

Come Look at This: Supporting Fluent Transitions between Tightly and Loosely Coupled Collaboration in Social Virtual Reality

Pauline Bimberg, Daniel Zielasko *IEEE*, Benjamin Weyers *IEEE*, Bernd Froehlich *IEEE*, Tim Weissker



Abstract—Collaborative work in social virtual reality often requires an interplay of loosely coupled collaboration from different virtual locations and tightly coupled face-to-face collaboration. Without appropriate system mediation, however, transitioning between these phases requires high navigation and coordination efforts. In this paper, we present an interaction system that allows collaborators in virtual reality to seamlessly switch between different collaboration models known from related work. To this end, we present collaborators with functionalities that let them work on individual sub-tasks in different virtual locations, consult each other using asymmetric interaction patterns while keeping their current location, and temporarily or permanently join each other for face-to-face interaction. We evaluated our methods in a user study with 32 participants working in teams of two. Our quantitative results indicate that delegating the target selection process for a long-distance teleport significantly improves placement accuracy and decreases task load within the team. Our qualitative user feedback shows that our system can be applied to support flexible collaboration. In addition, the proposed interaction sequence received positive evaluations from teams with varying VR experiences.

Index Terms—Virtual Reality, 3D User Interfaces, Multi-User Environments, Social VR, Groupwork, Collaborative Interfaces

1 INTRODUCTION

VIRTUAL Reality (VR) allows users to overcome the limitations of the real world, for example, by instantly changing their position within the virtual environment independent of the distance covered by that change. In theory, this spatial flexibility should enable group members to collaborate, even while mainly working in different virtual locations. This is especially interesting in situations where users discover something a user at another location might be interested in, have some advice on, or could actively help with. However, current VR applications do not seem to equip their users with the tools that are needed to utilize this potential for effective group work by including frequent switches between loosely coupled collaboration, where users work towards a shared goal without directly

interacting or needing to be near each other, and tightly-coupled collaboration, which is more closely coordinated and typically happens in the same place [1]–[3]. Instead, similar to real-world situations, users have to either stay together at all times to be able to consult each other or be willing to completely give up on their current activity and travel to another user’s location for face-to-face collaboration. The work presented in this paper addresses this problem by providing users with tools to support seamless transitions between different levels of tightly to loosely coupled collaboration within the virtual space and to return to their previous working context when an exchange has ended.

In order to achieve this, the paper presents a theoretical framework detailing different consecutive phases of cooperation by teams that are distributed in a virtual space and the design choices that have to be made when implementing these phases. Specifically, we suggest connecting the states of loosely coupled collaboration from different places in the virtual environment and tightly coupled collaboration in the same environment through three additional collaboration phases. In the overview phase, users are enabled to gain an overview of collaborators’ positions and potential activities and states in the shared virtual space. The check-in phase lets users consult each other through asymmetric collaboration metaphors without either person leaving their current position. A visit marks a transfer of one user to the position of their collaborator, which can be concluded by returning to their position of origin or extended by joining their partner for a longer phase of face-to-face collaboration. To illustrate how this framework can be applied in a real task scenario to support collaboration across large virtual distances, we have developed an example implementation that we evaluated in a user study with 32 participants in teams of two. Our overarching research question asked how our interface choices aligned with user expectations and if our suggested framework supported collaborating teams sufficiently. Since the change between working in separate places and face-to-face collaboration is a central aspect of our framework, we also asked the more specific research question of which effects differently designed long-distance teleportation mechanisms have on user performance, understanding, and well-being. To approach this more constrained topic, we have realized two alternative

Pauline Bimberg, Daniel Zielasko, and Benjamin Weyers are with Trier University, Germany. Bernd Froehlich is with Bauhaus-Universität Weimar, Germany. Tim Weissker is with RWTH Aachen University, Germany. E-Mail: bimberg@uni-trier.de, daniel.zielasko@rwth-aachen.de, weyers@uni-trier.de, bernd.froehlich@uni-weimar.de, me@tim-weissker.de
Manuscript received June 20, 2023; revised March 11, 2024; accepted May 30, 2024.

long-distance teleportation methods within our framework implementation, in which users can either position themselves via a world in miniature (WIM) metaphor or be positioned by their partner at the target location. Overall, our work makes the following contributions:

The conceptual introduction of different phases of collaboration and their connection in a theoretical framework, including the design decisions that have to be made for each phase.

An example implementation of the suggested phases of collaboration and their connection within one fluid interaction sequence.

The development of two target selection techniques for long-distance teleportation to the position of a collaborator.

Results of a comparative user study between two different target-planning mechanisms showing that distributing control for the preparation of a long-distance teleport significantly improves placement accuracy and decreases task load within the team.

Results of an exemplary framework evaluation indicating that the suggested stages of collaboration can improve spatial flexibility during distributed teamwork of user pairs with different experience levels.

Overall, our results encourage the support of flexible forms of collaboration in social virtual reality systems as well as the use of our suggested framework to conceptualize, discuss, and evaluate these support mechanisms.

2 RELATED WORK

Although social virtual reality spaces are gaining a larger audience for both socialization [4]–[6] and domain-specific collaboration [7]–[11], research into the design of virtual environments for collaboration predates this development by decades [5], [12]–[14]. Ongoing technical developments allowed the main focus of research to shift from solving the technical challenge of connecting multiple users within one virtual space [13], [14] towards the design of interaction techniques for fluent and intuitive collaboration [12]. In the following sections, we will discuss approaches toward supporting virtual collaboration in related work, starting with a short overview of different types of collaboration in Section 2.1. We will then specifically discuss two prominent aspects if multi-user collaboration in the literature that particularly influenced our work: the use of multiple viewpoints during multi-user interaction (see Section 2.2) and the role that navigation plays in the support of groups within large virtual spaces (see Section 2.3).

2.1 Types of Collaboration

Computer-Supported Collaborative Work (CSCW) can be split up with respect to the axes of space (users being either in the same place or in different places during collaboration) and time (users working on a problem simultaneously or at different times) [12], [15]. In recent years, most new collaborative applications in extended reality have been developed for users who enter a virtual environment from different real-world places at the same time [12]. Such a synchronous collaboration in shared virtual environments

can be symmetric or asymmetric, both when looking at the devices used by the collaborating parties and when looking at the roles and responsibilities that each partner carries [16], [17]. Fully symmetric setups supply users with comparable knowledge in the targeted domain with identical hardware and interaction possibilities. Examples of asymmetric collaboration include a user wearing an Augmented Reality device interacting with a user using a VR device [18]–[20]. Similarly, a user wearing a VR Device might be guided by a user surveying the environment from a different perspective on a desktop screen [21]. Finally, collaboration can represent a tight or loose coupling of activities between collaborating users or parties [22]. In tightly-coupled collaboration, users are working together directly, while loosely-coupled collaboration has users split up and work on different aspects of a shared goal on their own or in sub-groups. Often, collaborative tasks are made up of phases of both tightly- and loosely-coupled collaboration and collaborative systems need to support both forms of work as well as transitioning between them [23]. Focusing on the prevalent use case of synchronous collaboration between partners occupying different workspaces in the real world, our work aims to support VR users in the fluid transition between different phases of tightly- and loosely-coupled collaboration.

2.2 Multiple Views for Collaboration

While the most popular examples of using different views during collaboration can be found in asymmetric collaboration scenarios, providing all users with a selection of different perspectives onto a common working environment can also be useful for partners with symmetric capabilities. This idea of using multiple simultaneous views during collaboration has long been a central concept for cooperative work using traditional 2D screens [24]. Depending on the task, this can be done by equipping a team with multiple viewing surfaces or by splitting a shared surface into different zones [24]–[26]. For a group search task in a 3D environment presented on a desktop screen, Dodds and Ruddle equipped users with multiple simultaneous views of the environment and let users teleport to a group member’s position when needed. In their study, these options increased both the amount of communication within the group and the distance that group members collaborated over [27], [28].

In immersive virtual reality, examples of the use of multiple views without the static distribution of user roles can be found in the work of Kunert et al., where users could open smaller Photoportals for individual exploration within a shared large 3D display [29]. Building on this approach, the authors later gave teams of co-located users simultaneous access to a large 3D display wall showing a first-person view of the environment and a smaller 3D tabletop display showing a bird’s eye view of the environment [22]. In their user study, they were able to show advantages of this multi-screen approach concerning user comfort, presence, and performance in a shared search task. Similarly, Schroeder et al [30] later introduced an analysis framework for what they describe as *transitional interfaces*. In their example use case, two collaborators can each freely change between viewing the shared working environment via an HMD, a tablet and a desktop setup.

Seeing that the utilization of exible viewpoints for different tasks in a shared virtual environment is more prevalent in desktop systems, we believe that there is still unexplored potential for the support of such mechanisms in immersive VR. In our work, we suggest providing users with different views onto the virtual environment to support different collaborative interactions. In this regard, we present our users with ways to overlook the whole collaborative environment, see more detailed representations of a collaborator's location and share the viewpoint of another user onto their surroundings in the form of portal surfaces.

2.3 Support of Collaboration in Navigation

Another option for users to gain different perspectives of a virtual environment is navigation. In distributed setups, single-user navigation techniques are still the most prevalent [31]. Popular examples of virtual travel techniques used to bridge short distances include steering methods, where users continuously specify the speed and direction of their movement [32]–[34], and teleportation in vista space, where users specify a target position and/or rotation directly in the virtual environment and are transported to this point instantaneously [32], [35], [36]. For longer distances, users of virtual environments can either select a new target for themselves using smaller proxies of the virtual environment as it is done in the World in Miniature and Voodoo Doll metaphors [37], [38], choose a target through the positioning of a portal [29], or simply trigger a transition to a predefined goal position [6]. The option to travel through the virtual environment together using techniques that move the whole group as a unit has mainly been suggested for co-located setups [31], [39], [40]. Nevertheless, the combined use of single-user navigation and group navigation metaphors has shown promising results in related work. Examples include the work presented by Weissker et al. on the mechanisms of forming and leaving groups for the joint navigation of distributed teams. Through these mechanisms, the authors extended their multi-ray jumping technique for co-located users and combined it with single-user navigation capabilities [41], [42]. Results showed that this enabled users to split up for phases of individual exploration before joining together as a group to present their findings to their collaborator [42]. A similar approach was also presented by Beck et al., where two user groups using separate large 3D displays could temporarily join together for discussions or joint exploration (tightly coupled collaboration) and split up for individual activities (loosely coupled collaboration) [43]. A different approach towards combining joint and individual travel is the addition of long-distance group teleports to spaces that mainly offer single-user navigation techniques. Examples of this include the use of Photoportals in the work of Kunert et al. to take user groups to a position suggested by one group member and the use of party portals in the AltSpaceVR application, which allow groups of users to change between spaces dedicated to different activities together [6], [29].

