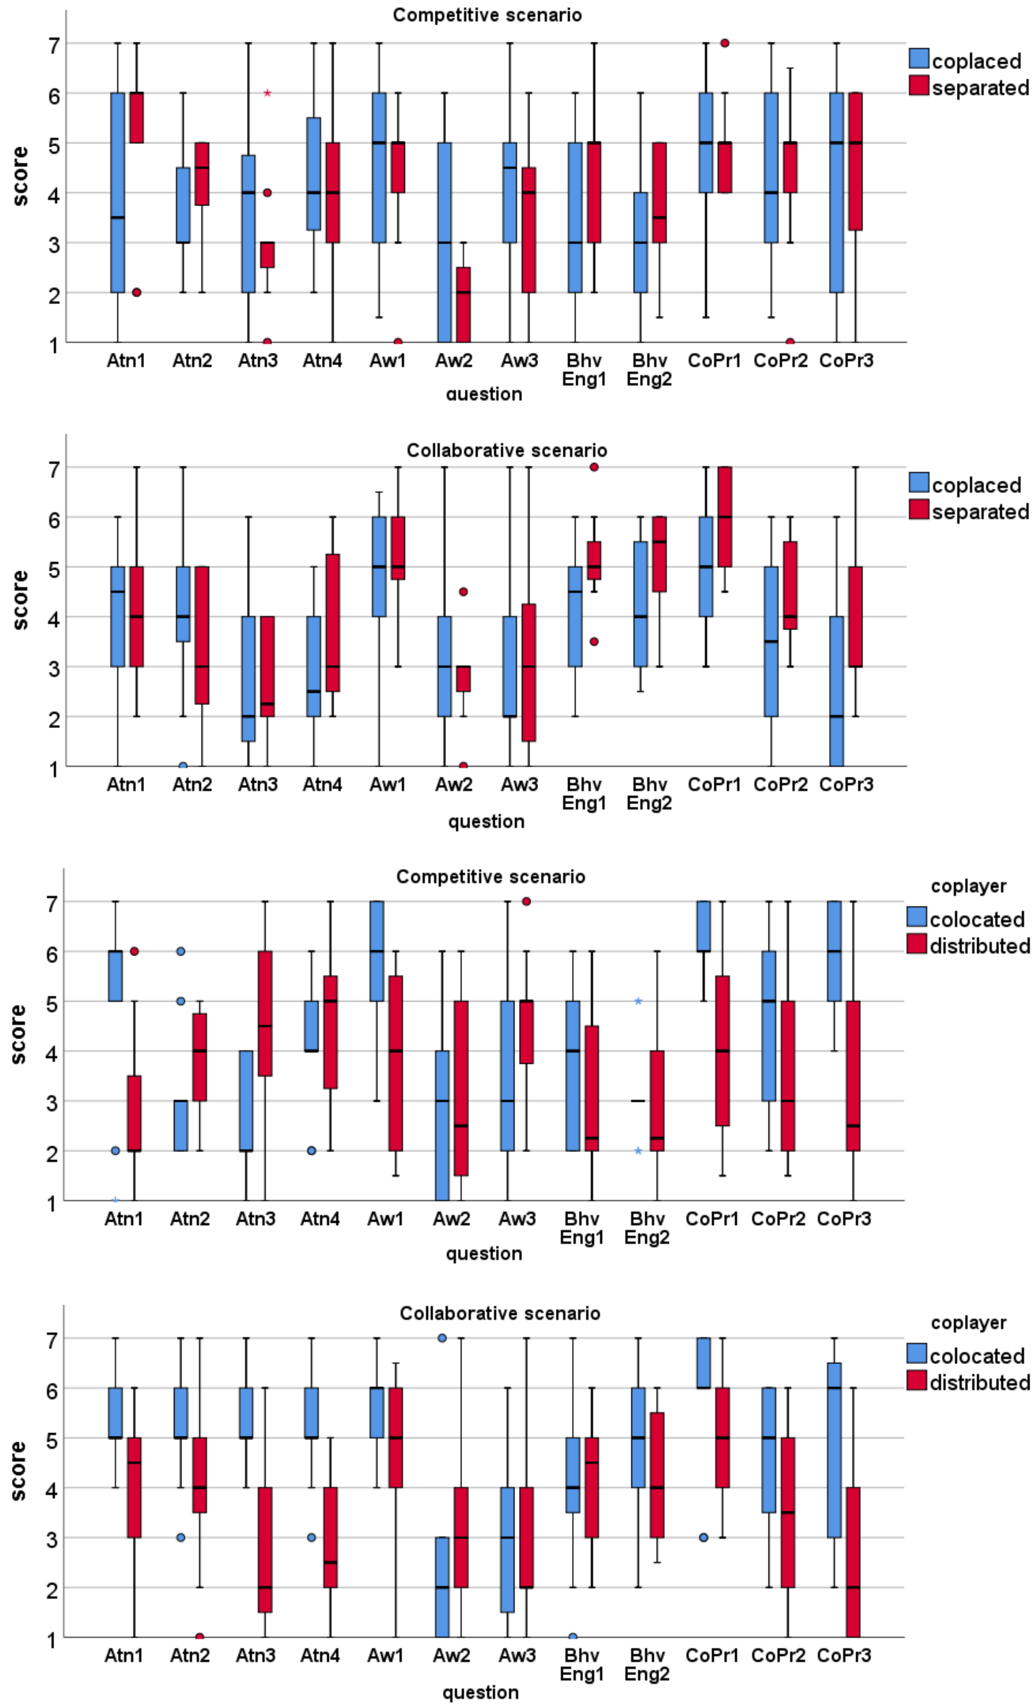

Figure 6.29: Scores of the perceived copresence in the *CollisionInformed* group.

Table 6.12: Results of the statistically significant between-placement comparisons of the scores of questions about distributed co-players answered by the participants from the *CollisionIgnorant*. Mann-Whitney U test was used; test statistics, p-values and effect sizes r are shown.

question	U	z	p	r
Competitive scenario				
CoPr1	165.0	2.925	.003	.53
CoPr2	156.5	2.524	.011	.46
CoPr3	159.0	3.305	$<$ .001	.61
Aw1	170.0	3.133	.001	.57
Aw2	35.5	-2.863	.003	-.52
Atn4	30	-3.112	$<$ .001	-.57
Collaborative scenario				
CoPr1	164.0	2.887	.04	.53
CoPr2	153.5	2.426	.017	.44
CoPr3	132.5	2.037	.044	.38
Aw1	158.5	2.632	.008	.48
Aw2	26.5	-3.301	$<$ .001	-.60
Atn1	143.5	2.549	.01	.47
Atn3	43	-2.600	.011	-.47
BhvEng1	148.0	2.148	.035	.39

groups. However gender had a significant effect in a number of questions, mostly in the collaborative scenario: in the separated placement, female participants reported higher copresence scores than male participants. Corresponding boxplots are shown in Figure 6.30, the details of statistically significant comparisons are summarized in Table 6.15

For separated participants, the type of scenario (competitive vs collaborative) produced significant differences in the copresence scores of only three questions. For co-placed participants from both *CollisionIgnorant* and *CollisionInformed* groups, only some of the copresence scores in regard of their distributed co-players were significantly different between the competitive and collaborative scenarios.

