



Partially Blended Realities: Aligning Dissimilar Spaces for Distributed Mixed Reality Meetings

Jens Emil Grønbaek
Aarhus University, Denmark
jensemil@cs.au.dk

Ken Pfeuffer
Aarhus University, Denmark
ken@cs.au.dk

Eduardo Velloso
University of Melbourne, Australia
eduardo.velloso@unimelb.edu.au

Morten Astrup
Aarhus University, Denmark
201705289@post.au.dk

Melanie Sønderkær Pedersen
Aarhus University, Denmark
202003153@post.au.dk

Martin Kjær
Aarhus University, Denmark
201709372@post.au.dk

Germán Leiva
Aarhus University, Denmark
leiva@cc.au.dk

Hans Gellersen
Lancaster University, United Kingdom
Aarhus University, Denmark
h.gellersen@lancaster.ac.uk

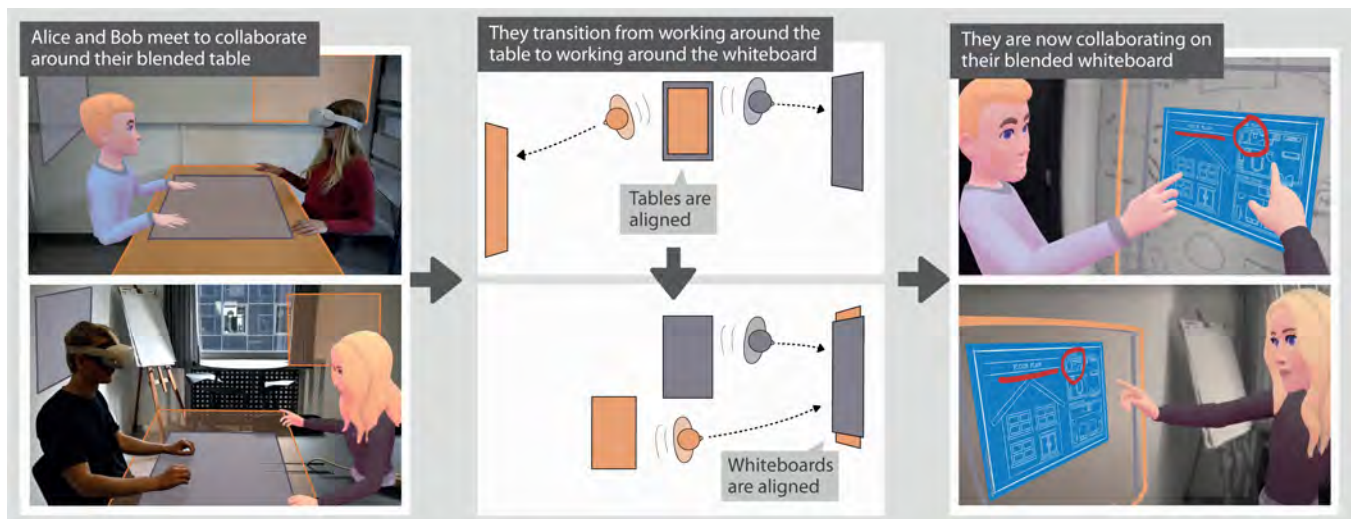


Figure 1: In Partially Blended Realities, remote collaborators meet in a distributed Mixed Reality space composed of their local surfaces. Dissimilar rooms will only be partially blended. This creates a need for realigning when collaborators move from one surface to another. We developed RealityBlender to study the user experience and design space of partial alignment techniques.

ABSTRACT

Mixed Reality allows for distributed meetings where people's local physical spaces are virtually aligned into blended interaction spaces. In many cases, people's physical rooms are dissimilar, making it challenging to design a coherent blended space. We introduce the concept of Partially Blended Realities (PBR) — using Mixed Reality to support remote collaborators in partially aligning their physical

spaces. As physical surfaces are central in collaborative work, PBR supports users in transitioning between different configurations of tables and whiteboard surfaces. In this paper, we 1) describe the design space of PBR, 2) present RealityBlender to explore interaction techniques for how users may configure and transition between blended spaces, and 3) provide insights from a study on how users experience transitions in a remote collaboration task. With this work, we demonstrate new potential for using partial solutions to tackle the alignment problem of dissimilar spaces in distributed Mixed Reality meetings.