Fig. 1. An overview of the different phases of collaboration discussed in this paper and their connection proposed as part of our interaction framework. In most other collaborative VR systems, users have to solely rely on individual navigation to change from loosely- to tightly-coupled collaboration and vice versa.

3 STAGES OF COLLABORATION IN DISTRIBUTED VIRTUAL TEAMS

Based on collaboration methods suggested by related publications in the realm of CSCW, we have derived different stages of collaboration relevant to different types of collaborative work performed by distributed users. The proposed model builds upon the trade-off between staying at one's own current position in the virtual environment, and leaving it behind to join another user by adding different steps between these two extreme cases of loosely- and tightly-coupled collaboration. Specially, users can gain a general impression of the virtual space in the Overview stage (see Section 3.1), consult each other on tasks in the Check-In stage (see Section 3.2), join together temporarily for virtual face-to-face collaboration in the Visit stage (see Section 4.3), and finally have the option to Return to their previous working context or Join their partner for a longer phase of closely-coupled collaboration (see Section 3.4). The interplay and transitions between the proposed stages are visualized in Figure 1 and will be described in more detail in the following sections. Our theoretical model, as it is discussed here, is not meant to be an explicit guide for one correct implementation of each phase of collaboration, but it can serve as a basis for many but it can serve as a basis for the systematic design and evaluation of different collaborative interaction techniques. One exemplary implementation of this framework, which was also tested in our reported user studies later in the paper, will be described in Section 4. In the following descriptions, we will refer to the user giving up their current position in the virtual environment to collaborate in a different location as the visiting user and call the user who remains at their current location for the collaboration the hosting user.

3.1 Overview

The overview phase allows users to gain an overall impression of the layout of the virtual environment as well as the positions of other users, which has been a relevant feature in previous 2D interaction systems [27], [28]. This enables users to update their broad awareness of interactions happening within the environment, to find points of interest where other users are gathered, and to decide whether they want to further engage with any of them. It has been shown that equipping users with an overview map of their virtual environment can partially counteract the

Fig. 2. The Overview stage supports the user in gaining an overall impression of the environment in which they interact with their group members. For this, a suitable representation of the current environment as well as the other users within the environment has to be chosen.

loss of spatial orientation that is often introduced by long-distance teleportation [44]–[46]. The concrete design decisions required to implement the overview phase consist of choosing an appropriate representation of the environment in which the collaboration takes place as well as suitable user representations (see Figure 2).

3.1.1 Environment Representation

The chosen representation of the virtual environment highly depends on the use case. For smaller environments, a map or world-in-miniature view seems to be the most intuitive way to represent all relevant features of a virtual space together with the distribution of other users [37], [45], with potential design additions having to be made to account for necessary scaling [47], occlusion [48] or multi-level environments [49]. However, collaboration can also happen in contexts where a smaller-scale representation of a whole virtual scene does not seem feasible. Larger worlds might be hard to represent at a scale where they can be looked over entirely without losing all features of interest. Here, it might be more feasible to use a varying scale for landmarks or places that are currently occupied by other users by relying on techniques known from the field of information visualization such as Overview + Detail or Focus + Context [50]. Collaboration might also happen between environments that are entirely unconnected, places where the spatial connection is not relevant to the collaboration, or places where spatial connections cannot be represented within a map. Examples of such spaces include separate virtual spaces for activities, different museums on the same topic, or different versions of the same building design [4], [7]. Here, other approaches to overview representation like graph views [51], [52] could be explored.

3.1.2 User Representation

Users should be visualized in a way that makes them easily visible within the context of the chosen environment representation. In larger worlds, showing full user representations to size within the chosen environment representation would lead to them being easy to miss or even hard to perceive when found. Here, simplified representations of the users' avatars that communicate their (relative) position and direction present an alternative. As social formations can usually be interpreted from a distance [53], [54], this information should be sufficient to enable the viewer to gain an impression of current groupings and their level of engagement. Similar to mini-maps used in video games, additional information like a user's current talking status,

explicitly formed groups, or availability for engagement could also be presented with the user's position marker [55].

3.2 Check-In

Fig. 3. Design decisions to be made in the Check-In phase include the choice of a forming mechanism and the selection of features used to support collaboration.

If the overview alerts users to an interesting activity or if they have found something in their working context that they would like to discuss, they can start a check-in with another user. The check-in stage represents an asymmetric, remote collaboration scenario, where the user checking in receives information about the actions and direct environment of their collaborator(s) without leaving their own position in the virtual environment. In essence, the check-in phase is most similar to asymmetric collaborations between a remote supporter and an on-site worker, as described in Section 2.1. Since the check-in phase changes the nature of the collaboration from a general awareness of every user's actions to a direct exchange with one specific user or user group, a check-in also represents the start of an explicit collaborative exchange. Therefore, both the design of a check-in itself as well as the mechanisms needed to engage in a check-in are design decisions that have to be made depending on the intended usage context (see Figure 3).

3.2.1 Forming Mechanism

Similar to the group forming process in multi-user navigation [42], the mechanisms chosen for check-in should be representative of the working relationship between the included parties. Users could be able to start a check-in automatically, only when the other user has given general permission to be engaged with or only with explicit confirmation by both parties. Here, two collaborators of similar rank might want to rely on a system of mutual confirmation while a teacher overlooking several groups of working students might be given the right to check in at any point. When groups enter a check-in with each other, a full confirmation of every included party might also be bothersome such that a need for other agreement mechanisms might arise.

3.2.2 Collaboration during a Check-in

Since the check-in phase should give visiting users the ability to relate to the hosting user's current surroundings, the representation of the whole environment in the overview phase should be swapped for a more concentrated representation of the engaged user at this point in the collaboration. Similar to measures seen in AR/VR collaboration, this might happen by enlarging the previous overview representation

of the virtual environment such that additional details are visible [47], [56] or by sharing the view of the hosting user with the visiting user to be able to discuss the hosting user's activities from the same point of view [57]. As check-ins require direct communication between the involved users, this would also be the point where a voice connection between users becomes more relevant and user representations might become more detailed to include gestures or even expressions, depending on what is offered by the technology in use. In addition, context information from the visiting user like their field of view [58], or additional ways of communication, like letting them place hints in the environment of their partner [49], might make the interaction more fluid. Depending on the intended interaction, the communication features can be chosen symmetrically, with both users receiving information about each other's environment, or asymmetrically, setting the focus on one user's surroundings that are to be discussed [56].

3.3 Visit

Fig. 4. To realize the Visit stage, the four phases of teleportation have to be implemented and communicated within both the visiting and hosting user's environments.

At the point where two users feel the need to collaborate at a level of directness that exceeds the interaction in the check-in phase, a visiting user could decide to temporarily travel to the hosting user's location. In contrast to the overview and check-in stages which let users gain insights about their team members' activities while remaining at their current location, the visit phase marks the change from virtually remote to virtually co-located interaction. This change requires the visiting user to leave their current position and activities completely for the duration of the face-to-face interaction. One of the aims of our framework is the fluid change between loosely coupled collaboration in two different places and tightly coupled collaboration in the same place. The most efficient way of performing this change is using a target-based travel technique or teleportation [32]. Depending on the environment, this can, of course, be complemented by other travel techniques offered for short-distance travel for activities where users are working near each other or if traversing the route between the two places in question is of importance. Similar to the framework proposed by Weissker et al. for short-distance teleportation [36], our process of teleporting a user to the location of a team member can be split into the phases of target selection, pre-travel information, transition to the target, and post-travel feedback, with our process not only involving the visiting user but also integrating the hosting user who is already at the goal position.

3.3.1 Target Selection

While the general target area of the teleport for collaboration purposes is defined by the hosting user's position, the concrete point at which the visiting users should arrive still needs to be selected to avoid unpleasant overlaps with scene objects or avatars. Conceptually, this selection could be performed by either the visiting user using the information given in the check-in phase, the hosting user in their immediate surroundings, or automatically by the system based on predictions of user intentions and available space. Apart from the selection responsibility, it can also be decided what input method is used for the selection and if the specified target should define a user's position or include their orientation.

3.3.2 Pre-Travel Information

The information of where a visiting user is supposed to arrive in the virtual environment is not only relevant to the visiting user themselves but also to the hosting user, since a sudden appearance could lead to confusion or even colliding movements in the moment of travel. Therefore, pre-travel information about the user's future position and orientation needs to be conveyed to the visiting user at their current position, for example by centering the information used for the check-in above their new position, and to the hosting user at the goal position.

3.3.3 Transition

The change between two different positions in a virtual environment can be instantaneous or rely on a transition period, which might either animate the change between the two positions or require active user interaction like the use of a portal to be walked through or a preview sphere that is brought to the user's face [59]. As usual, this transition should be chosen with the overall use case in mind, considering a trade-off between the user experience of the chosen transition and efficiency [60]. At the same time, the hosting user at the goal position should be informed about the changing status. If a longer transition is chosen during which the user is not aware of changes being made at the target position, this should also be communicated to the collaborating users.

3.3.4 Post-Travel Feedback

While users should be prepared for their position change by the given pre-travel information, additional post-travel feedback at the goal position can be used to help the user orient themselves in relation to the other present parties. This might be especially relevant if the target orientation was not pre-planned to fit the cause of the visit or if a longer animation in the travel phase might lead to changes in the environment being made between the visiting user receiving pre-travel information and them arriving at the target. Additionally, since visits can be made with the intention to later return to the previous location, post-travel feedback may also be needed at a user's potential point of return. Depending on the use case, post-travel feedback might represent the user's return position, their current status or position in the environment, or even information about what they were working on before starting a visit to another position.