6.2.4 Discussion

In the *CollisionIgnorant* group, placement of participants in the tracking area had strong influence on their self-report of copresence. Co-placed participants gave higher scores in respect to their colocated co-players than their distributed co-players in all sub-categories of the copresence and social presence assessment. Similarly, participants from *CollisionIgnorant* group judged their distributed co-players more copresent if they themselves were alone in a separate tracking space. Significant differences were found in the direct self-report of copresence as well as in mutual awareness (that largely contributes

Table 6.13: Significant differences in the scores of the perceived copresence of colocated and distributed co-players resulting from the answers given by co-placed participants from the *CollisionIgnorant* group. Wilcoxon matched-pair signed-rank test is used; test statistics, p - values and the effect sizes r are shown.

question	T	z	p	r
Competitive scenario				
CoPr1	27.0	-2.939	.003	-.46
CoPr2	6.0	-2.614	.009	-.41
CoPr3	16.0	-3.043	.002	-.48
Aw1	4.0	-3.468	.001	-.55
Aw2	136.0	3.545	<.001	.56
Aw3	96.0	2.058	.040	.33
Atn1	15.5	-3.221	.001	-.51
Atn2	24.0	-2.889	.004	-.46
Atn3	190.0	3.848	<.001	.61
Atn4	134.5	2.769	.006	.44
BhvEng1	9.0	-3.210	.001	-.51
BhvEng2	24.0	-2.294	.022	-.36
Collaborative scenario				
CoPr1	5.0	-2.684	.007	-.42
CoPr2	4.50	-3.209	.001	-.51
CoPr3	2.5	-3.738	<.001	-.59
Aw1	9.0	-2.792	.005	-.44
Aw2	136.0	3.594	<.001	.57
Atn1	2.5	-3.293	.001	-.52
Atn2	14.0	-2.219	.026	-.35
Atn3	89.0	3.063	.002	.48
Atn4	113.0	2.366	.018	.37
BhvEng1	0.0	-3.557	<.001	-.56
BhvEng2	10.5	-3.009	.003	-.48

Table 6.14: Significant differences in the scores of the perceived copresence of colocated and distributed co-players given by co-placed participants from the *CollisionInformed* group. The data is analyzed with Wilcoxon matched-pair signed-rank test;test statistics, p - values and the effect sizes r are shown.

question	T	z	p	r
Competitive scenario				
CoPr1	4.0	-2.448	.013	-.50
CoPr3	0.0	-2.825	.005	-.58
Aw2	55.0	2.871	.004	.59
Aw3	55.0	1.974	.048	.40
Atn3	39.5	2.034	.042	.42
Collaborative scenario				
CoPr2	0.0	-2.719	.007	-.51
CoPr3	3.0	-2.991	.003	-.57
Atn1	2.5	-2.584	.010	-.49

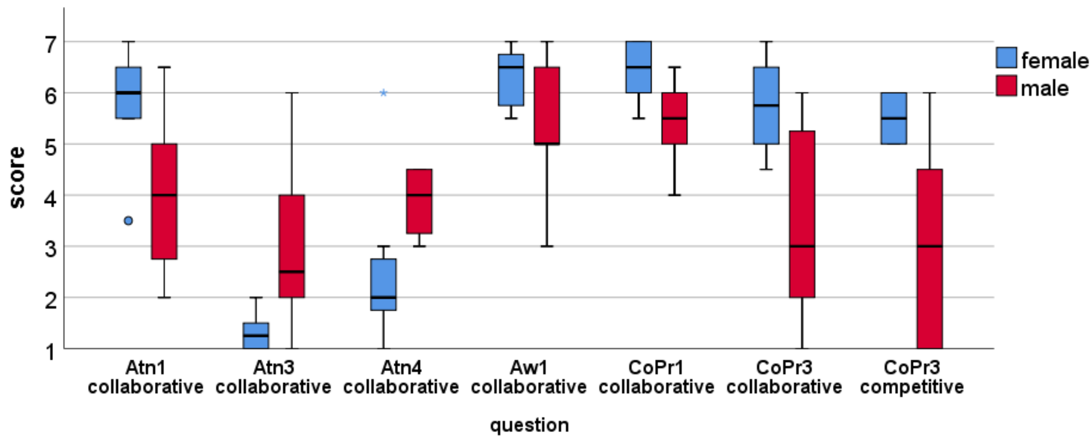

Figure 6.30: Statistically different scores resulting from the answers of female and male participants in the separated placement.

Table 6.15: Significant between-gender effects on the scores of separated participants in the Mann-Whitney U test. Median values, test statistics, *p*-values and effect sizes are shown.

question	*U*	*z*	*p*	*r*
Competitive scenario				
CoPr3	37.0	2.321	.020	.58
Collaborative scenario				
CoPr1	52.50	2.272	.023	.57
CoPr3	40.5	2.142	.032	.54
Aw2	12.0	-2.119	.034	-.53
Atn1	47.5	2.267	.023	.56
Atn3	27.0	-2.525	.012	-.63
Atn4	37.0	-2.219	.027	-.55

Table 6.16: Significant effects of the scenario type on the copresence scores of separated participants in the Wilcoxon matched-pair signed-rank test.

question	Md competitive	Md collaborative	z	p	r
CoPr1	5	6	-2.048	.041	.39
Aw1	5	5.5	-2.337	.019	.44
BhvEng2	3	4	-2.078	.038	.39

Table 6.17: Significant effects of the scenario type on the copresence scores of co-placed participants in regard of their distributed co-players in the Wilcoxon matched-pair signed-rank test.

question	Md competitive	Md collaborative	z	p	r
CollisionIgnorant group					
Aw1	4	5	-2.627	.009	-.42
Aw3	4	3	2.069	.039	.33
Atn1	3	4	-2.391	.017	-.38
Atn3	4	3	2.121	.034	.34
CollisionInformed group					
Aw2	5	3	2.124	.034	.43
Atn3	4	2	2.057	.040	.42
Atn4	4.5	2	2.561	.010	.52
BhvEng1	3	5	-1.979	.048	-.40
BhvEng2	2.5	5	-2.567	.010	-.52

to copresence according to Biocca et al. [21]) and attention, especially in the collaborative scenario. This is an indication that the difference in the perception of co-players might be caused by the a-priori knowledge that some co-players are right physically close and some are not. In the *CollisionInformed* group, the difference in self-report of experienced copresence by co-placed participants in regard of their colocated and distributed co-players was not as prominent. The effect of the placement of the participants giving responses disappeared as well. Taking into account that responses of separated participants did not differ between the *CollisionIgnorant* and *CollisionInformed* groups, this is an indication of an increase in the perceived copresence of distributed co-players among the co-placed participants from the *CollisionInformed* group.

In both groups, co-placed participants reported high levels of worrying about possible collisions with their colocated co-players. The scores of awareness of real events were higher and self-report of the experienced presence lower (although both measures respectively rather low and high in absolute terms) among co-placed participants than separated participants. It is possible that the novelty of VR experience, especially in respect to walking around in the virtual world, and the fast-paced nature of the scenarios contributed to higher alertness of co-placed participants.

Observed proxemics behavior only partially supported the differences in reported perceived copresence of distributed co-players between the two groups. Even though co-placed participants from the *CollisionIgnorant* group often walked through avatars of their distributed co-players whereas co-placed participants from *CollisionInformed* group where explicitly forbidden to do so, the average clearance between colocated co-players was larger than between distributed co-players in both *CollisionIgnorant* and *CollisionInformed* groups. The scenario-placement-group interaction was the only reflection of the difference between two groups in the locomotor trajectories. This interaction can be interpreted in the following way: participants got close to their test partners during the competitive task as they had to be quick to win and could not pay too much attention on collision avoidance; when this need disappeared in the collaborative scenario they kept larger distances for safety. Because distributed co-players were perceived as copresent less strongly in the *CollisionIgnorant* group there was no need to keep larger distances to them during the collaborative task.

Motion patterns displayed by the participants extend the results of the first experiment described in this chapter. As in the previous experiment, the average walking speeds did not differ between the distributed and colocated setup. In the previous experiment, we found a priming effect where participants did not walk through each others' avatars in the distributed setup if they first performed the test in the colocated setup. In this user study, such priming condition was absent as participants never switched between co-placed and separated placement. As a result we observed participants from the *CollisionIgnorant* group regularly walk through the avatars of their distributed co-players.