CCS CONCEPTS

- Human-centered computing → Mixed / augmented reality; Collaborative and social computing systems and tools; Empirical studies in collaborative and social computing.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9421-5/23/04...\$15.00
<https://doi.org/10.1145/3544548.3581515>

KEYWORDS

mixed reality, augmented and virtual reality, remote collaboration, blended realities, proxemic transitions

ACM Reference Format:

Jens Emil Grønþæk, Ken Pfeuffer, Eduardo Velloso, Morten Astrup, Melanie Sønderkær Pedersen, Martin Kjær, Germán Leiva, and Hans Gellersen. 2023. Partially Blended Realities: Aligning Dissimilar Spaces for Distributed Mixed Reality Meetings. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3544548.3581515>

1 INTRODUCTION

For teleconferencing tools to truly support the future of work, they must go beyond simply enabling people to talk to each other and effectively support them in accomplishing tasks together. Previous research on Computer-Supported Cooperative Work (CSCW) has highlighted the importance of physical surfaces for coordinating activities, referencing information, and co-creating content [8, 10, 19, 27, 29, 32, 34]. Still, most videoconferencing applications fail to incorporate the physical environment around users into the collaborative experience. In this context, by combining digital and physical environments, Mixed Reality (MR) offers opportunities for addressing these limitations. Local users can access their space's physical features while communicating with distant users embodied as avatars. For example, it is possible to envision a scenario where users sit at their local physical desks individually but share the same desk in MR (Figure 1 left). Towards this goal, previous works have explored approaches that carefully arrange physical spaces to support this experience, e.g., by arranging spaces as two halves of the same room [4] or ensuring that both rooms are identical to overlay them in MR [26].

However, the rigidity of the spatial requirements in these solutions is incompatible with the dynamic nature of collaborative work, which is often organised around multiple physical surfaces [7, 27, 29, 32, 34], transitioning between these surfaces and multiple environments [1, 8, 22, 28, 35]. Specifically, solutions must enable people to work together in MR while leveraging the surfaces in the physical environment, *even when these environments are different*. Consider the scenario of two people working remotely across a table and a whiteboard (see Figure 1). Each collaborator has access to these surfaces in their own physical environment. However, they vary in size, relative orientation, and proximity to each other, e.g., while sitting at their desks, Bob's whiteboard is behind him, while Alice's whiteboard is on her right. In practice, this creates a conflict when they move from their desks to their whiteboards. Because of the different environment configurations, only one pair of surfaces can be aligned at any time. This means that if both desks are aligned and Bob moves to the whiteboard, Alice will not perceive him as walking towards her own whiteboard but to a different virtual one (see Figure 1B). This creates a challenge if Alice wants to join Bob and work together around a shared whiteboard while leveraging her own physical whiteboard.

Previous approaches for addressing this issue include warping the environments or trying to find a single optimal alignment between them. In contrast, we focus this paper on a class of solutions that received little attention in previous work—*Partially Blended*

Realities (PBR). The main idea behind PBR is to only align the elements in the physical space relevant to the task, e.g., surfaces and avatars, instead of trying to align the entirety of the physical space. This approach gives users a 1:1 mapping between the two spaces in the area near the surface used for the current task. Users can then use a different alignment between the spaces centered around the new surface as the task moves to it. Going back to our example, when Alice and Bob decide to move to the whiteboard, they can use a new alignment between the two spaces, now centered around the whiteboard (Figure 1 right).

Multiple interaction techniques can instantiate the idea of partially blended realities. The realignment between the environments can be triggered manually or automatically; it can happen instantly or through animations. To understand this design space, there are several unanswered questions: (1) How does the partial alignment of spaces affect the collaborative user experience? (2) When do users need the environments to be realigned? (3) How do users' mental models change as they shift between different alignments?

To explore PBR and address these questions, we developed a system called *RealityBlender* and conducted a user study. *RealityBlender* is a system for remote collaboration that enables users to conduct meetings in MR by co-creating a blended interaction space composed of the individual collaborators' physical surfaces in their local environments. We conducted a user study with 24 participants completing a collaborative task in pairs with two different alignment techniques: (1) *REALIGNMENT*, manually triggered by the users, and (2) *OVERLAY* with both possible alignments (around a table and a whiteboard) visible at all times. We qualitatively compared these approaches to their spatially consistent counterpart, in which participants completed the same task across spaces with the same physical arrangement.

We found that incorporating physical surfaces was effective for creating a feeling of being together, that *REALIGNMENT* worked better when triggered before than after user movement, and that users needed attention support by fading between layers in *OVERLAY*. These findings are discussed and operationalized into a design space including recommendations for how to design partially blended realities.