3.4 Return or Join

Fig. 5. To give the visiting user a choice between joining the host permanently or returning to their previous working context, corresponding mechanisms have to be selected for each choice.

After engaging in a short visit to the position of another user, a visitor might want to return to their previous working context. They might, however, also discover that their new environment is interesting enough for an extended phase of collaboration or be simply done with their work at their previous location. In this case, they could decide to completely join the other user in their given environment or stay with them for a longer period of time (see Figure 5).

3.4.1 Return

While a target selection phase is not necessary for returning to one's original position, users should again receive pre-travel information to see the goal of the teleportation step, undergo a transition upon confirming the travel, and receive post-travel feedback at their goal position.

3.4.2 Join

Since a user's previous position should still be displayed as post-travel feedback to other people at this location, the choice to not travel back to the position of origin should also be made explicitly to update others and remove unnecessary information at a user's origin. When a user is currently on a visit to one of their collaborators, it is a sensible choice to display a marker at their previous position as post-travel feedback to communicate that they might still return. However, if the user permanently joins their collaborator or plans to remain with them for a longer time, this form of additional communication must be removed or their stay must be reflected in their marker. On one hand, this avoids cluttering the collaborative environment with indicators for absent users, on the other hand, it avoids the confusion of third parties who might otherwise wait for a user's expected return. Depending on the use case in question, the mechanism to permanently join or communicate a longer absence could be done explicitly by the user or implicitly by system mechanisms. Examples of this could be prompting a user to join when they want to interact with scene objects at the target position when they exceed a certain movement radius, or when a certain amount of time has been spent at the target.

4 SYSTEM IMPLEMENTATION

Based on our suggested phases of collaboration described in Section 3, we have followed an iterative design process to develop an interaction system that allows distributed VR users to enter and seamlessly switch between these phases (see Figure 6), drawing inspiration from previous

VR interaction concepts. Our system implementation is built as a prototype to explore the usefulness of the different types of collaboration as well as the idea of their user connection in one interaction sequence. In addition, we were especially interested in the possibilities that virtually distributed collaborators bring to long-distance teleportation since the ability to rapidly change between positions of collaborators is a central idea that is required independent of the concrete application scenario. This led us to choose exemplary features for all stages while putting a particular focus on two alternative methods for the transition between the Check-In and Visit stages.

For simplicity, our scenario of this paper takes place in one connected environment with a traversable area of 125m*125m that can be fully explored with conventional navigation methods. A full image of the outdoor city environment used in our implementation can be seen in Figure 9. In this setting, our system serves as a complement to facilitate collaboration over longer distances in the virtual environment. In addition, we have built the interactions without any assumed hierarchies between the partners.

In our implementation, each user was equipped with one HTC Vive Pro headset and one controller. The single-user navigation technique offered in our system consists of a simple short-distance teleportation technique activated by the trigger button [32], [35], [36] and added snap-turns of 30 degrees around the user's own axis activated by pressing down on the left and right quadrant of the touchpad [61]. A pointing ray for communication and selection could be activated with the bottom quadrant of the touchpad.

In the following sections, we will explain the implementation of each proposed stage in our system.

4.1 Overview

Based on our initial choice of environment, the world can be represented by a world-in-miniature (WIM) with one height level [37]. To make the user position within the WIM salient, we have chosen an abstract user representation that has been increased in size for easier identification. The abstract avatars communicate the position and orientation of the user through placement and shape as well as the users' identities through their color. The scale factor of the overview-WIM in our implementation was 0.01, which led to a WIM size of around 1.2 m.

In our implementation, the overview-WIM can be opened with the controller's menu button (see Figure 6, left). If users decide that they want to further engage with another teammate or if they have entered the overview phase with the intention to consult another user, they can request a check-in with this user by selecting their avatar within the miniature representation of the virtual environment using a raycast. Selecting their own avatar signals that they are in need of assistance and should be contacted by one of the other users, who in turn are prompted to open their WIM to start further interaction. As an alternative to the selection on the WIM, users can also be provided with virtual buttons attached to the WIM in order to request check-ins or ask for assistance.

Fig. 6. Our implementation of the different phases of user collaboration in spatially distributed virtual reality. From left to right: A user gets an overview of the virtual environment in the form of a world in miniature (WIM), checks in with their partner, receives a placement suggestion for a long-distance teleport, and joins their partner for face-to-face collaboration.

4.2 Check-In

In our implementation, the forming mechanism for the check-in requires the consent of the other user. Using the metaphor of receiving a phone call, a user who is requested for a check-in will be alerted by their controller rumbling and can decide to accept or refuse the request using controller buttons. On acceptance of the check-in, the WIM seen by the visiting user will zoom in while simultaneously being centered around the partner's position (see Figure 6, center-left). The zoom factor of the check-in WIM in our implementation was set to 7, resulting in an overall scale factor of 0.07. After this transition, the visiting partner will see a miniature representation of the hosting user's avatar as well as all of their surroundings within a 5m distance in the environment, which corresponds to a radius of 35cm on the newly scaled WIM representation. They will also be able to see a 2D first-person view of their partner's current viewing perspective in the virtual space. This selection of features represents two communication methods often seen in related work on asymmetric collaboration, where virtual clones of the user's environment, as well as viewpoint sharing, are often utilized to assist teams [19], [43], [57], allowing us to explore the use of both viewing modes in our study. Should users wish to fully leave their current location based on what they saw in the preview, they can join their partner for face-to-face interaction. For this, they have to prepare a visit to the other user by either selecting a target location on the WIM or requesting their partner to suggest a target for them in the environment. Otherwise, they can simply return to their previous work by ending the call.

4.3 Visit

A visit to another user's position requires a long-distance teleport by the visiting user. In the following, we will describe our chosen implementation for each of the phases of teleportation.

4.3.1 Target Selection

In our work, we focus on a comparison of two user-driven approaches for target selection, which we call WIM placement and local placement, respectively. In summary, these approaches differ in the collaborator responsible for selecting the target position (visiting versus hosting user) and the medium that is used for accomplishing this task (WIM

versus immediate surroundings). This introduces a trade-off between a potential heightened placement accuracy for local placement due to the increased size of the target environment compared to WIM placement and a potential overhead introduced by distributing control between two collaborators. Our study described in Section 5 therefore put an explicit emphasis on a more detailed comparison of these two placement methods. For a valid comparison, both techniques used similar specification work flows inspired by the anchored teleportation method by Bimberg et al. [61]. The operating user first selects a target position within the corresponding medium using a parabolic pick ray, locks this position by button press, and then moves the pick ray away from the locked position into the desired viewing direction.

During a check-in, users can choose their preferred placement technique by either selecting the floor of the virtual environment to place themselves or the avatar of the other user to ask them to suggest a placement. In addition, the two placement techniques can be activated by utilizing corresponding virtual buttons that can be shown on the side of the preview window. Any currently activated placement technique (either WIM or local placement) is operated using the bottom quadrant of the touchpad. Here, the placement ray is activated by touching the assigned area, the selected position is confirmed by pressing down on the button and the subsequent rotation specification is confirmed when letting go of the touchpad.

4.3.2 Pre-Travel Information

While specifying a teleportation target using the parabolic pick ray with either technique, a ghost avatar is displayed to communicate the currently selected position and orientation. For improved mutual comprehensibility, a synchronized copy of this avatar is presented both on the WIM and in the real environment (see Figure 7). The idea of suggesting a target position by placing a copy of another user's avatar has already been used in the collaborative design environment described by Xia et al., but it was never evaluated explicitly [7]. In our work, the ghost avatar in the WIM is used together with a preview window, communicating the user's future view onto the target environment (see Figure 6, center-right).

4.3.3 Transition

The transition to the target position is done once the visiting user confirms their travel destination by selecting the preview window with their controller or selecting a confirmation button next to the window. Since the long-distance teleport represents a complete cut from the previous point of view, we decided on a fade-to-black transition to prevent unpleasant feelings introduced by abrupt viewpoint changes [59].

4.3.4 Post-Travel Feedback

Since the visiting user will be able to engage with their local partner upon arrival, we decided against adding explicit post-travel feedback at the target position. We do, however, leave a ghost avatar of the visiting user at their original position and on the WIM representation to communicate to other users in the scene that they could return there after their visit.

4.4 Join or Return

In our implementation, the return option can be activated by opening the overview WIM again and selecting the representation of the ghost avatar at the previous location. Users are then shown a preview window of their point of view at the previous position and are teleported back on confirmation. Instead of returning after a visit, users can also join their partner at the target location permanently and resume working from there. This might happen because their work at the previous location is done or because the task calls for more extensive joint collaboration. For the joining mechanism, we have implemented an automatic, distance-based approach, where the ghost avatars and therefore the return options are deactivated once the corresponding user leaves a marked radius of 10m around the position that they were teleported to.

5 COMPARISON OF WIM PLACEMENT AND LOCAL PLACEMENT

In order to evaluate our proposed stages of distributed group work within a collaborative setting, we invited 32 participants, who were asked to sign up in groups of two. They were asked to complete two distinct tasks, which were developed with a focus on different parts of our system. The first task, discussed in this section, focused on the comparison of two different long-distance teleport metaphors. The second task, discussed in section 6 is focused on an exemplary evaluation of all implemented phases of collaboration.

During all phases of the study, users were divided into two different tracking spaces located in the same room and could thus talk to each other directly at all times. This decision had to be made in order to be able to keep hygiene and safety protocols in place with only one experimenter present, and we did not feel like the resulting spatial audio mismatch during conversations would negatively affect our abilities to address our research questions. During the study, each user was equipped with an HTC Vive Pro headset, which was connected to a dedicated workstation running the Unity application through a wireless module, and one pro controller. Over the course of the study, participants

Fig. 7. Procedure diagram of a single trial in the first study task. Teams start in the check-in phase and have to select a pre-defined target with either the WIM or the local placement technique. Upon confirmation, a teleport to the selected target is triggered, where users are asked to solve a spatial orientation task – first without and then with seeing the surroundings.

where always able to hear each other clearly without having to raise their voices.