Behavioral markers such as the maintenance of appropriate interpersonal distances are thought of as a demonstration of the perceived copresence [8, 9]. However, our results show that certain proxemics patterns (in our case, forced necessity to avoid colliding with avatars), which could also be perceived as more respectful interaction, could in their turn prime users to feel more copresent with others. The assignment of avatar colors was not known to participants from the *CollisionInformed* group in the beginning of each game. It is possible that they paid more attention to each other trying to figure out which test partner was physically close and this attention effort led to equalized subjective copresence responses for both colocated and distributed co-players. We expect that co-placed participants would experience even lesser difference in the perception of their colocated and distributed participants if they would not know the type of their own placement (alone or sharing a tracking space). Such an experiment however, apart from being difficult to organize, would hardly correspond to any possible real scenario.

In a real shared VR application with mixed physical colocation, user actions like walking through or trying to touch avatars are unlikely to be forbidden. It is however likely that users would only discover the appearance of others' avatars gradually and would need to spend time to figure out which avatars represent distributed and which colocated co-players. Moreover, situations are possible where some users join a scenario late and some leave while their co-players are still immersed. For users already immersed in the scenario, the real location of late-joining users could stay unknown for a period of time. In some scenarios, avatars might even be designed to look similar. This way, the results of our both test groups are relevant for out-of-the-lab experiences.

The experiment showed clear gender differences, with female participants demonstrating more careful behavior and reporting higher illusion of copresence in separated placement in the collaborative game. Slightly higher ratings of presence and body ownership might contribute to this effect. As in the previous experiment described in this chapter, higher scores among female participants might be correlated with their lesser reported familiarity with VR and computer games. However, this hypothesis would need to be tested in a separate dedicated experiment.

We expected to see the increase of social presence in the collaborative scenario, however the influence of the scenario type was not conclusive. Larger clearance distances in the collaborative scenario might reflect greater social presence but can be also explained by greater effort of collision avoidance when completion time was not crucial, especially among colocated co-players. Higher scores of mutual attention and behavioral engagement would reflect the increase of self-report of social presence. Such increase was partially observed in the scores of co-placed participants from both *CollisionIgnorant* and *CollisionInformed* groups but only in respect to their distributed co-players. The reason could have been, in part, the necessity to pay more attention to colocated co-players during the competitive game to not run into them, while the slower pace of the collaborative game allowed more equal distribution of attentional resources among both co-players.

Many participants wanted to be able to talk during the collaborative scenario. The absence of speech communication might have been an obstacle for the collaborative game. However we did not want to introduce additional factors that would contribute to the difference in perception of colocated and separated co-players. Using microphones is the only solution for communicating with distributed co-players. However, our experience of using voice communication in ImmersiveDeck showed that it is virtually impossible to filter out real voices of colocated co-players completely. The restriction on communication was a possible reason for the ambiguous results in between-scenario analysis. The other possible reason however was our underestimation of the social nature of the competitive game.

6.3 Summary

This chapter presented two experiments on proxemics and copresence in shared VR in three different conditions of physical colocation: colocated, distributed and mixed. The experimental conditions of both presented studies required large amounts of walking in the virtual space. In the first experiment, pairs of participants performed a walking task potentially requiring mutual collision avoidance in a colocated and in a distributed setup. In the second experiment, triples of participants performed a task also potentially requiring collision avoidance. In this study, one participant from a triple had a tracking space of their own whereas the remaining two shared another tracking space, the whole setup thus having a mixed colocation configuration. The described experiments lead to the following findings.

- **Collision avoidance is more demanding in shared VR compared to real environments, especially in a colocated setup.** Participants were more careful when walking together in VR than in the real environment. In accordance with the results of previous experiments, walking speeds are lower and path curvatures are bigger in VR. When users choose to avoid colliding with each other in a shared VR scenario (in the colocated setup and in some cases in the distributed setup) they do it with more caution than in the real world keeping larger clearance distances. According to participants self-report in both experiments, users are worried of possible collisions with their colocated co-players and a large portion of their attention is dedicated to avoiding them.

- **Proxemics behavior is influenced by physical colocation.** Without any specific instructions or priming influence, many participants of the both described experiments walked through avatars of distributed co-players. Walking through avatars of distributed co-players may be explained by low sense of copresence towards these co-players. Participants themselves often explained their proxemics behavior by the absence of the possibility of actual collisions and also by the curiosity to see what would happen. However, even when interpersonal distances were maintained and participants avoided their distributed co-players clearance distances were smaller between distributed co-players than between colocated ones.

- **Priming can influence proxemics behavior.** We observed priming effects manifested by the proxemics behavior of participants in a shared distributed VR setup. In particular, participants avoided mutual collisions in the shared distributed

VR setup after having performed the experimental task in the shared colocated VR setup where they had to avoid mutual collisions to not physically run into each other. This result demonstrates that priming could be intentionally used in shared distributed VR scenarios that require the maintenance of conventional interpersonal distances.

- **Proxemics behavior can influence self-report of copresence.** Proxemics behavior, and in particular the maintenance of interpersonal distances, is considered to be a behavioral marker of copresence. The results of the second experiment described in this chapter suggest that certain proxemics behavior itself, even if it is dictated by the rules of the experiment, could lead to the increase in the reported copresence. This conclusion is drawn from the fact that those co-placed participants who were instructed to avoid walking through avatars of others reported higher copresence towards their distributed co-players than those participants who did not receive such instruction and as a result did not avoid collisions with their distributed co-players. Similarly, separated participants who were instructed to not walk though others' avatars reported higher copresence towards their distributed co-players than co-placed participants who did not receive such instruction.
- **There may be gender effects on the elicitation of the sense of presence and copresence. Alternatively, presence and copresence may be correlated with familiarity with VR and computer games.** In both experiments described in this chapter, female participants reported stronger illusion of presence and higher copresence in shared distributed VR than male participants. In addition, female participants reported stronger sense of body ownership over their avatars in the second experiment. Finally, female participants always avoided colliding with avatars of others as well as with unanimous objects in the playground in the second experiment. Such behavior preserving the personal space is normally attributed to strong illusion of presence and copresence. In both experiments, female participants reported lower familiarity with VR and computer games than male participants. Therefore, the stronger experience of presence and copresence may be correlated rather with low experience of computer-generated worlds than with gender. A dedicated experiment with comparable numbers of female and male participants, each group including participants who had different levels of exposure to VR and computer games, would help to investigate the effect further.

CHAPTER 7

Conclusion

7.1 Summary

The research presented in this dissertation has focused on walkable large-scale multi-user VR environments. First, we presented ImmersiveDeck, our multi-user VR platform that allows several users at a time to explore a shared or separate virtual worlds while walking in a large-scale tracking area. Within the scope of this dissertation, ImmersiveDeck was used to conduct four experiments on walkable multi-user VR.

First two experiments addressed research questions connected to navigation by walking performed simultaneously by users of colocated non-shared VR. Specifically, we first studied whether colocated users can be aware of each other's proximity within the shared tracking area and then explored methods of preventing imminent collisions in colocated non-shared VR. Our results show that HMD-based VR can produce immersion so strong that users do not notice others being present in their immediate proximity, at times even in the events of light collisions. This fact makes the development of collision prevention methods the crucial research goal in the area of colocated non-shared VR. We suggested a method for preventing user-on-user collisions by displaying avatars at respective positions of users for short amounts of time when users are within a possible collision range. The results of our conducted user study suggest that this method is effective for preventing imminent collisions. Further, user preference for the visual appearance of a notification avatar depends on the application scenario: human avatars are preferred in exploration tasks whereas an abstract non-distracting bounding box is reported to better suit scenarios with specific goals and interactive content. The detailed summary of our findings concerning colocated non-shared VR can be found in the end of Chapter 5.