With this work, we make the following research contributions:

- A novel approach for distributed MR meetings—*Partially Blended Realities*—which enables remote collaborators to transition between different ways of working around physical surfaces.
- A prototype MR system—*RealityBlender*—that supports users in creating and transitioning between multiple partial alignments.
- A design space and recommendations based on a user study with *RealityBlender* on how to develop interaction techniques for navigation in partially aligned collaboration spaces.

2 BACKGROUND AND RELATED WORK

We provide an overview of solutions and theories related to the concept of blending distributed collaborative spaces.

2.1 From Blended Interaction Spaces to Blended Realities

Collaborative MR systems enable distributed users to meet virtually while retaining presence in their physical space [2, 6, 16, 20, 30, 31, 36]. The interaction metaphor in most of these MR systems is that *they bring remote collaborators into the user's physical space* by warping them into the local space [18, 38].

Recent research on distributed collaboration has explored approaches for creating *Blended Interaction Spaces* [27]. In contrast to the metaphor above, this new metaphor not only brings remote collaborators into the user's local space but also elements of their physical space that can be shared. For example, through such blending, users can sit at their own physical desks and have the sense that remote collaborators are sitting at the same desk. Often such solutions are constrained to environments specifically designed for this purpose. For example, by creating environments that are either overlaid (e.g., Holoportation [26]) or brought together with displays that act as portals (e.g., Cisco Telepresence or HP Halo [27]) creating the illusion that the two rooms are two halves blended into the same extended space.

With the global switch to remote and hybrid work, future solutions for distributed collaboration will need to support people meeting from anywhere across rooms that are dissimilar. This introduces new challenges for blending physical spaces: rooms can have different sizes, furniture, shapes, layouts, etc. Yet, with the emergence of MR headsets that can scan local environments, a new class of solutions—which we term *Blended Realities*—enable distributed collaboration across a variety of physical environments. To cope with room disparities, researchers have developed prototype systems that demonstrate various approaches to aligning distributed spaces. Solutions involve discretizing environments into functional spaces [45], warping the avatar's deictic gestures based on landmarks [44], adapting avatar movements [5, 14, 43]), and computing an optimal partial alignment [21]. We revisit these solutions in more detail in section 3.

2.2 Proxemics and Collaboration Spaces

Our work builds on proxemics, pioneered by Edward Hall [10], which denotes the study of spatial relations between people and features of the physical environment.

2.2.1 Physicality and Collaboration around Information Surfaces. A principal idea in proxemics is that we cannot consider interpersonal relations in isolation from the environment. Dimensions of interpersonal space such as proximity should be considered within a material world, where fixed and semifixed features of the physical environment (such as furniture, tables, walls) condition people's actions [10]. This idea is evidenced through CSCW studies, showing that co-located collaborators leverage physical information surfaces for organizing themselves for work. Collaborative information surfaces can be physical (walls, whiteboards, and physical tables) or digital (such as tabletops [34] or wall-sized displays [32]), and collaborative work often relies on the physicality of information distributed in the environment (e.g., [7, 22, 29]). Collaborators organize themselves in different spatial patterns known as F-formations around shared physical information surfaces [24],

use multiple information surfaces in the vicinity for juxtaposing important information resources [29], or switch between horizontal (e.g., tabletops or floors) and vertical (e.g., boards, screens, or walls) surfaces for utilizing their different collaborative affordances [7, 8, 32]. A line of MR research further motivates the perceptual benefits of harnessing the user's local environment. Studies have shown that passive haptics supports a better sense of presence in the virtual environment [13, 37], which can be used in systems for aligning physical and virtual environments [11, 37, 41]. Moreover, providing landmarks in the physical environment can enhance the ability to spatially reference shared objects [25, 44].

In conclusion, the research literature provides extensive evidence for the value of bringing physical surfaces into collaborative activities, instead of purely virtual counterparts.

2.2.2 Flexibility and Proxemic Transitions. Another insight from proxemics is that the use of space in social encounters is highly dynamic. Over time, people will organize themselves in various facing formations depending on the people, activities, and the surrounding environment [7, 8, 12, 17, 22, 35]. The concept of Proxemic Transitions [8, 19] formalizes the bodily patterns of collaboration as the transitions in the spatial configurations of people and content. Collaborators may either *move between different facing formations* around the surrounding fixed/semifixed features or *reconfigure the physical features* to optimize the spatial relations for the task at hand. It has been demonstrated how such transitions can be supported through variations of explicit and implicit interactions with physically dynamic furniture, such as interactive tabletops [9] or vertical displays [42].