All study tasks took place in the same virtual environment, which can be seen in Figure 9. To match the low-poly environment and to keep the avatars inclusive without having to add a character customization phase to the study, users were represented as simple avatars consisting of a robot head, shirt, and tracked controllers. Users were assigned either the orange or the blue version of this robot avatar. The assignment of corresponding colors within the team was done randomly and defined the sub-tasks each user was responsible for during the study procedure described in Section 5.2.

5.1 Study Design

Our first study task was developed to compare the use of the two placement variants described in Section 4.3. Our comparison is based on the discussed trade-off between the added organizational overhead and the potential increase in performance. In addition, we wanted to evaluate the suitability of our chosen travel preview mechanisms for long-distance teleportation. This leads to the following specific questions to investigate in detail:

RQ1.1 What are the advantages and disadvantages of the two target selection methods for preparing a long-distance teleportation to a collaborator?

RQ1.2 How does the choice of placement method influence the task load that the visiting and hosting user experience?

RQ1.3 Do the suggested travel preview mechanisms and preview features enable users to orient themselves at their new position after a long-distance teleport?

5.1.1 Task Design

To study the use of the two placement techniques in isolation, we omitted the overviewjoin and return features in the first study task, which leaves the features of the check-in and visit phase to be investigated in detail. Thus, each sub-task during the study started in the check-in phase, was followed by a target selection and teleportation, and was concluded by a spatial orientation task at the teleport target. A general task item consisted of two location markers, one where

the hosting user was spawned and one where the visiting user should teleport to (see Figure 7). The markers were arranged in one of four selected formations that are typical for joint work, namely side-by-side, face-to-face, facing each other across an object of interest, and L-shaped [53]. In addition, each target location contained a pink amingo geometry that was used for the visitor's spatial orientation task. Overall, there was one set of four tasks used for the tutorial and four sets of eight tasks used during the recorded trials. Each of the task sets contained the same number of each target formation. In more detail, the task phases, which can be seen in Figure 7, looked as follows:

1. Check-In Phase: Both users are automatically placed at their starting positions. A WIM in the check-in state is opened directly in front of the visitor, who then is asked to position themselves in such a way that they can see the environment and the pointing target well and start the target selection phase.

2a. WIM Placement: The visitor uses the target selection technique on the WIM to place their miniature avatar at the target position and direction as indicated by circular geometries in the WIM.

2b. Local Placement: The hosting user gets alerted via controller vibration that the visitor has started the target selection process. They then place a ghost avatar of their partner into the suggested target position and rotation to suggest a teleportation target.

3. Target Confirmation: When the visitor is satisfied with the placement and has prepared for the orientation task, they can confirm it, which will automatically trigger the next task phase.

4. Orientation Task Blind: After the visiting user triggers the teleport, both users are transported into a featureless room. There, the visiting user is asked to point in the direction where they expect the previously seen amingo geometry to appear after the teleport to test their spatial awareness of the target environment. At the same time, the hosting user is asked to point in the direction where the visiting user will appear after teleporting to test their awareness of the visitor's future location. Blindly pointing at the estimated location of a target object is a common task to evaluate spatial orientation [62]–[64].

5. Orientation Task Seeing: Directly after logging their pointing direction via the trigger button, both users are placed back in the virtual environment, with the visiting user having arrived at the previously selected target. There, they repeat the previous pointing task while seeing their goal objects to provide a baseline. After both users have logged their pointing direction in the seeing orientation task, the next task item gets activated.

5.1.2 Measures

During each trial, we recorded the placement and rotation specification errors made during each step as well as the users' pointing errors in the orientation tasks. After completing all tasks with one of the techniques, users were asked to answer a questionnaire. Here, their well-being was measured using a single-item discomfort score ("On a scale from 0-10, 0 being how you felt coming in, 10 is that

you want to stop, where are you now?") [65], [66]. After this, users filled in the Raw TLX to report the task load they experienced when fulfilling the role of the visiting and hosting user, respectively [67], [68]. After having completed both conditions, users were additionally asked to choose a favorite out of the two techniques that they had tested and to explain their preference.

5.1.3 Hypotheses

According to our intentions when designing both placement techniques as well as the research questions named at the beginning of this section, we have derived hypotheses concerning the results of our user study. We expect the local placement to be easier for the users since they are operating the placement techniques at 1:1 as opposed to a miniature scale. We, therefore, formulated the following hypotheses concerning the user performance with each placement technique, which was the focus of RQ1.1:

H1.1 Users' position selection will be more precise in the local than in the WIM placement technique.

H1.2 Users' rotation selection will be more precise in the local than in the WIM placement technique.

Addressing RQ 1.2, we expected the task loads experienced by our users to reflect the amount of effort that they had to invest in each of the placement tasks, depending on which user had to perform the placement:

H2.1 The visiting user will experience a higher task load with the WIM placement technique than with the local placement technique.

H2.2 The hosting user will experience a higher task load with the local placement technique than with the WIM placement technique.

H2.3 The task load experienced by the hosting user in the local placement technique will be lower than the task load experienced by the visiting user in the WIM placement technique.

Since our system does feature several travel preview mechanisms, we expected our users to be able to orient themselves at the target position, regardless of the technique that was used during the target selection phase. Addressing RQ 1.3 thus brings us to the following hypotheses regarding the pointing precision of our users after a long-distance teleport location has been confirmed by the visiting user:

H3.1 There will not be a difference in the blind pointing accuracy of the visitor between the placement techniques

H3.2 There will not be a difference in the blind pointing accuracy of the hosting user between the placement techniques

5.2 Study Procedure

After the study was approved by our ethics board, participants were recruited via internal mailing lists as well as the digital notice boards of our town's universities. All participants received an expense allowance of 15 Euros for their participation in the study. Upon arrival in our laboratory, users were given some information about the general purpose of the study and filled out an informed consent

form. After giving their consent and asking any additional questions, users were randomly assigned either to the role of the orange or the blue user. They were then shown a video introducing the first study task and the method that they would use for their first trial round. To avoid order or task effects, the local and WIM placement methods were presented to the users in a counter-balanced order while the overall task order remained the same. After having gained a first impression of the task, users were helped with putting on their HMD and allowed to familiarize themselves with the single-user navigation technique and the virtual environment. They were then guided through four tutorial tasks with the blue user fulfilling the role of the visitor and the orange user fulfilling the role of the host before completing eight recorded tasks for which they did not receive any further guidance. After the first round of recorded tasks, the roles were swapped and the users underwent another four tutorials and eight recorded tasks. After completing the tasks with each participant fulfilling both roles, participants filled out an intermediate questionnaire and had the option to take a break before repeating the procedure with the second placement method. After filling out the final questionnaire for the first study task, users moved on to the second study task, described in Section 6.

5.3 Study Results

Overall, our sample consisted of 32 participants (13 female, 19 male) between the ages of 20 and 54 ($M = 27.5$; $SD = 7.3$). When asked about their experience with VR and HMDs, 5 reported to be first-time users, 18 had used them a few times, 6 several times, and 3 regularly. 12 users reported playing 3D video games never or seldomly while the remaining 20 were playing regularly or several times a week.

Based on our hypotheses, we analyzed our data using IBM SPSS Statistics. If both data series of a comparison were approximately normally distributed as indicated by a non-significant Shapiro-Wilk test, we conducted a paired-samples t-test. In the case of non-normality, the non-parametric Wilcoxon signed-rank test was performed instead. For both tests, we also computed the effect size r and applied the threshold values introduced by Cohen to identify small ($r > 0.1$), medium ($r > 0.3$), and large ($r > 0.5$) effect sizes [69]. For the remaining variables without corresponding hypotheses, only descriptive analyses were performed. In addition to the measures defined by our hypotheses, we recorded some exploratory data, of which we will present a descriptive evaluation.

5.3.1 Placement Accuracy

The distribution of position and angular placement accuracy scores is visualized in Figure 8 (left). The mean position placement accuracy scores were not normally distributed for local placement ($W(32) = 0.639$, $p < 0.001$). The results of the Wilcoxon signed-rank test indicated that the central tendencies of WIM placement ($Med = 0.132m$) and local placement ($Med = 0.050m$) were significantly different ($z = 4.619$, $p < 0.001$, $r = 0.820$). This supports H1.1 with a large effect size.

The mean angular placement accuracy scores were not normally distributed for both placement methods (both

Fig. 8. Boxplots illustrating the distribution of position and angular placement accuracy for both placement methods (left) as well as the distribution of task load scores for both placement methods and user roles (right).

$W(32) < 0.901$, $p < 0.007$). The results of the Wilcoxon signed-rank test indicated that the central tendencies of WIM placement ($Med = 4.974$) and local placement ($Med = 2.763$) were significantly different ($z = 3.815$, $p < 0.001$, $r = 0.674$). This supports H1.2 with a large effect size.

5.3.2 Task Load

The distribution of task load scores is visualized in Figure 8 (right). The mean task load scores of the visiting user were approximately normally distributed for both placement methods (both $W(32) > 0.944$, $p > 0.100$). The results of the paired-samples t-test indicated that the means of WIM placement ($M = 33.333$) and local placement ($M = 25.339$) were significantly different ($t(31) = 2.474$, $p = 0.019$, $r = 0.406$). This supports H2.1 with a medium effect size.

The mean task load scores of the hosting user were not normally distributed for both placement methods (both $W(32) < 0.919$, $p < 0.020$). The results of the Wilcoxon signed-rank test indicated that the central tendencies of WIM placement ($Med = 4.167$) and local placement ($Med = 8.750$) did not differ significantly ($z = 1.721$, $p = 0.085$, $r = 0.304$). We therefore cannot confirm H2.2.

Since the mean task load scores of the hosting user were not normally distributed, another Wilcoxon signed-rank test was performed for the comparison of task loads when executing the user placement. The results indicated that the central tendencies of the hosting user operating local placement ($Med = 8.750$) and the visiting user operating WIM placement ($Med = 33.333$) were significantly different ($z = 4.573$, $p < 0.001$, $r = 0.808$). This supports H2.3 with a large effect size.