The following two experiments investigated locomotion and co-presence in shared VR in the colocated, distributed and mixed physical colocation setups. The main goal of the first of these two experiments was to compare locomotor behavior displayed by users of shared VR to that of the real world. In the second experiment, we investigated locomotor behavior and reported experience of copresence of those participants who used a tracking space alone and of those who shared a tracking space within a three-user mixed colocation setup. The results showed, first, that users display more cautious behavior while walking together in colocated shared VR than in the real world: they walk slower and take larger detours around each other during collision avoidance. Participants of the both experiments also reported being worried about colliding with each other in the colocated setup. In the distributed setup, participants often walked through each other's avatars. Conventional collision avoidance behavior was however preserved if the experiment in the distributed setup was preceded by a test run in the colocated setup. In the three-user experiments with mixed physical colocation, assigned colocation conditions affected locomotor and proxemics behaviour of participants. In this case, again, participants displayed more cautious collision avoidance behavior in respect to their colocated co-players than in respect to their distributed co-players. Physical colocation also influenced self-report of the experienced copresence given by the participants: when participants were able to walk through avatars of their distributed co-players these distributed co-players were perceived as less copresent as the colocated ones. Our experimental findings concerning shared VR setups are summarized in the end of Chapter 6.

The experiments presented in this thesis were set in conditions designed to resemble possible application scenarios as closely as possible. To this end, all four experiments were conducted in a multi-user setup and not with simulated users. To our knowledge, our user studies are among the first ones to include multiple users walking in VR. We hope that their findings will be useful for designers of multi-user VR applications and contribute to the development of comfortable and effective experiences.

7.2 Open questions

The following are open questions connected to the contributions described in this dissertation that we plan to investigate in the future research.

Scalability of the experimental results

The presented experiments were conducted with two or three users at a time. Further experiments with larger amount of users would be beneficial to verify whether our findings can be extended to situations where larger groups of users share a physical and/or virtual space.

The question of scalability applies to all multi-user VR setups. In colocated non-shared VR, levels of mutual awareness might be somewhat higher than those observed in our experiment if the tracking space is shared with a large number of physically colocated users. Similarly, our proposed avatar notification strategy might have different effects if notification avatars appear much more often because of the increased amount of physically colocated users. In shared VR scenarios, the increase of the amount of users with whom the virtual world is shared might lead to the increase of the difficulty of avoiding collisions. In this case, users might skip collision avoidance at higher rates in distributed shared VR. However, it is not the large amount of users itself that might produce the described effects but the increased user density in the real and virtual spaces.

Long-term effects of immersion

Our studies were mostly conducted with unexperienced VR users. Even though some participants had previous experience of immersion with HMDs or even navigating VR worlds by walking they were not truly advanced users of walkable VR. Longitudinal studies are needed to see whether our results persist with time. Especially, the effect of prolonged exposure to walkable VR on locomotion and proxemics within VR worlds is of interest. Our results showed that users display more careful locomotion behavior in VR, especially in the colocated shared VR setup, than in the real world. They walk slower and take larger detours to avoid colliding with each other. These findings extend the results of previous research on walking and avoiding static obstacles in VR. One of the possible explanations of this more cautious behavior in VR is that users are not used to walking in virtual worlds. It is possible that when users get a better understanding of the capabilities of the mediating technologies with time their locomotion patterns would become closer to those exhibited in the real space. However, research is needed to verify this hypothesis.

Further research on priming effects

We observed a priming effect in the experiment on mutual collision avoidance in shared VR: having performed collision avoidance in the colocated shared VR setup, participants continued doing so in the subsequent distributed shared VR condition. When the distributed shared VR condition was the first in the experiment, participants did not display conventional collision avoidance behavior.

Shared VR scenarios can potentially benefit from the use of similar priming effects, especially in cases where priming leads to the increase in the perception of copresence and social presence experienced by users. Additionally, the question arises whether situations could happen where users adapt to the absence of collisions after a long-term immersion in a distributed shared VR setup and subsequently stop being alert about collision avoidance in colocated shared VR. Research is needed to gain a better understanding of situations that provoke priming, to make use of it in leveraging copresence and potentially to prevent dangerous for users situations.

Effects of pre-existing social relations between users of VR

Another research direction is to study how social relations between users of shared walkable VR affect user experience in terms of copresence and proxemics behavior.

Shared VR scenarios can be used by social groups and by strangers alike. For example, groups of friends or families visit location-based VR entertainment centers to spend time together over a group activity, in the same way as people spend time playing laser tag, bowling, escape rooms etc. Another scenario may include strangers meeting for the first time in VR. For example, distributed teams of employees of the same company or members of the same organization may have joint rehearsals or team training activities in VR.

Would groups of friends and strangers display similar proxemics behavior in VR? For example, would walking through avatars of distributed co-players that we observed in both our experiments in shared VR be repeated in case of strangers training together in VR? These questions have practical importance for the design of shared VR experiences.

Bibliography

[1] J. Achenbach, T. Waltemate, M. E. Latoschik, and M. Botsch. Fast generation of realistic virtual humans. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*, VRST '17, pages 12:1–12:10, New York, NY, USA, 2017. ACM.

[2] D.B. Anderson, J.W. Barrus, J.H. Howard, C. Rich, C. Shen, and R.C. Waters. Building multiuser interactive multimedia environments at merl. *IEEE MultiMedia*, 2(4):77–82, 1995.

[3] F. Argelaguet Sanz, A.-H. Olivier, G. Bruder, J. Pettré, and A. Lécuyer. Virtual Proxemics: Locomotion in the Presence of Obstacles in Large Immersive Projection Environments. In *IEEE VR 2015*, pages 75 – 80, 2015.

[4] M. Azmandian, T. Grechkin, M. Bolas, and E. Suma. Automated path prediction for redirected walking using navigation meshes. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 63–66, March 2016.

[5] M. Azmandian, T. Grechkin, and E. S. Rosenberg. An evaluation of strategies for two-user redirected walking in shared physical spaces. In *2017 IEEE Virtual Reality (VR)*, pages 91–98, March 2017.

[6] E. R. Bachmann, M. Zmuda, J. Calusdian, X. Yun, E. Hodgson, and D. Waller. Going anywhere anywhere: Creating a low cost portable immersive ve system. In *2012 17th International Conference on Computer Games (CGAMES)*, pages 108–115, July 2012.

[7] J.N. Bailenson, E. Aharoni, A.C. Beall, R.E. Guadagno, A. Dimov, and J. Blaskovich. Comparing behavioral and self-report measures of embodied agents social presence in immersive virtual environments. In *Proceedings of Presence 2004*, October 2004.

[8] J.N. Bailenson, J. Blascovich, A.C. Beall, and J.M. Loomis. Interpersonal distance in immersive virtual environments. *Personality and Social Psychology Bulletin*, 29(7):819–833, 2003.

[9] J.N. Bailenson, K. Swinth, C. Hoyt, S. Persky, A. Dimov, and J. Blascovich. The independent and interactive effects of embodied-agent appearance and behavior on self-report, cognitive, and behavioral markers of copresence in immersive virtual environments. *Presence: Teleoper. Virtual Environ.*, 14(4):379–393, August 2005.

[10] J.N. Bailenson and N. Yee. A longitudinal study of task performance, head movements, subjective report, simulator sickness, and transformed social interaction in collaborative virtual environments. *Presence: Teleoper. Virtual Environ.*, 15(6):699–716, December 2006.

[11] J.N. Bailenson, N. Yee, D. Merget, and R. Schroeder. The effect of behavioral realism and form realism of real-time avatar faces on verbal disclosure, nonverbal disclosure, emotion recognition, and copresence in dyadic interaction. *Presence*, 15:359–372, 2006.

[12] D. Banakou, R. Groten, and M. Slater. Illusory ownership of a virtual child body causes overestimation of object sizes and implicit attitude changes. 110(31):12846–12851, 2013.

[13] D. Banakou and M. Slater. Body ownership causes illusory self-attribution of speaking and influences subsequent real speaking. *Proceedings of the National Academy of Sciences*, 111(49):17678–17683, 2014.

[14] B.J. Basten, S.E. Jansen, and I. Karamouzas. Exploiting motion capture to enhance avoidance behaviour in games. In *Proceedings of the 2Nd International Workshop on Motion in Games*, MIG '09, pages 29–40, Berlin, Heidelberg, 2009. Springer-Verlag.