For Blended Realities, it is challenging to support such dynamic use of space, because users often move around dissimilar spaces. A popular approach for collaborative virtual environments is to teleport remote users' representations to another location in the local virtual environment [3, 40]. However, as an instant transition, it often results in disorientation and the need for guidance to let users reorient themselves [3, 15]. Proposed alternatives include multiple overlaid perspectives from different locations in the same environment [33] or redirected walking [14, 39, 45].

3 BLENDED REALITIES: PROBLEM DEFINITION

Before describing Partially Blended Realities, we specify the space of existing solutions to the problem of *aligning* two dissimilar spaces such that distributed users can experience shared co-presence. An alignment is a linear or non-linear transformation (translation, rotation, and scaling) that maps points in one space to points in the other space. Alignments can be characterized as either *discrete* or *continuous* and *global* or *partial* alignment (Figure 2).

Global alignments aim to blend the entirety of both environments, i.e. every point in one environment has a corresponding point in the other. The most straightforward global alignment approach is **PHYSICAL ALIGNMENT**: to physically arrange the rooms similarly. If both environments have the same physical layout, they can simply be overlaid. This is achieved with Blended Interaction Spaces [26, 27]. The benefit of these approaches is that they build upon users' natural intuitions for where people and objects are, but they require careful design of physical spaces, and as a consequence




Figure 2: Discrete, continuous, and partial alignment for the same dissimilar environments. In this figure, the user walks a short distance in the smaller environment, but the avatar must traverse a much larger distance in the larger environment. **CONTINUOUS ALIGNMENT** approaches map every point in one environment to one in the other environment, potentially creating distortions in the user’s movement. In **DISCRETE ALIGNMENT** approaches, landmarks are conceptually mapped across environment, without necessarily matching in spatial layout. As users walk between landmarks in their physical space, their avatar teleports between them, but only upon arrival at the landmark. **PARTIAL ALIGNMENT** approaches maintain the proportions and distances of the user’s physical environment, but limit their movement in the remote environment. In the example, the user’s movement is equidistant in both spaces, but when the user reaches a wall in their local environment, there is a need for realigning them.

they cannot work as a general purpose solution. In cases where rooms cannot be physically aligned, MR solutions must virtually align them. **CONTINUOUS ALIGNMENT** solutions create a continuous linear or non-linear mapping between points in one space to points in the other, without necessarily preserving proportions, e.g., avatar arm posture warping (for consistent deixis) [44] or avatar motion adaptation [5, 14, 43]. However, these approaches can create uncanny distortions when the environments are substantially different. For instance, if one environment is twice the size of the other, every movement in that environment is twice as large in the other (Figure 2). Note that this is a simplified illustration of the approach that only warps by scaling along one dimension. Rotating to align surfaces, such as in Jo et al. [14] and Congdon et al. [5], would cause further directional distortions on the user’s walking paths or body poses. Though these approaches can be versatile in mapping different environments, Sra et al. found that the distortions created by scaling them led to a reduced sense of social presence and togetherness as compared to other mapping techniques [39].

DISCRETE ALIGNMENT solutions [45] identify points of interest in each environment and as users move between these landmarks in their physical environment, their avatars are teleported to their counterparts in the other. The benefit is that interactions are anchored around physical features that are shared in both environments (e.g., users can move between different pieces of furniture regardless of how they are arranged). However, users’ movements are not preserved until they reach the set landmarks — avatars can only teleport between them (Figure 2).

The above limitations show that global approaches are difficult to scale as the complexity of blended spaces increases with more than two distributed spaces, users, and/or surface pairs. **Partial**

alignments, on the other hand, can potentially scale with increasing complexity as they do not require globally resolving the avatar pose. Such solutions blend *parts* of both physical environments around an anchor point to create a shared area, while the region outside the area may not be fully aligned. For example, if the remote space is larger than the physical space of the user, there will be inaccessible areas (Figure 2). The advantage is that users can move around in vicinity to the anchor point without any distortion, but only to a certain extent before requiring realignment. In a partial alignment setup, if one environment is smaller than the other, one can be placed inside the other or even partially overlapped if that makes more sense for the task at hand—two conditions explored by Sra et al. [39].

Previous work on partial alignment—which we term **OPTIMAL PARTIAL ALIGNMENT**—has used optimization techniques to maximize desired properties of the alignment, such as the amount of free shared floor space [21]. In the general case, however, no single alignment between dissimilar environments will be perfect. For example, consider a case in which two collaborators work remotely, each in a room with a desk and a whiteboard, but with different sizes and configurations (Figure 3). The optimal alignment depends on around which surface the work is taking place — when working around the table, it is useful to align the tables; when working on the whiteboards, it is useful to align the whiteboards, but both cannot be achieved at once with a single isometric transformation due to the different configurations. As such, moving between the table and whiteboard creates the need for a realignment of the spaces.