5.3.3 Spatial Awareness

The mean blind pointing accuracy scores of the visiting user were not normally distributed for both placement methods (both $W(32) > 0.785$, $p < 0.006$). The results of the Wilcoxon signed-rank test indicated that the central tendencies of WIM placement ($Med = 15.934$) and local placement ($Med = 19.292$) did not differ significantly ($z = 0.879$, $p = 0.379$, $r = 0.067$). While the absence of a significant effect does not automatically indicate that the opposite

is true, the strongly overlapping 95% confidence intervals around the median (WIM placement [14:273 ; 20:448], local placement [12:853 ; 26:296]) provide initial indications to support H3.1.

The mean blind pointing accuracy scores of the hosting user were not normally distributed for both placement methods (both $W(32) < 0:893$, $p < 0:004$). The results of the Wilcoxon signed-rank test indicated that the central tendencies of WIM placement (Med = 2:208) and local placement (Med = 2:814) did not differ significantly ($z = 0:841$, $p = 0:400$, $r = 0:149$). As before, the strongly overlapping 95% confidence intervals around the median (WIM placement [1:682 ; 2:954], local placement [1:804 ; 3:459]) provide initial indications to support H3.2 nonetheless.

Since we did not have any hypotheses concerning the seeing-pointing score and were mostly interested in the overall precision our users achieved here, we only looked at the descriptive data for this task phase. The mean seeing pointing accuracy scores using WIM placement were 3.356 ($\sigma = 2:135$) for the visiting user and 2.846 ($\sigma = 1:055$) for the hosting user. For the local placement method, these values were 2.887 ($\sigma = 1:417$) for the visiting user and 2.784 ($\sigma = 1:060$) for the hosting user.

The mean seeing pointing duration using WIM placement was 1.732s ($\sigma = 0:922$ s) for the visiting user and 1.469s ($\sigma = 1:340$ s) for the hosting user. For the local placement method, these values were 1.629s ($\sigma = 0:448$ s) for the visiting user and 1.674s ($\sigma = 1:150$ s) for the hosting user.

5.3.4 Discomfort (Descriptive Only)

The discomfort scores were generally very low and therefore did not show any difference trends based on the placement technique and user role. Among all 64 scores measured in the study, 53 (90.6%) fell into the low-end range between 0 and 2. 10 (15.6%) more observations were located in the following range between 3 and 5, all of which were recorded in the WIM placement condition. A single outlier reported a score of 9 in the WIM placement condition. On further inquiry by the experimenter, they reported that this was not due to simulator sickness but due to them being worried about their performance with the technique and that they wished to continue with the study nonetheless. After the second study condition, where they were presented with the local placement method, the same user reported a score of 0.

5.3.5 Technique Preferences (Descriptive Only)

As the visiting user, 18 participants preferred using the local placement and 16 preferred using the WIM placement. The most given reason for preferring the WIM placement, named by 8 users, was the feeling that they could orient themselves at the target more easily when specifying their own target position. Five users reported that they liked the increased control or responsibility they had with this technique. Other reasons given by one user each were the enjoyment of the increased interaction with the system, the increased challenge, and that they felt that they were more precise with this technique. The most named reasons for preferring the local placement were users feeling more precise or confident using the technique (12 mentions) with two users mentioning that the local user had a better estimate of the target region,

seeing it true to size. Five users liked that the visitor could concentrate on other tasks during the target selection and four users felt that they had a better orientation at the target position during local placement. Other reasons named were that users found the technique more pleasant (2), felt more confident (1), and were faster (1).

When fulfilling the role of the hosting user, 25 users preferred the local placement and only 7 the WIM placement. Reasons given for preferring the WIM placement were the reduced dependency on the other player (3), and that they felt that it made their task more relaxing (3) or less complicated (1). Two users who picked the WIM placement technique also mentioned that they felt mostly neutral about the technique used as the hosting user. The main reason that users preferred the local placement technique as the hosting user was the increased engagement (14) and teamwork (4). Furthermore, twelve users reported that they were able to use the technique easily, with five users specifying that they were happy to support their team partner with an accurate placement. Four users mentioned that it was easier for them to know where their partner would arrive after placement and three users simply mentioned that they had more fun using the technique.

6 EVALUATION OF THE IMPLEMENTED INTERACTION SYSTEM

6.1 Study Design

While the first study task consisted of a mainly quantitative evaluation of the target planning techniques for long-distance teleportation, the second study task was planned to evaluate the overall concept behind our interaction system as well as the specific features chosen for the different phases of collaboration introduced in our framework. Since both study tasks were completed in one sitting, information about the overall study environment as well as the participant sample can be found in Section 5. In particular, we intended to evaluate the overall usability of our current system and identify possible improvements by asking the following research questions:

RQ2.1 Does our developed interaction system enable user groups to change between phases of tightly- and loosely-coupled collaboration while being distributed across the virtual environment?

RQ2.2 Are the suggested interface elements chosen appropriately to support the different stages of interaction?

RQ2.3 Are there features that are missing from our system that could further improve flexible user interaction during collaboration?

6.1.1 Task Design

To tackle these research questions, we chose a simplified scenario of two experts working on separate tasks within the virtual environment who have to consult with each other at some key points of their work. The study once again included the suggested button interface. Due to there only being two users we omitted the step of having to manually select the other user on the WIM as well as the "Join" option after the visit phase.

Fig. 9. Overview of the second study task. The left image shows the study environment with the routes that the blue and orange users completed. Along the routes, users found maintenance tasks meant either for one (single-color circles) or both (two-color circles) team members. The right side shows two of these maintenance tasks in detail, one from the perspective of the orange user (top) and one from the perspective of the blue user (bottom).

To avoid introducing our users to additional tasks or complex expert knowledge on top of trying out our system, we decided to use different semantic views of the virtual environment to simulate varying expertise. When immersed in the virtual environment, each user was shown a separate route that connected a series of different waypoints to be navigated to. Each of these waypoints contained a simplified maintenance task, consisting of pressing a button in the respective user color (orange or blue) to change it from dimmed to glowing. The routes that were shown to each user led away from each other in the virtual environment and did not overlap to encourage the use of virtual communication and travel techniques within the study. To represent both users having different areas of expertise, users were not able to see each other's routes. In addition, users could only see the current status of their own respective buttons. Buttons that had to be operated by the other user always appeared grey such that a user was not able to judge if they required any interference from their partner. On a route, some waypoints required a user to press one button in their color (indicated by single-color circles in Figure 9) and could therefore be solved alone. Other waypoints featured one button of each color (indicated by two-colored circles in Figure 9) and therefore required consultation by the remote collaborator. Upon arriving at such a waypoint, users alerted their team member, who could then use the check-in stage to determine if their button was in the correct state. Depending on their assessment, they could initiate a visit or inform their partner that no further action was required. In contrast to the previous study, visiting users could choose between both placement options for each visit by selecting the respective menu option. Once they had fulfilled a practice task, users were presented with the study task shown in Figure 9 that consisted of nine-way points for each route, of which four contained one button for each user.

6.1.2 Measures

The main purpose of this more open task was to test out the developed system's features and to gather feedback on whether participants found it useful and what could be

improved for future uses of a similar system. Therefore, we encouraged participants to form an opinion of the presented interaction techniques in the system while we observed them working on the joint task. After completing the task, both team members were given digital questionnaires about their experiences which they filled in at individual workstations.

In the questionnaires, they were asked to rate the individual features of our system on a seven-point scale from very disturbing (1) to very useful (7). In addition, they were asked if there were any further features that would have helped them during their use of the system. Users were also asked which of the two offered placement techniques was used by their group for the completion of this more open task and to rate their ability to continue with their own work after returning from a visit to the other user. To further supplement the data gathered with the questionnaires, the experimenter noted down any comments made by participants during the task. Since the second task's purpose was to gain an initial impression of the use of our proposed stages of collaboration, we did not formulate any concrete hypotheses.

6.2 Study Procedure

After completing the first study task as described in Section 5 and taking a break, our user groups were introduced to the second study task using a short video of the technique and task scenario. They then had the opportunity to try out the system in a tutorial task consisting of two way-points which both contained two inactivated buttons, such that each user had to undergo all stages of collaboration for the task to be solved. After indicating that they understood the task, users moved on to the recorded route task shown in Figure 9. After the task was completed successfully, users moved on to the last set of questionnaires ending with the collection of some demographic data. Before leaving, users received their expense allowance and were debriefed appropriately.

6.3 Study Results

Overall, all 16 teams were able to successfully solve the given task while using the presented techniques to split up during their work in the environment and consult each other when needed before returning to their individual sub-tasks. Since the second study task had a more exploratory character without formulated hypotheses, our analysis of the resulting data will remain descriptive and focus on the feedback users gave on our collaboration methods via the questionnaires after their joint interaction in the virtual space as well as their intermediate comments to the experimenter and to each other during the study task.

6.3.1 Chosen Placement Method

Out of the 16 groups, 13 reported that they used both placement techniques during the task, two groups exclusively used the local placement technique and one group only used the WIM placement technique. When asked for their reasoning for using both techniques, 4 users answered that this was due to different preferences within the group. The remaining 22 users reported that they wanted to try out both

Overview	O1	World in Miniature (WIM) for Overview
	O2	Abstract Representation of Partner on the WIM
	O3	Abstract Representation of Oneself on the WIM
	O4	Opportunity to Contact Partner
	O5	Opportunity to Ask Partner for Contact
Check-In	C1	WIM Zoomed in on Partner
	C2	Preview Window of Partner's Perspective
	C3	Local Placement Method
	C4	WIM Placement Method
	C5	Option to Teleport to Partner
Visit	V1	Ghost Avatar on WIM Previewing the Teleportation Target
	V2	Preview Window of Perspective at Teleportation Target
	V3	Full-Sized Ghost Avatar at Teleportation Target
Return	R1	Option to Return to Previous Position
	R2	Seeing the Previous Position on the Overview WIM
	R3	Preview Window of Perspective at Previous Position

Fig. 10. Top: Boxplots illustrating the distributions of feature ratings from 1 (very disturbing) through 4 (neither disturbing nor useful) to 7 (very helpful). Bottom: List of abbreviations of the surveyed features.

techniques again. Out of the users who gave further explanations, 7 concluded that they preferred the local placement technique overall, 3 preferred the WIM placement technique and 4 users preferred switching between techniques for variance or depending on the current task situation. The two groups who exclusively used the local placement technique explained their choice with the technique feeling more cooperative and precise while the group who exclusively used the WIM placement technique found it easier to choose their own target position.