[15] S. Beck, A. Kunert, A. Kulik, and B. Froehlich. Immersive group-to-group telepresence. *IEEE Transactions on Visualization and Computer Graphics*, 19(4):616–625, April 2013.

[16] S. Benford, J. Bowers, L.E. Fahlén, and C. Greenhalgh. Managing mutual awareness in collaborative virtual environments. In *Proceedings of the Conference on Virtual Reality Software and Technology*, VRST '94, pages 223–236, River Edge, NJ, USA, 1994. World Scientific Publishing Co., Inc.

[17] S. Benford, J. Bowers, L.E. Fahlén, C. Greenhalgh, and D. Snowdon. User embodiment in collaborative virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '95, pages 242–249, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

[18] S. Benford and C. Greenhalgh. Introducing third party objects into the spatial model of interaction. In *Proceedings of CSCW'97: The Fifth European Conference on Computer Supported Cooperative Work*, pages 189–204, Lancaster, UK, 1997. Kluwer Academic Publishers.

[19] G. Bente, S. Ruggenberg, B. Tietz, and S. Wortberg. Measuring behavioral correlates of social presence in virtual encounters. In *Proceedings of Presence 2004*, October 2004.

[20] F. Biocca. The cyborg's dilemma: Progressive embodiment in virtual environments [1]. *Journal of Computer-Mediated Communication*, 3(2):JCMC324, 1997.

[21] F. Biocca, C. Harms, and J.K. Burgoon. Toward a more robust theory and measure of social presence: Review and suggested criteria. *Presence: Teleoper. Virtual Environ.*, 12(5):456–480, October 2003.

[22] F. Biocca, C. Harms, and J. Gregg. The networked minds measure of social presence: Pilot test of the factor structure and concurrent validity. 2001.

[23] J. Blascovich. The social life of avatars. chapter Social Influence Within Immersive Virtual Environments, pages 127–145. Springer-Verlag, Berlin, Heidelberg, 2002.

[24] A. Bönsch, S. Radke, H. Overath, L.M. Asche, J. Wendt, T. Vierjahn, U. Habel, and T. Kuhlen. Social vr: How personal space is affected by virtual agents' emotions. *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 199–206, 2018.

[25] A. Bönsch, B. Weyers, J. Wendt, S. Freitag, and T. W. Kuhlen. Collision avoidance in the presence of a virtual agent in small-scale virtual environments. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 145–148, March 2016.

[26] M. Botvinick and J. Cohen. Rubber hands 'feel' touch that eyes see. *Nature*, February 1998.

[27] S.T. Bulu. Place presence, social presence, co-presence, and satisfaction in virtual worlds. *Comput. Educ.*, 58(1):154–161, January 2012.

[28] M.E. Caplan and M. Goldman. Personal space violations as a function of height. *The Journal of Social Psychology*, 114(2):167–171, 1981.

[29] C. Carlsson and O. Hagsand. Dive a multi-user virtual reality system. In *Virtual Reality Annual International Symposium, 1993., 1993 IEEE*, pages 394–400. IEEE, 1993.

[30] C. Carlsson and O. Hagsand. Dive—a platform for multi-user virtual environments. *Computers and Graphics*, 17(6):663 – 669, 1993.

[31] J. Casanueva and E. Blake. The effects of group collaboration on presence in a collaborative virtual environment. In *Virtual Environments 2000*, pages 85–94, Vienna, 2000. Springer Vienna.

[32] G. Cirio, M. Marchal, T. Regia-Corte, and A. Lécuyer. The magic barrier tape: A novel metaphor for infinite navigation in virtual worlds with a restricted walking workspace. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology*, VRST '09, pages 155–162, New York, NY, USA, 2009. ACM.

[33] S.H. Creem-Regehr, J.K. Stefanucci, W.B. Thompson, N. Nash, and M. McCardell. Egocentric distance perception in the oculus rift (dk2). In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception*, SAP '15, pages 47–50, New York, NY, USA, 2015. ACM.

[34] S.H. Creem-Regehr, P. Willemsen, A.A. Gooch, and W.B. Thompson. The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual indoor environments. *Perception*, 34(2):191–204, 2005.

[35] J.E. Cutting, P.M. Vishton, and P.A. Brare. How we avoid collisions with stationary and moving objects. *Psychological Review*, 102, 1995.

[36] J.M. Dabbs and N.A. Stokes. Beauty is power: The use of space on the sidewalk. *Sociometry*, 38(4):551–557, 1975.

[37] S. Daher, K. Kim, M. Lee, R. Schubert, G. Bruder, J. Bailenson, and G. Welch. Effects of social priming on social presence with intelligent virtual agents. In *Intelligent Virtual Agents*, pages 87–100, Cham, 2017. Springer International Publishing.

[38] N. David, A. Newen, and K. Vogeley. The "sense of agency" and its underlying cognitive and neural mechanisms", journal = "consciousness and cognition. 17(2):523 – 534, 2008. Social Cognition, Emotion, and Self-Consciousness.

[39] T. Ducourant. Timing and distance characteristics of interpersonal coordination during locomotion. *Neurosci. Lett.*, 389:6–11, 2005.

[40] T. Duval, T.T.H. Nguyen, C. Fleury, A. Chauffaut, G. Dumont, and V. Gouranton. Improving awareness for 3d virtual collaboration by embedding the features of users' physical environments and by augmenting interaction tools with cognitive feedback cues. *Journal on Multimodal User Interfaces*, 8(2):187–197, June 2014.

[41] H.H. Ehrsson. The experimental induction of out-of-body experiences. *Science*, 317 5841:1048, 2007.

[42] L.E. Fahlén, C.G. Brown, O. Ståhl, and C. Carlsson. A space based model for user interaction in shared synthetic environments. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, CHI '93, pages 43–48, New York, NY, USA, 1993. ACM.

[43] B.R. Fajen and W.H. Warren. W.: Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psycholology: Human Perception Performance*, 2003.

[44] P. Fink, P. Foo, and W. Warren. Obstacle avoidance during walking in real and virtual environments. *ACM Trans. Appl. Percept.*, 4(1), 2007.

[45] C. Fleury, A. Chauffaut, T. Duval, V. Gouranton, and B. Arnaldi. A generic model for embedding users' physical workspaces into multi-scale collaborative virtual environments. In *ICAT 2010 (20th International Conference on Artificial Reality and Telexistence)*, Adelaide, Australia, December 2010.

[46] J. Fox, S.J. Ahn, J.H. Janssen, L. Yeykelis, K.Y. Segovia, and J.N. Bailenson. Avatars versus agents: A meta-analysis quantifying the effect of agency on social influence. *Hum.-Comput. Interact.*, 30(5):401–432, September 2015.

[47] E. Frécon, G. Smith, A. Steed, M. Stenius, and O. Ståhl. An overview of the coven platform. *Presence*, 10(1):109–127, Feb 2001.

[48] E. Frécon, C. Greenhalgh, and M. Stenius. The divebone—an application-level network architecture for internet-based cves. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST '99, pages 58–65, New York, NY, USA, 1999. ACM.

[49] E. Frecon and O. Hagsand. The dive client interface. Technical report, 1995.

[50] E. Frecon and O. Hagsand. The dive/tcl behavior interface. Technical report, 1995.

[51] E. Frécon and M. Stenius. Dive: A scaleable network architecture for distributed virtual environments. *Distributed Systems Engineering*, 5(3):91, 1998.

[52] R. Fribourg, F. Argelaguet, L. Hoyet, and A. Lécuyer. Studying the Sense of Embodiment in VR Shared Experiences. In *VR 2018 - 25th IEEE Conference on Virtual Reality and 3D User Interfaces*, pages 1–8, Reutlingen, Germany, March 2018. IEEE.

[53] M. Garau. *The Impact of Avatar Fidelity on Social Interaction in Virtual Environments*. PhD thesis, 2003.

[54] M. Garau, D. Friedman, H. Widenfeld, A. Antley, A. Brogni, and M. Slater. Temporal and spatial variations in presence: Qualitative analysis of interviews from an experiment on breaks in presence. *Presence: Teleoper. Virtual Environ.*, 17(3):293–309, June 2008.