This situation opens a new opportunity for partial alignment: one in which users can benefit from having multiple partial alignments




Figure 3: Partial alignment problem: One user has their desk and whiteboard on opposite walls, while the other has them on adjacent walls. This means that there are two possible partial alignments depending on whether they need to share the table or the whiteboard. When the spaces are aligned around one surface, notice that the other cannot be blended without realigning.

at their disposal and switching between them depending on the current task. We call this approach *Partially Blended Realities*.

4 PARTIALLY BLENDED REALITIES

We introduce Partially Blended Realities (PBR) – an approach to Blended Realities that enables partially blending dissimilar spaces for distributed collaboration through partial alignment of physical surfaces in MR. We now describe the concept and explore two different partial alignment solutions for PBR.

4.1 Blending Physical Surfaces

We view collaboration through a proxemics lens, considering room furniture and information surfaces as core fixed/semifixed features for conditioning spatial relations between people and content. Partially Blended Realities enables users to take virtual representations of their local physical surfaces as proxies of collaborative surfaces such as whiteboards and meeting tables. When two or more users join a meeting, the virtual representation of their surfaces will be aligned so that users appear to be working around the same surface. To incorporate fixed/semifixed features (e.g., tables, whiteboards, or walls) into the blended collaboration space, an MR system can automatically detect them or users can manually define them. When initiating the blended collaboration space, the defined surfaces are shared; the surface pair that is regarded as the *active blended surface* will be aligned such that each user sees their local active surface with the remote counterpart virtually overlaid directly on top of it. In this state, the two surfaces in the pair are considered *aligned*.

4.2 Proxemic Transitions and Partial Alignment

Our focus is on the design of partial alignment techniques to support people’s movements during collaborative work. Hereby, we explore a simplified collaborative scenario with two surface pairs: a table and a whiteboard at each location. In this scenario, the spatial layout of local surfaces at each location is dissimilar (Figure 4 Physical Setups). Collaborators can engage in *proxemic transitions* [8] where they move and reconfigure the blended space as the activities shift focus between different information surfaces. There are two classes of solutions for supporting such proxemic transitions through partial alignment; one partial alignment is shown at a time, or all are overlaid on each other.

In the *REALIGNMENT* technique (Figure 4), users can actively change how their workspaces are partially aligned as they move around their physical surfaces (e.g., tables and whiteboards) in their respective local spaces. This means only one of the defined surface pairs is initially set as the *active surface*, becoming the point-of-reference for the blended state. The realignment can then, for instance, be animated by interpolating between the start and end state of the remote space in relation to the local. This provides visual feedback to the local user that the remote user’s space is being realigned to the local space, and that this changes how they face one another. Alternatively, this transition could be instantaneous (i.e., a discrete realignment), though in our pilot studies participants found this experience jarring and disorienting due to the lack of visual feedback in the transition.

In the *OVERLAY* solution (Figure 4), the local user sees two avatar replicas of the same remote user and two replicas of each remote surface, i.e., the remote space is effectively overlaid twice on the local space in the two possible partial alignments. Seeing the same space and people from multiple perspectives is similar to the experience of *OVRlap* [33]. As both alignment layers are equally visible, the user mentally shifts their attention between the two layers of the remote user. When both users move from table to whiteboard in synchrony, they can switch focus from one avatar at the table to the other that will come within proximity to the whiteboard as the pair moves relative to their local surfaces.

5 REALITYBLENDER

To explore the design opportunities and challenges around PBR, we developed *RealityBlender* – an MR system for distributed collaboration. We use RealityBlender to explore the design of partial alignment techniques through collaborative application scenarios.

5.1 Implementation

We implemented RealityBlender using Unity3D for the Meta Quest 2. The prototype integrates the Oculus XR Interaction Toolkit for hand tracking and interaction, the Meta Avatars SDK for the rendering of avatars, and the Photon Pun 2 framework for multi-user networking. Below we elaborate on its features.