6.3.2 Feature Ratings and User Suggestions

The obtained rating scores regarding the features of our system are shown in Figure 10. The five features regarding the overview phase all received a favorable median score of 6/7. The most and least appreciated features as indicated by the data ranges without outliers were the opportunity to contact the other user (O4) and the opportunity to ask the other user for contact (O5), respectively. When asked for features that they missed or changes that they wanted to be made in the overview stage, two users mentioned wanting a shortcut for the calling functionality, especially for asking the other user to call them. Another feature mentioned was the setting of target suggestions at points of interest that can be used by team members later without direct interaction.

The median scores of the five features available during the check-in and target selection stage were also entirely located at 6. Here, the option to actually travel to the other user (C5) was rated best while having the option of activating the WIM- and local placement methods (C3, C4) were most controversial depending on user preference.

The three visit features also showed median scores of 6, with the WIM and ghost avatar (V1) being rated most favorably and the ghost avatar at the target location (V3) being rated most controversially. When asked about missing features to the check-in and target selection phases, one user mentioned that they would have liked more awareness cues for the hosting user during calls. Further suggestions were to be able to tilt the WIM for a better view and to add facial expressions to the avatars.

Finally, the option to return to the previous position (R1) received the overall highest median score of 7 with only a single outlier at 3. However, the visualization of the previous position on the WIM (R2) as well as the preview window of the corresponding view (R3) only achieved medians of 5 with scores ranging down to 2 and 3, respectively. With respect to missing features, users mentioned that they would have liked a button to directly activate the return to their previous position without activating the target preview first. One user suggested leaving a beacon at the original location to help with task continuation and another user asked for the option to take their partner along when traveling back to their original position, for cases where both users needed the other's help.

7 DISCUSSION OF USER STUDY RESULTS

7.1 Comparison of Placement Techniques

Overall, the results of our quantitative comparison of the two placement techniques fulfilled our expectations, showing that a transfer of the target selection responsibility to the local user led to an increase in precision, a decrease in task load for the visiting user, and an overall decrease in the maximum task load experienced within the team. These qualities were reflected in the feedback of users who preferred the local placement technique with precision and ease of use being mentioned repeatedly. Furthermore, no significant difference was found between the task load of the hosting user when placing the visiting user and when waiting for the visiting user to place themselves. While this does not prove the absence of a difference, it provides some further indication for the ease of use of the local placement method.

One concern of ours going into the study had been that the distribution of control introduced by the local placement method would result in bothersome overheads that might irritate or distract both collaborators. Interestingly, enjoyment of the interaction with their partner was repeatedly mentioned when asked for hosting user's preference, with only a few users being content to relax or not wanting to take responsibility for their partner's placement. While a large majority of users preferred the local placement technique as the hosting user, preferences were more split from the perspective of the visiting user. Although the data did not show an overall advantage for one of the two techniques in terms of pointing accuracy at the target, both had some users who felt that they had a better spatial orientation during their use. Apart from this aspect, opinions were mainly split between users wanting to retain full control and users who were happy to defer the placement to their partner. Some users felt that the distribution of control led

to better results and enjoyed the interaction with their partner from both perspectives. Others liked being in control over their own performance or wanted to interact with the system as much as possible. When looking at the second study phase, where groups could choose how to select their teleport target, nearly all of them used both of the placement techniques. Some users arrived at a preference for one of the two and others enjoyed the freedom to choose depending on their current situation. Given this split in preferences, we recommend offering both placement techniques in the context of free collaboration between equal partners and letting users choose depending on their current needs and preferences.

7.2 Feature Ratings and User Suggestions

Looking at the implemented system, all of our interface features received positive evaluations from a majority of the participants. This suggests an overall benefit of the offered phases of collaborative interaction in our two-user scenario. These results are also supported by the fact that user groups with varying levels of experience in VR were able to use the system to collaborate with each other and to solve the given study task after a short introduction phase. In the following sections, we will discuss the user feedback we received for each stage and devise recommendations on how to improve our system for the use in collaborative virtual environments.

7.2.1 Overview

In the overview stage, users were most convinced by the ability to call other users and least convinced by the ability to request a call. This might be due to the fact that our rather simple task did not require a lot of engagement outside of exchanging information with the other user, making the double confirmation obsolete. In addition, some of the user groups just talked to each other instead of requesting a call, since there were no other people to disturb in the shared space. Since this might change with more involved tasks and larger user groups who might not uphold a constant audio connection, we would still suggest not completely discarding this feature. Nevertheless, larger groups might benefit from being able to manage a form of call etiquette by changing their status between being readily available to talk, wanting to be asked first, and blocking external distractions.

7.2.2 Check-In and Travel Preparation

While the features provided to the visiting user during the check-in stage were generally well received, some users missed further feedback for the hosting user during this stage, who was only alerted that a call was going on but did not know exactly what the other user was seeing on the WIM. To resolve this problem, different approaches from related work on asymmetric collaboration could be employed, ranging from subtle awareness cues indicating the collaborator's field of view to more explicit collaboration mechanisms [58]. Since the ghost avatar at the target location felt erratic to some hosting users when used to reflect the visitor's target selection with the WIM placement technique, future implementations of the system should consider smoothing these movements. This same consideration should also be made for any added asymmetrical feedback

mechanism. One idea that was mentioned by a few users was the option to prepare teleport goals that could be used by their partner at a later time. This might be an interesting approach towards combining the spatial exibility offered by our current system with ways to support asynchronous interaction.

7.2.3 Return

Overall, the option to return to one's target location was the most positively rated feature while simultaneously having the worst ratings for the corresponding interface elements. Since some users suggested adding a shortcut for this option, we believe that the level of complexity offered by our system was not needed for the given study task. In general, our users reported that they had no problems continuing with their tasks after returning to the previous location. Nevertheless, two users mentioned that they would have liked post-travel feedback at the return target. A combination of a return shortcut with post-travel feedback to ease orientation at the previous position might be an alternative to the current return process. Finally, some users also suggested that they would have liked the option to take their partner along during a return teleportation. This would be an interesting way to combine the spatial exibility given by our system with the ability of group navigation techniques to continue tightly coupled collaboration while exploring a virtual environment together.

8 DISCUSSION OF POTENTIAL USAGE SCENARIOS

As discussed in Section 3, the concrete implementation of the suggested stages of collaboration and the connected challenges are largely use-case dependent. We therefore want to explore additional requirements that result from two example use cases. To do this, we will first discuss the interaction of visitors in museums and digital heritage sites as an example where users in the same role but with different interests and goals experience a shared space. Second, we will discuss the scenario of a teacher or instructor in the context of education and training to explore the potential in use of special user roles on the framework.

8.1 Visitors in Museums and Digital Heritage Sites

Our first use-case from related work that presents many interesting challenges for the support of different forms of collaboration is the exploration of virtual museums and digital heritage sites through visitor groups. Users of such exhibition spaces can follow different approaches regarding their use of the exhibits, depending on their character, company, and the size of the virtual space [70] and can be introduced to content and ideas through exploration, presentations, and interactive challenges [71].

In such a context, the overview phase could help users find interesting exhibits by identifying where clusters are forming in the virtual space. In addition, the overview might display information on the current availability of others. If users find an especially interesting exhibit or want to help other users, they can initiate a Check-Into see if something discovered by them is interesting and new enough to someone else to warrant a face-to-face exchange. If this

is the case, a Visit presents the opportunity to engage with other users at the target location, taking a closer look at an interesting exhibit or engaging with an interactive part of the exhibition together. Based on our study results, we would recommend letting the users themselves choose their preferred placement method for the necessary long-distance teleport, depending on each interacting partner's comfort level with VR technology and need for independence. Similar to the Join and Return differentiation proposed in our framework, it should be made clear if a user is expected to return to their original position in the near future and is not expected to join their partner for a longer period. If this is the case, other users might refrain from changing an interactive exhibit or wait for the other person to return if they want to engage with something together. If users decide for joining a longer activity at the location they are visiting, they still might be allowed to retain several return points, but they should not block other users' views of the exhibits with life-sized preview avatars.

Interesting additional challenges in the museum space lie in managing how much information different users want to receive from others in the virtual space during different phases of collaboration. Users who arrive together with others might be mainly interested in a certain sub-group of people that they already know. Some users who are exploring alone might prefer to keep to themselves without engaging others much but still want to see other users' activities in some way, to orient themselves, and to become aware of points of interest. On the other end of the spectrum, users might want to find new people who share their interests or to do activities together. Each of these preferences, which might even change throughout the exhibition, calls for different interaction patterns to support them. Another topic that could be promising to explore in this context is the inclusion of asynchronous collaboration metaphors. Similar to the different phases of interaction proposed, previous users' activities could be integrated into an overview of the space, or people could actively leave notes, drawn hints, or even recordings of themselves (as seen in [72]) in locations of interest that could be engaged with similar to a Check-In or Visit.

8.2 Instructors in Education and Training

In the previous scenario, we assume users to have different wishes or interests, but broadly the same task (exploring the space in a way that benefits them best) and role (one visitor of a virtual space). In contrast to this, applications used for teaching and training often feature a group of learners being managed by one or more people who might need different mechanisms to support them in overseeing and managing activities [73]. While the students themselves might use the system in a similar way to the museum visitors that were discussed previously, we want to highlight the differences in use that might exist for someone in the instructor role.