[55] M. Garau, M. Slater, D. Pertaub, and S. Razzaque. The responses of people to virtual humans in an immersive virtual environment. *Presence: Teleoper. Virtual Environ.*, 14(1):104–116, 2005.

[56] M. Garau, M. Slater, V. Vinayagamoorthy, A. Brogni, A. Steed, and A.M. Sasse. The impact of avatar realism and eye gaze control on perceived quality of communication in a shared immersive virtual environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 529–536, New York, NY, USA, 2003. ACM.

[57] M. Gerin-Lajoie, C. Richards, and B. J. McFadyen. The negotiation of stationary and moving obstructions during walking: Anticipatory locomotor adaptations and preservation of personal space. 9:242–69, 08 2005.

[58] E. Goffman. *Relations in public: microstudies of the public order*. Basic Books, New York, 1971.

[59] M. Gonzalez-Franco and T.C. Peck. Avatar embodiment. towards a standardized questionnaire. *Frontiers in Robotics and AI*, 5:74, 2018.

[60] C. Greenhalgh. Dynamic, embodied multicast groups in massive-2. Technical report, 1996.

[61] C. Greenhalgh and S. Benford. Massive: A collaborative virtual environment for teleconferencing. *ACM Trans. Comput.-Hum. Interact.*, 2(3):239–261, September 1995.

[62] C. Greenhalgh, J. Purbrick, and D. Snowdon. Inside massive-3: Flexible support for data consistency and world structuring. In *Proceedings of the third international conference on Collaborative virtual environments*, pages 119–127. ACM, 2000.

[63] O. Hagsand. Interactive multiuser ves in the dive system. *IEEE MultiMedia*, 3(1):30–39, January 1996.

[64] E.T. Hall. *The Hidden Dimension*. Anchor Books, 1966.

[65] J.J. Hartnett, K.G. Bailey, and C.S. Hartley. Body height, position, and sex as determinants of personal space. *The Journal of Psychology*, 87(1):129–136, 1974.

[66] D. Henry and T. Furness. Spatial perception in virtual environments: Evaluating an architectural application. In *Proceedings of IEEE Virtual Reality Annual International Symposium*, pages 33–40, Sept 1993.

[67] E. Hodgson, E. Bachmann, D. Vincent, M. Zmuda, D. Waller, and J. Calusdian. Weavr: a self-contained and wearable immersive virtual environment simulation system. *Behavior research methods*, 47(1):296–307, 2015.

[68] E. Hodgson, E. Bachmann, and D. Waller. Redirected walking to explore virtual environments: Assessing the potential for spatial interference. *ACM Trans. Appl. Percept.*, 8(4):22:1–22:22, December 2008.

[69] J. Holm. Collision prediction and prevention in a simultaneous multi-user immersive virtual environment. Master's thesis, Miami University, 2012.

[70] A. Kalckert and H. Ehrsson. Moving a rubber hand that feels like your own: A dissociation of ownership and agency. *Frontiers in Human Neuroscience*, 6:40, 2012.

[71] H. Kaufmann, D. Schmalstieg, and M. Wagner. Construct3d: a virtual reality application for mathematics and geometry education. *Education and information technologies*, 5(4):263–276, 2000.

[72] R.S. Kennedy, N.E. Lane, K.S. Berbaum, and M.G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology*, 3(3):203–220, 1993.

[73] K. Kilteni, R. Groten, and M. Slater. The sense of embodiment in virtual reality. *Presence: Teleoper. Virtual Environ.*, 21(4):373–387, December 2012.

[74] E. Kokkinara, K. Kilteni, and M. Slater. First person perspective of seated participants over a walking virtual body leads to illusory agency over the walking. *Scientific Reports*, 6, 2016.

[75] E. Kokkinara and M. Slater. Measuring the effects through time of the influence of visuomotor and visuotactile synchronous stimulation on a virtual body ownership illusion. *Perception*, 43(1):43–58, 2014.

[76] D. M Krum, S.-H. Kang, and T. Phan. Influences on the Elicitation of Interpersonal Space with Virtual Humans. In *Proceedings of the 2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, Tuebingen/Reutlingen, Germany, mar 2018. IEEE.

[77] R. Kuck, J. Wind, K. Riege, M. Bogen, and S. Birlinghoven. Improving the avango vr/ar framework: Lessons learned. In *Workshop Virtuelle und Erweiterte Realität*, pages 209–220, 2008.

[78] B.R. Kunz, L. Wouters, D. Smith, W.B. Thompson, and S.H. Creem-Regehr. Revisiting the effect of quality of graphics on distance judgments in virtual environments: A comparison of verbal reports and blind walking. *Attention, Perception, & Psychophysics*, 71(6):1284–1293, Aug 2009.

[79] J. Lacoche, N. Pallamin, T. Boggini, and J. Royan. Collaborators awareness for user cohabitation in co-located collaborative virtual environments. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*, VRST '17, pages 15:1–15:9, New York, NY, USA, 2017. ACM.

[80] B. Laha, J.N. Bailenson, A. Won, and J.O. Bailey. Evaluating control schemes for the third arm of an avatar. *Presence: Teleoper. Virtual Environ.*, 25(2):129–147, November 2016.

[81] E. Langbehn, E. Harting, and F. Steinicke. Shadow-avatars: A visualization method to avoid collisions of physically co-located users in room-scale vr, 03 2018.

[82] M.E. Latoschik, J.-L. Lugrin, and D. Roth. Fakemi: A fake mirror system for avatar embodiment studies. In *Proceedings of the 22Nd ACM Conference on Virtual*

Reality Software and Technology, VRST '16, pages 73–76, New York, NY, USA, 2016. ACM.

[83] M.E. Latoschik, D. Roth, D. Gall, J. Achenbach, T. Waltemate, and M. Botsch. The effect of avatar realism in immersive social virtual realities. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*, VRST '17, pages 39:1–39:10, New York, NY, USA, 2017. ACM.

[84] K.-C. Lin. Dead reckoning and distributed interactive simulation. In *Distributed Interactive Simulation Systems for Simulation and Training in the Aerospace Environment: A Critical Review*, volume 10280, page 1028004. International Society for Optics and Photonics, 1995.

[85] J. Loomis and J. Knapp. Visual perception of egocentric distance in real and virtual environments, 1999.

[86] J.-L. Lugrin, J. Latt, and M. Latoschik. Anthropomorphism and illusion of virtual body ownership. In *ICAT-EGVE 2015*, pages 1–8, 2015.

[87] S. Lynch, J. Pettré, J. Bruneau, R. Kulpa, A. Cretual, and A.-H. Olivier. Effect of virtual human gaze behaviour during an orthogonal collision avoidance walking task. In *IEEE VR 2018 - 25th IEEE Conference on Virtual Reality*, Reutlingen, Germany, Mar 2018. IEEE.

[88] S. D. Lynch, R. Kulpa, L. A. Meerhoff, J. Pettré, A. Crétual, and A.-H. Olivier. Collision avoidance behavior between walkers: Global and local motion cues. *IEEE Transactions on Visualization and Computer Graphics*, 24(7):2078–2088, July 2018.

[89] Gomez Maureira M.A. and Verbeek F. The impact of co-located play on social presence and game experience in a vr game. In *DiGRA/FDG &#3916 - Proceedings of the First International Joint Conference of DiGRA and FDG*, Dundee, Scotland, August 2016. Digital Games Research Association and Society for the Advancement of the Science of Digital Games.

[90] S. Mallaro, P. Rahimian, E.E. O'Neal, J.M. Plumert, and J.K. Kearney. A comparison of head-mounted displays vs. large-screen displays for an interactive pedestrian simulator. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*, VRST '17, pages 6:1–6:4. ACM, 2017.

[91] A. Maselli and M. Slater. The building blocks of the full body ownership illusion. *Frontiers in Human Neuroscience*, 7:83, 2013.