5.1.1 Defining local surfaces and matching distributed surface pairs. Users can turn their physical surfaces into virtual planes. Before

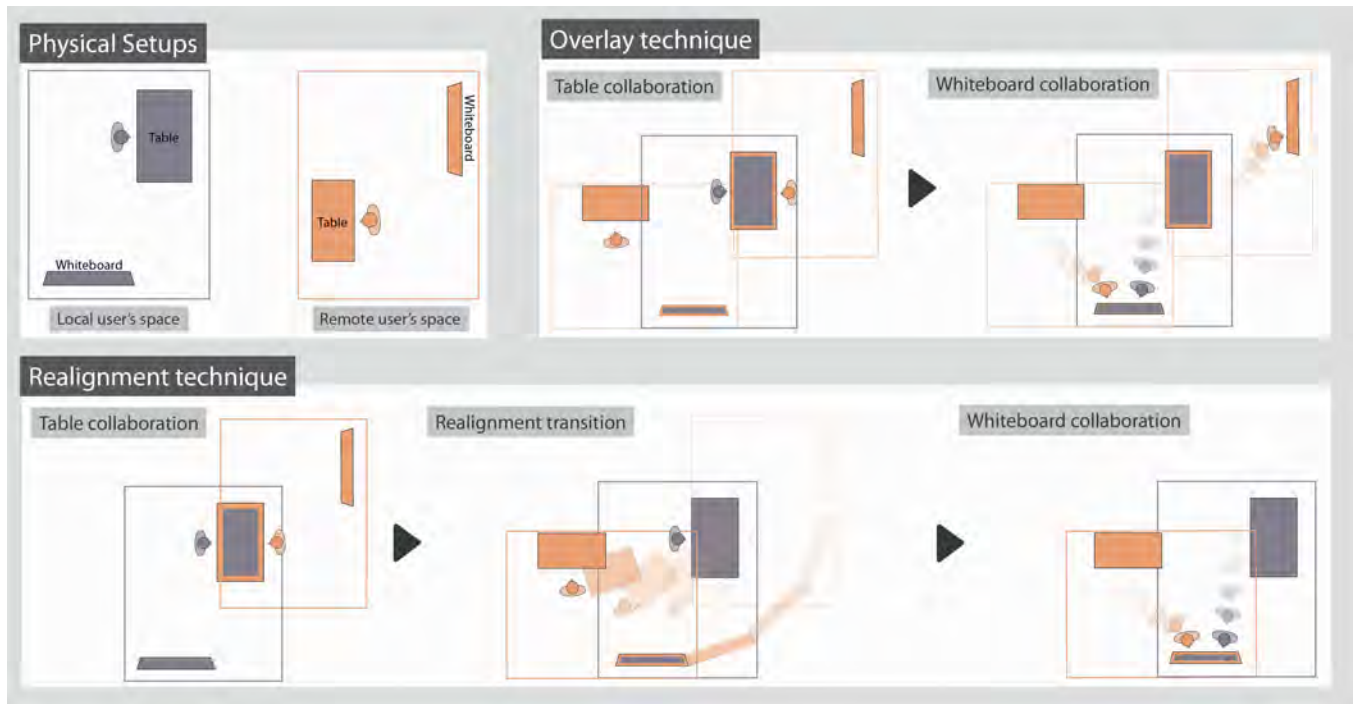


Figure 4: The alignment problem for two rooms with dissimilar physical setups and our novel partial alignment techniques, REALIGNMENT and OVERLAY. Grey colors are used for the local and orange for the remote spatial layouts.

defining the surface in the physical space, the user first selects from a menu whether they want to define a horizontal or a vertical plane. The surface is then defined in the physical space by pinching the bottom corners of the surface (to set the width of the plane) and by sliding their finger (to set the height)—Figure 5. The virtual plane is initially private, i.e., it is only visible to the local user, but can be shared with remote users later on.

Users can, in principle, define as many surfaces as they want in their physical space to share in the blended meeting space (only limited by computational power of the headset or visual clutter in the user interface). Once the meeting connection is established, the first horizontal surfaces defined in each space are matched (and aligned) as the *initial active surface pair*. Everything else is placed in reference to these surfaces.

5.1.2 Aligning planes of a distributed surface pair. Because surfaces may have different dimensions (width, height, orientation, etc.), there must be a rule to determine how two distributed surfaces are positioned, rotated, and scaled in relation to each other. RealityBlender applies a simple rule: two surfaces of different dimensions are aligned around their center. The real-world scale of each surface is preserved. This means that surfaces of different sizes will not fully occlude each other when aligned, and this is visualized via the surface outlines (Figure 5). It is then up to the application to define how shared content is laid out (e.g. bound within or overflowing edges). Regarding rotation, the system assumes that users sit at the side of the horizontal surface from which they pinch in the corners nearest themselves. In this way, the surface pair is rotated such that users initially appear to be sitting face-to-face on opposite sides.