Here, the overview stage might be used primarily to gauge the overall activities of the students and not to select one's independent activity and task. Instructors might use this overview to decide who needs help or closer supervision or simply when to end a shared activity. Here, activity indicators that could be easily collected by the system such

as showing people as talking, interacting, or idle might be especially useful. Check-ins could, as usual, be an opportunity to answer smaller questions by any of the students, but could also be used for teachers making rounds during an activity or wanting to check in on a group based on their previously observed activities. One interesting question here is finding appropriate agreement mechanisms for a check-in based on the target group. Visits could be used for intermediate presentations or more extensive help on a problem that students have. While the basic Join or Return logic might not make the most sense for a supervisor without their active working place, awareness mechanisms about their current availability might still be needed for students waiting for help.

Interesting additions to the scenario might be the creation of mechanisms to gather all or sub-groups of students at one group's location or in a shared new environment to present results to each other or to start new activities. In addition, specific support measures and communication rules for different activities such as individual work, separate groups, or exchanges between the whole class could be implemented.

9 CONCLUSION AND FUTURE WORK

Building on solutions from related work on supporting collaboration, we presented users with a system which directly supports different phases between tightly- and loosely-coupled collaboration. Our work shows that VR users can benefit from and enjoy using an interface which directly supports their fluent transition between these phases. Furthermore, our comparison of virtual placement techniques demonstrated that the distribution of task responsibility based on current user status can improve team performance and reduce mental load induced by long-distance teleportation.

The feedback gathered in our user study also inspired some interesting approaches for future work like the use of group navigation techniques in combination with our system or the support of asynchronous collaboration for further flexibility. Other areas that are still left to explore are the specific support of different user roles with different responsibilities and the adaptation of the suggested system to different environments and tasks. Finally, the most interesting challenge for future work lies in the further investigation of the proposed phases of collaboration by extending the presented system for the use by larger teams and in different usage contexts like the ones described in Section 8.

Overall, we believe that collaboration in social VR has yet to reach its full potential. Going beyond the limitations imposed by the real world, the concepts presented in traditional CSCW research provide us with interesting inspirations for novel VR adaptations, which shows great promise to further shape the way in which VR users can collaborate over large virtual and physical distances.

ACKNOWLEDGMENTS

This work was funded in part by the Ministry for Science and Health of Rhineland-Palatinate as part of the research

training group (Forschungskolleg) "Immersive Extended Reality for Physical Activity and Health" (XR-PATH). This work was also partially funded by the German Ministry of Education and Research (BMBF) under grant 16SV8716 (project Goethe-Live-3D) and the Thuringian Ministry of Economic Affairs, Science and Digital Society (TMWDDG) under grant 5575/10-5 (project MetaReal).

REFERENCES

- [1] C. Gutwin and S. Greenberg, "Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness," in *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* 1998, pp. 207–216.
- [2] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale, "Collaborative Coupling over Tabletop Displays," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* New York, NY, USA: Association for Computing Machinery, 2006, pp. 1181–1190.
- [3] T. Neumayr, H.-C. Jetter, M. Augstein, J. Friedl, and T. Luger, "Domino: A Descriptive Framework for Hybrid Collaboration and Coupling Styles in Partially Distributed Teams," *Proceedings of the ACM on Human-Computer Interaction* vol. 2, no. CSCW, 2018.
- [4] J. McVeigh-Schultz, A. Kolesnichenko, and K. Isbister, "Shaping Pro-Social Interaction in VR: An Emerging Design Framework," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* New York, NY, USA: Association for Computing Machinery, 2019, p. 1–12.
- [5] D. Maloney and G. Freeman, "Falling Asleep Together: What Makes Activities in Social Virtual Reality Meaningful to Users," in *Proceedings of the Annual Symposium on Computer-Human Interaction in Play* New York, NY, USA: Association for Computing Machinery, 2020, p. 510–521.
- [6] A. Kolesnichenko, J. McVeigh-Schultz, and K. Isbister, "Understanding Emerging Design Practices for Avatar Systems in the Commercial Social VR Ecology," in *Proceedings of the 2019 Designing Interactive Systems Conference* New York, NY, USA: Association for Computing Machinery, 2019, p. 241–252.
- [7] H. Xia, S. Herscher, K. Perlin, and D. Wigdor, "Spacetime: Enabling Fluid Individual and Collaborative Editing in Virtual Reality," in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* New York, NY, USA: Association for Computing Machinery, 2018, p. 853–866.
- [8] D. Ververidis, S. Nikolopoulos, and I. Kompatsiaris, "A Review of Collaborative Virtual Reality Systems for the Architecture, Engineering, and Construction Industry," *Architecture* vol. 2, no. 3, pp. 476–496, 2022.
- [9] S. Philippe, A. D. Souchet, P. Lameris, P. Petridis, J. Caporal, G. Coldeboeuf, and H. Duzan, "Multimodal teaching, learning and training in virtual reality: a review and case study," *Virtual Reality & Intelligent Hardware* vol. 2, no. 5, pp. 421–442, 2020.
- [10] Y. Li, E. Ch'ng, S. Cai, and S. See, "Multiuser Interaction with Hybrid VR and AR for Cultural Heritage Objects," in *2018 3rd Digital Heritage International Congress* 2018, pp. 1–8.
- [11] J. Torous, S. Bucci, I. Bell, L. Kessing, M. Faurholt-Jepsen, P. Whelan, A. Carvalho, M. Keshavan, J. Linardon, and J. Firth, "The growing field of digital psychiatry: current evidence and the future of apps, social media, chatbots, and virtual reality," *World Psychiatry* vol. 20, no. 3, pp. 318–335, 2021.
- [12] B. Ens, J. Lanir, A. Tang, S. Bateman, G. Lee, T. Piumsomboon, and M. Billinghamurst, "Revisiting Collaboration through Mixed Reality: The Evolution of Groupware," *International Journal of Human-Computer Studies* vol. 131, pp. 81–98, 2019.
- [13] E. F. Churchill and D. Snowdon, "Collaborative Virtual Environments: An Introductory Review of Issues and Systems," *Virtual Reality*, vol. 3, no. 1, pp. 3–15, 1998.
- [14] S. Benford, J. Bowers, L. E. Fahren, J. Mariani, and T. Rodden, "Supporting Cooperative Work in Virtual Environments," *The Computer Journal* vol. 37, no. 8, pp. 653–668, 1994.
- [15] R. Johansen, *GroupWare: Computer Support for Business Teams* USA: The Free Press, 1988.
- [16] M. Billinghamurst, S. Bee, J. Bowskill, and H. Kato, "Asymmetries in Collaborative Wearable Interfaces," in *Digest of Papers. Third International Symposium on Wearable Computing* 1999, pp. 133–140.
- [17] N. Feld and B. Weyers, "Mixed Reality in Asymmetric Collaborative Environments: A Research Prototype for Virtual City Tours," in *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW'21)*, pp. 250–256.
- [18] T. Piumsomboon, G. A. Lee, and M. Billinghamurst, "Snow Dome: A Multi-Scale Interaction in Mixed Reality Remote Collaboration," in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* New York, NY, USA: Association for Computing Machinery, 2018, p. 1–4.
- [19] T. Piumsomboon, G. A. Lee, J. D. Hart, B. Ens, R. W. Lindeman, B. H. Thomas, and M. Billinghamurst, "Mini-Me: An Adaptive Avatar for Mixed Reality Remote Collaboration," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* New York, NY, USA: Association for Computing Machinery, 2018, p. 1–13.
- [20] C. G. Fidalgo, Y. Yan, H. Cho, M. Sousa, D. Lindlbauer, and J. Jorge, "A Survey on Remote Assistance and Training in Mixed Reality Environments," *IEEE Transactions on Visualization and Computer Graphics* vol. 29, no. 5, pp. 2291–2303, 2023.
- [21] B. Sousa Santos, P. Dias, A. Pimentel, J.-W. Baggerman, C. Ferreira, S. Silva, and J. Madeira, "Head-Mounted Display versus Desktop for 3D Navigation in Virtual Reality: A User Study," *Multimedia Tools and Applications* vol. 41, no. 1, p. 161–181, 2009.
- [22] A. Kunert, T. Weissker, B. Froehlich, and A. Kulik, "Multi-Window 3D Interaction for Collaborative Virtual Reality," *IEEE Transactions on Visualization and Computer Graphics* vol. 26, no. 11, pp. 3271–3284, 2020.
- [23] P. Dourish and V. Bellotti, "Awareness and Coordination in Shared Workspaces," in *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work* New York, NY, USA: Association for Computing Machinery, 1992, p. 107–114.
- [24] C. Gutwin and S. Greenberg, "The Effects of Workspace Awareness Support on the Usability of Real-Time Distributed Groupware," *ACM Transactions on Computer-Human Interaction* vol. 6, no. 3, p. 243–281, 1999.
- [25] E. Tse, J. Histon, S. Scott, and S. Greenberg, "Avoiding Interference: How People Use Spatial Separation and Partitioning in SDG Workspaces," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work* 2004, pp. 252–261.
- [26] J. R. Wallace, S. D. Scott, T. Stutz, T. Enns, and K. Inkpen, "Investigating teamwork and taskwork in single- and multi-display groupware systems," *Personal and Ubiquitous Computing* vol. 13, no. 8, pp. 569–581, 2009.
- [27] T. Dodds and R. Ruddle, "Using Teleporting, Awareness and Multiple Views to Improve Teamwork in Collaborative Virtual Environments," in *14th Eurographics Symposium on Virtual Environments* B. Mohler and R. van Liere, Eds. Eurographics Association, 2008, pp. 81–88.
- [28] T. J. Dodds and R. A. Ruddle, "Mobile Group Dynamics in Large-Scale Collaborative Virtual Environments," in *2008 IEEE Virtual Reality Conference* 2008, pp. 59–66.
- [29] A. Kunert, A. Kulik, S. Beck, and B. Froehlich, "Photoportals: Shared References in Space and Time," in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* New York, NY, USA: Association for Computing Machinery, 2014, p. 1388–1399.
- [30] J.-H. Schöder, D. Schacht, N. Peper, A. M. Hamurculu, and H.-C. Jetter, "Collaborating Across Realities: Analytical Lenses for Understanding Dyadic Collaboration in Transitional Interfaces," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* New York, NY, USA: Association for Computing Machinery, 2023.
- [31] T. Weissker, P. Bimberg, and B. Froehlich, "An Overview of Group Navigation in Multi-User Virtual Reality," in *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, 2021, pp. 363–369.
- [32] D. A. Bowman, E. Kruijff, J. LaViola, Joseph J., and I. Poupyrev, "An Introduction to 3-D User Interface Design," *Presence: Teleoperators and Virtual Environments* vol. 10, no. 1, pp. 96–108, 02 2001.
- [33] D. Zielasko, S. Horn, S. Freitag, B. Weyers, and T. W. Kuhlen, "Evaluation of Hands-Free HMD-Based Navigation Techniques for Immersive Data Analysis," in *2016 IEEE Symposium on 3D User Interfaces (3DUI)* 2016, pp. 113–119.
- [34] D. Zielasko, Y. C. Law, and B. Weyers, "Take a Look Around – The Impact of Decoupling Gaze and Travel-direction in Seated and Ground-based Virtual Reality Utilizing Torso-directed Steering," in *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 2020, pp. 398–406.