[92] B.J. McFadyen and D.A. Winter. Anticipatory locomotor adjustments during obstructed human walking. *Neuroscience Research Communications*, 1991.

[93] B.E. Mennecke, J.L. Triplett, L.M. Hassall, Z.J. Conde, and R. Heer. An examination of a theory of embodied social presence in virtual worlds. *Decision Sciences*, 42(2):413–450.

[94] R. Messing and F.H. Durgin. Distance perception and the visual horizon in head-mounted displays. *ACM Trans. Appl. Percept.*, 2(3):234–250, July 2005.

[95] B.J. Mohler, H.H. Bülthoff, W.B. Thompson, and S.H. Creem-Regehr. A full-body avatar improves egocentric distance judgments in an immersive virtual environment. In *Proceedings of the 5th Symposium on Applied Perception in Graphics and Visualization*, APGV '08, pages 194–, New York, NY, USA, 2008. ACM.

[96] B.J. Mohler, S.H. Creem-Regehr, W.B. Thompson, and H.H. Bülthoff. The effect of viewing a self-avatar on distance judgments in an hmd-based virtual environment. *Presence: Teleoperators and Virtual Environments*, 19(3):230–242, 2010.

[97] M. Mori, K.F. MacDorman, and N. Kageki. The uncanny valley. *IEEE Robot. Automat. Mag.*, 19:98–100, 2012.

[98] A. Mossel and H. Kaufmann. Wide area optical user tracking in unconstrained indoor environments. In *Proceedings of the The 23rd International Conference on Artificial Reality and Telexistence*, pages 108–115. IEEE, 2013.

[99] S. Narang, A. Best, and D. Manocha. Simulating movement interactions between avatars and agents in virtual worlds using human motion constraints. In *IEEE VR 2018 - 25th IEEE Conference on Virtual Reality*, Reutlingen, Germany, Mar 2018. IEEE.

[100] T. Nescher and A. Kunz. Analysis of short term path prediction of human locomotion for augmented and virtual reality applications. In *Cyberworlds (CW), 2012*, pages 15–22, 2012.

[101] T. Nescher and A. Kunz. *Using Head Tracking Data for Robust Short Term Path Prediction of Human Locomotion*, pages 172–191. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[102] J.-M. Normand, E. Giannopoulos, B. Spanlang, and M. Slater. Multisensory stimulation can induce an illusion of larger belly size in immersive virtual reality. *PLOS ONE*, 6, 01 2011.

[103] J.-M. Normand, B. Spanlang, F. Tecchia, M. Carrozzino, D. Swapp, and M. Slater. Full body acting rehearsal in a networked virtual environment-a case study. *Presence: Teleoper. Virtual Environ.*, 21(2):229–243, April 2012.

[104] V. Normand, C. Babski, S. Benford, A. Bullock, S. Carion, Y. Chrysanthou, N. Farcet, E. Frécon, J. Harvey, N. Kuijpers, et al. The coven project: Exploring applicative, technical, and usage dimensions of collaborative virtual environments. *Presence: teleoperators and virtual environments*, 8(2):218–236, 1999.

[105] K. L. Nowak and F. Biocca. The effect of the agency and anthropomorphism on users' sense of telepresence, copresence, and social presence in virtual environments. *Presence*, 12(5):481–494, Oct 2003.

[106] A.-H. Olivier, J. Bruneau, R. Kulpa, and J. Pettre. Walking with virtual people: Evaluation of locomotion interfaces in dynamic environments. *IEEE Transactions on Visualization and Computer Graphics*, 06 2017.

[107] A.-H. Olivier, A. Marin, A. Crétual, A. Berthoz, and J.Pettré. Collision avoidance between two walkers: Role-dependent strategies. *Gait & Posture*, 38(4):751 – 756, 2013.

[108] A.-H. Olivier, A. Marin, A. Crétual, and J. Pettré. Minimal predicted distance:a common metric for collision avoidance during pairwise interactions between walkers. *Gait & Posture*, 36(3):399 – 404, 2012.

[109] A.-H. Olivier, J. Ondřej, R. Pettré, J.and Kulpa, and A. Crétual. Interaction between real and virtual humans during walking: Perceptual evaluation of a simple device. In *APGV*, pages 117–124, 2010.

[110] O. Oyekoya, R. Stone, W. Steptoe, L. Alkurdi, S. Klare, A. Peer, T. Weyrich, B. Cohen, F. Tecchia, and A. Steed. Supporting interoperability and presence awareness in collaborative mixed reality environments. In *Proceedings of the 19th ACM Symposium on Virtual Reality Software and Technology*, VRST '13, pages 165–174, New York, NY, USA, 2013. ACM.

[111] T.C. Peck, S. Seinfeld, S.M. Aglioti, and M. Slater. Putting yourself in the skin of a black avatar reduces implicit racial bias. *Consciousness and Cognition*, 22(3):779 – 787, 2013.

[112] A. Peer and K. Ponto. Evaluating perceived distance measures in room-scale spaces using consumer-grade head mounted displays. *2017 IEEE Symposium on 3D User Interfaces (3DUI)*, 00:83–86, 2017.

[113] S. Razzaque, Z. Kohn, and M.C. Whitton. Redirected walking. In *Proceedings of Eurographics*, 2001.

[114] D. Roth, J.-L. Lugrin, M.E. Latoschik, and S. Huber. Alpha ivbo - construction of a scale to measure the illusion of virtual body ownership. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '17, pages 2875–2883, New York, NY, USA, 2017. ACM.

[115] R. Ruddle and S. Lessels. The benefits of using a walking interface to navigate virtual environments. *ACM Trans. Comput.-Hum. Interact.*, 16(1):5:1–5:18, 2009.

[116] M. Rygol, S. Ghee, J. Naughton-Green, and J. Harvey. Technology for collaborative virtual environments. In *Proceedings of the CVE'96 workshop*, Nottingham, September 1996.

[117] D. Scaramuzza, A. Martinelli, and R. Siegwart. A new method and toolbox for easily calibrating omnidirectional cameras. In *Proceedings of IEEE International*

Conference on Intelligent Robots and Systems (IROS 2006), Beijing, China, Oct 2006.

[118] A. Scavarelli and R.J. Teather. Vr collide! comparing collision-avoidance methods between co-located virtual reality users. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '17, pages 2915–2921, New York, NY, USA, 2017. ACM.

[119] D. Schmalstieg, A. Fuhrmann, G. Hesina, Z. Szalavári, L.M. Encarnação, M. Gervautz, and W. Purgathofer. The studierstube augmented reality project. *Presence: Teleoper. Virtual Environ.*, 11(1):33–54, February 2002.

[120] R. Schroeder. Being there together and the future of connected presence. *Presence*, 15(4):438–454, Aug 2006.

[121] R. Schroeder, A. Steed, A.-S. Axelsson, I. Heldal, Å. Abelin, J. Wideström, A. Nilsson, and M. Slater. Collaborating in networked immersive spaces: As good as being there together? 25:781–788, 10 2001.

[122] N. Shibata and S. Yamamoto. Regular paper gpgpu-assisted subpixel tracking method for fiducial markers. *Journal of Information Processing*, 22(1), 2013.

[123] Z.D. Siegel and J.W. Kelly. Walking through a virtual environment improves perceived size within and beyond the walked space. *Attention, Perception, & Psychophysics*, 79(1):39–44, Jan 2017.

[124] M. Slater, J. McCarthy, and F.Maringelli. The influence of body movement on subjective presence in virtual environments. *Human Factors*, 40(3):469–477, 1998.

[125] M. Slater, D. Pérez Marcos, H. Ehrsson, and M. Sanchez-Vives. Towards a digital body: the virtual arm illusion. *Frontiers in Human Neuroscience*, 2:6, 2008.

[126] M. Slater, A. Sadagic, M. Usoh, and R. Schroeder. Small-group behavior in a virtual and real environment: A comparative study. *Presence: Teleoper. Virtual Environ.*, 9(1):37–51, February 2000.