5.2 Application and Use Cases

We developed an application to explore how Partially Blended Realities may enable new forms of MR collaboration with multimedia content. For this prototype, we focused on enabling interaction with virtual content by anchoring it to physical surfaces. Specifically, RealityBlender supports manipulating *multimedia sticky notes* that contain images or videos and *virtual sketching* directly on the physical surfaces (Figure 6). Users sketch on physical surfaces by pinching while touching the surface (as if holding a virtual pen), which draws a line until the user releases the pinch. Pinching is used for grabbing and placing sticky notes; when the hand is within close proximity to the surface, it snaps to the surface canvas. The application requires users to explicitly switch modes to disambiguate between pinching to grab and pinching to draw. Finally, the application must be integrated and distributed within a partially blended space (i.e., where local surfaces change between being aligned with their remote counterpart and being misaligned when the other is aligned). In our current implementation, sketches and sticky notes are positioned and arranged within a 2D canvas on top of the surface planes. The 2D canvases are synchronized across each surface replica (i.e., one synced canvas for the whiteboard and one for the table). This application enabled us to explore scenarios of collaborative work, such as ideation and brainstorming (with sticky notes and sketching) or sense making tasks (with prepared visual cards). For the user study, the application was adapted to a collaborative game with visual cards. Next, we will describe the user study and how the application was used in a remote collaboration task.




Figure 5: RealityBlender: how to start a meeting in PBR.



Figure 6: Multimedia collaboration in a workspace blended across two remote office rooms. RealityBlender supports virtual sketching on surfaces and moving digital sticky notes on and between surfaces.

6 USER STUDY

To better understand the user experience of our partial alignment techniques, we conducted a user study in which participants completed a collaborative MR task across a table and a whiteboard distributed across two different spaces. Through a qualitative analysis, we aim to explore the following research questions:

- RQ1 How does partial alignment affect the user experience? We investigate this question by prompting users to compare their experiences of partial alignment to a condition with optimal physical and virtual alignment.
- RQ2 When and how do users trigger realignment? We investigate this question by analysing different instances of users triggering realignment and how they consequently report on the experience of the REALIGNMENT condition in the interviews.
- RQ3 How do users manage and shift their attention when presented with multiple possible partial alignments? We investigate this question by asking users how they experience and make sense of the OVERLAY condition.

6.1 Conditions

- **REALIGNMENT:** Participants switch between which surface the environments should be aligned to with the press of a button accessible in multiple locations (located on the whiteboard, the table, and on the user's arm).
- **OVERLAY:** Both partial alignment layers are equally visible at all times (including the remote collaborator's avatar). Hence, the user can implicitly decide which partial alignment to pay attention to.
- **PHYSICAL ALIGNMENT:** A "perfect" baseline condition where the physical environments have the same configuration so that both surface pairs are virtually aligned at all times.

The conditions provide a blended collaboration space for users to switch focus from the whiteboard to the table, and back. A difference across conditions is in the spatial setup. The **PHYSICAL ALIGNMENT** condition is designed with a physically consistent setup, whereas **REALIGNMENT** and **OVERLAY** conditions are situated in a spatially inconsistent setup (Figure 7). The dissimilarities between the two spaces were controlled to be minimal yet still causing significant rotations on the blended environment during realignment. Based on pilot tests with different physical layouts, our experimental setup was chosen to present a range of interesting spatial differences via relative surface rotations without having to physically rearrange the space in every condition. The decision to reduce the need for rearranging between conditions was to draw focus toward experiencing the techniques, rather than learning new spatial layouts.

6.2 Task

Because the focus of the study is on how participants complete tasks collaboratively across multiple surfaces, we designed a task with four phases (Figure 8) that requires specific movements by participants: participants initially work around the table (A), then one participant moves up to the whiteboard (B), then the other participant joins (C), and finally, they go back together to the table (D). This allowed for observing a range of different proxemic transitions, with variations of individual (A to B to C) and joint transitions (C to D) between surfaces.

For this purpose, we adapted the board game *Mysterium* and incorporated it into our MR application. In this game, one player takes the role of a ghost who is trying to help a psychic (the other player) to find the person responsible for their murder as well as the location it happened. However, the ghost can only communicate with the psychic through visions in the form of picture cards. This game was chosen because it encourages non-verbal rather




Figure 7: Physical experimental setups for the three conditions. The thick borders indicate walls, which constrained the participants' local movements. *Similar Spaces*: for PHYSICAL ALIGNMENT, rooms were physically arranged for fully blended spaces. *Dissimilar Spaces*: for users to clearly experience the dissimilarity, the setup for REALIGNMENT and OVERLAY was constructed to create the appearance of users going in separate directions when moving from one surface to the other. While the room dissimilarities were minimal (only relative rotational differences), the transitions between alignments caused significant rotations (as illustrated in Figure 4).