- [35] E. Bozgeyikli, A. Raij, S. Katkoori, and R. Dubey, "Point & Teleport Locomotion Technique for Virtual Reality," in *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play*. New York, NY, USA: Association for Computing Machinery, 2016, p. 205–216.
- [36] T. Weissker, A. Kunert, B. Froehlich, and A. Kulik, "Spatial Updating and Simulator Sickness During Steering and Jumping in Immersive Virtual Environments," in *2018 IEEE Conference on Virtual Reality and 3D User Interfaces*, 2018, pp. 97–104.
- [37] R. Stoakley, M. J. Conway, and R. Pausch, "Virtual Reality on a WIM: Interactive Worlds in Miniature," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. USA: ACM Press/Addison-Wesley Publishing Co., 1995, p. 265–272.
- [38] J. S. Pierce, B. C. Stearns, and R. Pausch, "Voodoo Dolls: Seamless Interaction at Multiple Scales in Virtual Environments," in *Proceedings of the 1999 Symposium on Interactive 3D Graphics*. New York, NY, USA: Association for Computing Machinery, 1999, p. 141–145.
- [39] A. Kulik, A. Kunert, S. Beck, R. Reichel, R. Blach, A. Zink, and B. Froehlich, "C1x6: A Stereoscopic Six-User Display for Co-located Collaboration in Shared Virtual Environments," *ACM Transactions on Graphics*, vol. 30, no. 6, p. 1–12, 2011.
- [40] H. Salzmann, J. Jacobs, and B. Froehlich, "Collaborative Interaction in Co-Located Two-User Scenarios," in *Proceedings of the 15th Joint Virtual Reality Eurographics Conference on Virtual Environments*. Eurographics Association, 2009, p. 85–92.
- [41] T. Weissker, A. Kulik, and B. Froehlich, "Multi-Ray Jumping: Comprehensive Group Navigation for Collocated Users in Immersive Virtual Reality," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces*, 2019, pp. 136–144.
- [42] T. Weissker, P. Bimberg, and B. Froehlich, "Getting There Together: Group Navigation in Distributed Virtual Environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 5, pp. 1860–1870, 2020.
- [43] S. Beck, A. Kunert, A. Kulik, and B. Froehlich, "Immersive Group-to-Group Telepresence," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 4, pp. 616–625, 2013.
- [44] R. A. Ruddle, A. Howes, S. J. Payne, and D. M. Jones, "The Effects of Hyperlinks on Navigation in Virtual Environments," *International Journal of Human-Computer Studies*, vol. 53, no. 4, pp. 551–581, 2000.
- [45] D. Cliburn, S. Rilea, D. Parsons, P. Surya, and J. Semler, "The Effects of Teleportation on Recollection of the Structure of a Virtual World," in *Proceedings of the 15th Joint Virtual Reality Eurographics Conference on Virtual Environments*, 2009, p. 117–120.
- [46] S. Kerns, D. Cliburn, and M. O'Rourke, "Teleporting in Virtual Worlds while Learning Real World Places," in *20th International Conference on Artificial Reality and Telexistence*, 2010, pp. 126–130.
- [47] C. Wingrave, Y. Hachiahetoglu, and D. Bowman, "Overcoming World in Miniature Limitations by a Scaled and Scrolling WIM," in *3D User Interfaces (3DUI'06)*, 2006, pp. 11–16.
- [48] R. Trueba, C. Andujar, and F. Argelaguet, "Complexity and Occlusion Management for the World-in-Miniature Metaphor," in *Smart Graphics*, 2009, pp. 155–166.
- [49] T. T. H. Nguyen, T. Duval, and C. Fleury, "Guiding Techniques for Collaborative Exploration in Multi-scale Shared Virtual Environments," in *Proceedings of the International Conference on Computer Graphics Theory and Applications and International Conference on Information Visualization Theory and Applications*. SciTePress, 2013, pp. 327–336.
- [50] A. Cockburn, A. Karlson, and B. B. Bederson, "A Review of Overview+Detail, Zooming, and Focus+Context Interfaces," *ACM Computing Surveys*, vol. 41, no. 1, 2009.
- [51] D. A. Bowman, C. North, J. Chen, N. F. Polys, P. S. Pyla, and U. Yilmaz, "Information-Rich Virtual Environments: Theory, Tools, and Research Agenda," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2003, p. 81–90.
- [52] A. Drogemuller, A. Cunningham, J. Walsh, M. Cordeil, W. Ross, and B. Thomas, "Evaluating Navigation Techniques for 3D Graph Visualizations in Virtual Reality," in *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*, 2018, pp. 1–10.
- [53] A. Kendon, "Spacing and Orientation in Co-present Interaction," in *Development of Multimodal Interfaces: Active Listening and Synchrony*, 2010, pp. 1–15.
- [54] P. Marshall, Y. Rogers, and N. Pantidi, "Using F-Formations to Analyse Spatial Patterns of Interaction in Physical Environments," in *Proceedings of the ACM 2011 Conference on Computer Supported Co-operative Work*. New York, NY, USA: Association for Computing Machinery, 2011, p. 445–454.
- [55] S. A. Alharthi, G. E. Raptis, C. Katsini, I. Dolgov, L. E. Nacke, and Z. O. Toups, "Investigating the Effects of Individual Cognitive Styles on Collaborative Gameplay," *ACM Transactions on Computer-Human Interaction*, vol. 28, no. 4, 2021.
- [56] J. Herskovitz, Y. F. Cheng, A. Guo, A. P. Sample, and M. Nebeling, "XSpace: An Augmented Reality Toolkit for Enabling Spatially-Aware Distributed Collaboration," *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. ISS, 2022.
- [57] L. Poulliquen-Lardy, I. Milleville-Pennel, F. Guillaume, and F. Mars, "Remote collaboration in virtual reality: asymmetrical effects of task distribution on spatial processing and mental workload," *Virtual Reality*, vol. 20, no. 4, pp. 213–220, 2016.
- [58] S. Kim, G. Lee, W. Huang, H. Kim, W. Woo, and M. Billinghurst, "Evaluating the Combination of Visual Communication Cues for HMD-Based Mixed Reality Remote Collaboration," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–13.
- [59] M. Husung and E. Langbehn, "Of Portals and Orbs: An Evaluation of Scene Transition Techniques for Virtual Reality," in *Proceedings of Mensch Und Computer 2019*, 2019, p. 245–254.
- [60] N. Feld, P. Bimberg, B. Weyers, and D. Zielasko, "Keep It Simple? Evaluation of Transitions in Virtual Reality," in *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2023.
- [61] P. Bimberg, T. Weissker, A. Kulik, and B. Froehlich, "Virtual Rotations for Maneuvering in Immersive Virtual Environments," in *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2021.
- [62] C. C. Presson and D. R. Montello, "Updating after rotational and translational body movements: coordinate structure of perspective space," *Perception*, vol. 23, no. 12, pp. 1447–1455, 1994.
- [63] E. Langbehn, J. Wittig, N. Katzakis, and F. Steinicke, "Turn Your Head Half Round: VR Rotation Techniques for Situations With Physically Limited Turning Angle," in *Proceedings of Mensch Und Computer 2019*. New York, NY, USA: Association for Computing Machinery, 2019, p. 235–243.
- [64] D. Zielasko, J. Heib, and B. Weyers, "Systematic Design Space Exploration of Discrete Virtual Rotations in VR," in *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 2022, pp. 693–702.
- [65] L. Rebenitsch and C. Owen, "Individual Variation in Susceptibility to Cybersickness," in *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2014, p. 309–317.
- [66] P. Bimberg, T. Weissker, and A. Kulik, "On the Usage of the Simulator Sickness Questionnaire for Virtual Reality Research," in *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, 2020, pp. 464–467.
- [67] S. G. Hart and L. E. Staveland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," in *Advances in Psychology*, 1988, vol. 52, pp. 139–183.
- [68] S. G. Hart, "Nasa-Task Load Index (NASA-TLX); 20 Years Later," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, no. 9, 2006, pp. 904–908.
- [69] J. Cohen, "A power primer," *Psychological Bulletin*, vol. 112, no. 1, pp. 155–159, 1992.
- [70] A. Antoniou and G. Lepouras, "Modeling Visitors' Profiles: A Study to Investigate Adaptation Aspects for Museum Learning Technologies," *Journal on Computing and Cultural Heritage*, vol. 3, no. 2, 2010.
- [71] A. Puig, I. Rodríguez, J. L. Arcos, J. A. Rodríguez-Aguilar, S. Cebrián, A. Bogdanovych, N. Morera, A. Palomo, and R. Piqué, "Lessons learned from supplementing archaeological museum exhibitions with virtual reality," *Virtual Reality*, vol. 24, no. 2, p. 343–358, 2020.
- [72] K. Chow, C. Coyiuto, C. Nguyen, and D. Yoon, "Challenges and Design Considerations for Multimodal Asynchronous Collaboration in VR," *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, no. CSCW, 2019.
- [73] H. Bani-Salameh, J. Al-Gharaibeh, C. Jeffery, and Z. A. Al-Sharif, "Collaborative Education in a Virtual Learning Environment," *International Journal of Business Information Systems*, vol. 25, no. 4, pp. 474–489, 2017.