[127] M. Slater, B. Spanlang, M.V. Sanchez-Vives, and O. Blanke. First person experience of body transfer in virtual reality. *PLOS ONE*, 5:1–9, 05 2010.

[128] M. Slater and M. Usoh. Body centred interaction in immersive virtual environments. In *Artificial Life and Virtual Reality*, pages 125–148. John Wiley and Sons, 1994.

[129] M. Slater, M. Usoh, S. Benford, D. Snowdon, C. Brown, T. Rodden, G. Smith, and S. Wilbur. Distributed extensible virtual reality laboratory (devrl). In *Virtual Environments and Scientific Visualization '96*, pages 137–148, Vienna, 1996. Springer Vienna.

[130] M. Slater, M. Usoh, and A. Steed. Depth of presence in virtual environments. *Presence: Teleoper. Virtual Environ.*, 3(2):130–144, January 1994.

[131] M. Slater, M. Usoh, and A. Steed. Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Trans. Comput.-Hum. Interact.*, 2(3):201–219, September 1995.

[132] R.S. Sobel and N. Lillith. Determinants of nonstationary personal space invasion. *The Journal of Social Psychology*, 97(1):39–45, 1975.

[133] R. Sommer. *Personal Space: The Behavioral Basis of Design.* Spectrum Books, 1969.

[134] A. Steed, Y. Pan, F. Zisch, and W. Steptoe. The impact of a self-avatar on cognitive load in immersive virtual reality. In *2016 IEEE Virtual Reality (VR)*, pages 67–76, March 2016.

[135] A. Steed and R. Schroeder. Collaboration in immersive and non-immersive virtual environments. pages 263–282, 01 2015.

[136] A. Steed, W. Steptoe, W. Oyekoya, F. Pece, T. Weyrich, J. Kautz, D. Friedman, A. Peer, M. Solazzi, F. Tecchia, M. Bergamasco, and M. Slater. Beaming: An asymmetric telepresence system. *Computer Graphics and Applications, IEEE*, 32(6):10–17, Nov 2012.

[137] F. Steinicke, G. Bruder, K.H. Hinrichs, J. Jerald, H. Frenz, and M. Lappe. Real walking through virtual environments by redirection techniques. *Journal of Virtual Reality and Broadcasting*, 6(2), 2009.

[138] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe. Estimation of detection thresholds for redirected walking techniques. *IEEE Transactions on Visualization and Computer Graphics*, 16(1):17–27, Jan 2010.

[139] W. Steptoe, A. Steed, and M. Slater. Human tails: Ownership and control of extended humanoid avatars. *IEEE Transactions on Visualization and Computer Graphics*, 19(4):583–590, April 2013.

[140] E. A. Suma, Z. Lipps, S. Finkelstein, D. M. Krum, and M. Bolas. Impossible spaces: Maximizing natural walking in virtual environments with self-overlapping architecture. *IEEE Transactions on Visualization and Computer Graphics*, 18(4):555–564, 2012.

[141] E.A. Suma, S. Babu, and L. Hodges. Comparison of travel techniques in a complex, multi-level 3d environment. In *3DUI 2007*, 2007.

[142] Z. Szalavári, D. Schmalstieg, A. Fuhrmann, and M. Gervautz. “studierstube”: An environment for collaboration in augmented reality. *Virtual Reality*, 3(1):37–48, 1998.

[143] W.B. Thompson, P. Willemsen, A.A. Gooch, S.H. Creem-Regehr, J.M. Loomis, and A.C. Beall. Does the quality of the computer graphics matter when judging distances in visually immersive environments. *Presence: Teleoper. Virtual Environ.*, 13(5):560–571, October 2004.

[144] H. Tramberend. Avocado: A distributed virtual reality framework. In *Virtual Reality, 1999. Proceedings., IEEE*, pages 14–21. IEEE, 1999.

[145] J.G. Tromp, A. Steed, and J.R. Wilson. Systematic usability evaluation and design issues for collaborative virtual environments. *Presence: Teleoper. Virtual Environ.*, 12(3):241–267, June 2003.

[146] L.A. Vallis and B.J. McFadyen. Locomotor adjustments for circumvention of an obstacle in the travel path. *Experimental Brain Research*, 152(3):409–414, Oct 2003.

[147] C. Vassalo, A.-H. Olivier, P. Soueres, A. Cretual, O. Strasse, and J. Pettre. How do walkers avoid a mobile robot crossing their way? *Gait and Posture*, 51:97 – 103, 2017.

[148] K. Vasylevska and H. Kaufmann. Influence of vertical navigation metaphors on presence. In *In Challenging Presence - Proceedings of 15th International Conference on Presence (ISPR 2014)*, pages 205–212, 2014.

[149] K. Vasylevska and H. Kaufmann. Influence of path complexity on spatial overlap perception in virtual environments. In *Proceedings of International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*, pages 1–8. The Eurographics Association, 2015.

[150] K. Vasylevska and H. Kaufmann. Towards efficient spatial compression in self-overlapping virtual environments. In *Proceedings of IEEE Symposium on 3D User Interfaces*, pages 1–10. IEEE, 2017.

[151] K. Vasylevska, H. Kaufmann, M. Bolas, and E.A. Suma. Flexible spaces: Dynamic layout generation for infinite walking in virtual environments. In *3D User Interfaces (3DUI), 2013 IEEE Symposium on*, pages 39–42, March 2013.

[152] D. Waller, E. Bachmann, E. Hodgson, and A. Beall. The hive: A huge immersive virtual environment for research in spatial cognition. *Behavior Research Methods*, 39(4):835–843, 2007.

[153] T. Waltemate, D. Gall, D. Roth, M. Botsch, and M. E. Latoschik. The impact of avatar personalization and immersion on virtual body ownership, presence, and emotional response. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1643–1652, April 2018.

[154] R.C. Waters, D.B. Anderson, John W Barrus, David C Brogan, M.A. Casey, S.G. McKeown, T. Nitta, I.B. Sterns, and W.S. Yerazunis. Diamond park and spline: Social virtual reality with 3d animation, spoken interaction, and runtime extendability. *Presence: Teleoperators and Virtual Environments*, 6(4):461–481, 1997.

[155] C. Wienrich, K. Schindler, N. Dollinger, S. Kock, and O. Traupe. Social presence and cooperation in large-scale multi-user virtual reality – the relevance of social interdependence for location-based environments, 2018.

[156] L. Wilcox, R S. Allison, S. Elfassy, and C. Grelik. Personal space in virtual reality. 3:412–428, 10 2006.

[157] P. Willemsen, M.B. Colton, S.H. Creem-Regehr, and W.B. Thompson. The effects of head-mounted display mechanical properties and field of view on distance judgments in virtual environments. *ACM Trans. Appl. Percept.*, 6(2):8:1–8:14, March 2009.

[158] B. Witmer and M. Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoper. Virtual Environ.*, 7(3):225–240, 1998.

[159] B.G. Witmer and P.B. Kline. Judging perceived and traversed distance in virtual environments. *Presence: Teleoper. Virtual Environ.*, 7(2):144–167, April 1998.

[160] M. Wolff. *Notes on the behaviour of pedestrians. People in places: The sociology of the familiar.* 1973.

[161] Y. Yuan and A. Steed. Is the rubber hand illusion induced by immersive virtual reality? In *2010 IEEE Virtual Reality Conference (VR)*, pages 95–102, March 2010.

[162] M. Zank and A. Kunz. Optimized graph extraction and locomotion prediction for redirected walking. In *2017 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 120–129, March 2017.

[163] Z. Zhang. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(11):1330–1334, November 2000.

[164] S. Zhao. Toward a taxonomy of copresence. *Presence: Teleoper. Virtual Environ.*, 12(5):445–455, October 2003.

[165] K. Zibrek, E. Kokkinara, and R. Mcdonnell. The effect of realistic appearance of virtual characters in immersive environments - does the characters personality play a role? *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1681–1690, April 2018.

List of Figures

List of Tables