Figure 8: Instructions to participants were designed to investigate a selection of the possible proxemic transitions that may occur; we selected the subset of transitions that are initiated from movement by one or both in the pair.

than verbal communication. Six cards with pictures of potential murderers and six with potential murder locations spawned face up on both the table and the whiteboards in the application. Out of the 42 vision cards, seven random cards spawned face down in front of the ghost player, as these should not be visible to the psychic until the ghost has chosen the vision(s) they want to share. To elicit movement between the surfaces, we modified the rules of the game so that both players started at the table while the ghost chose the vision card and gave it to the psychic. Upon receiving the card, the psychic had to bring it to the whiteboard and select a suspect. They would then be joined by the ghost, who listened to their reasoning for picking this suspect. The ghost then explained their own rationale for picking that particular vision card and they both returned to the table. Through this structure, we replicate steps A–D described above.

6.3 Procedure and Participants

To induce dynamic spatial behavior where participants frequently move between surfaces, we designed the spatial layout of surface content such that it would allow for movement between surfaces

according to the task. We recruited 12 pairs of participants ($N=24$, 16 male, 8 female) with an average age of 28 ($SD=5.4$) from the local university campus and nearby companies. Upon arrival, participants listened to a brief introduction to the PBR concept, the study procedure, the rules of the game, and signed an informed consent form. The pair rated their familiarity with each other (from “not” to “very” familiar) and their individual experience with MR (from “no experience” to “expert”) on a 5-point (0–4) Likert scale. The average familiarity score in the pairs was 2.3 ($SD=1.6$), ranging from acquaintances to close friends, and the average MR experience score was 2.2 points ($SD=1.3$).

The three conditions were conducted across two rooms in the same lab but with walls that separated the pairs to create the experience of collaborating remotely in different local spaces. The sequence of actions in the task was controlled by a facilitator who instructed the pairs verbally and participated through RealityBlender via a third headset. The facilitator could see the participants from above, while the participants could only hear the voice of the facilitator giving instructions. Before each condition, the pair had a short training phase where they practiced the mechanics of the interactions relevant to the given condition. The conditions were counterbalanced to reduce the carry-over effect when interviewing pairs and prompting them to compare the user experience of the different conditions. Between conditions, the rooms were rearranged as specified in Figure 7, depending on whether the next condition (in the counterbalanced order) involved a similar or dissimilar setup. In each condition, the pair went through two rounds of the game (one playing as the ghost, and one as the psychic). With three different conditions, the pair played the game six rounds in total. After every condition, each individual user was given a custom co-presence questionnaire asking them to rate their agreement to four different statements on a 5-point Likert scale (Figure 9). Finally, after experiencing all three conditions, a semi-structured interview was conducted with the pairs in a focus group to stimulate discussion about the experience of each condition. The prepared questions

were designed to illuminate different aspects of the research questions, but experimenters supplemented them with spontaneous questions based on observations of interesting incidents during the deployed conditions.

6.4 Data Collection and Analysis

The data material for this work is primarily qualitative. Focus group interviews were recorded and transcribed. We video-recorded participants with a stationary camera in each room and screen-recorded streams from the Quest 2 headset displays. Notes were taken of interesting observations during game sessions. In the analysis phase, the transcribed interviews served as the primary data material for grouping data into themes related to the user experience. Two researchers split the data and coded quotes from participants while discussing code tags. Although prepared questions provided a structure for the interviews, they were not the basis for the analysis. Instead, participant quotes were coded independently of the questions to derive categories. During axial coding, the two researchers collaborated on grouping similar codes together into 19 categories. The individual codes were marked as either general comments or with which study condition the comment related to (REALIGNMENT, OVERLAY, or PHYSICAL ALIGNMENT) for later comparisons. From these categories, four themes were created. Quotes under each theme were then structured in relation to the research questions, RQ1-3. As the data from our co-presence questionnaire (Figure 9) did not show any clear patterns, our analysis focused on the focus-group interviews and video recordings.

7 STUDY FINDINGS

The analysis led to the themes **Asymmetry** (how PBR introduces asymmetries in the partners' experiences), **Control** (users' sense of control), **Mental Model** (how users formed mental models of PBR), and **Surface Ownership** (how physical surfaces were involved in the social interactions). Themes are based on quotes from the focus group interviews, supplemented with instances of user interactions from room and headset recordings. We structure the findings to answer our research questions RQ1-3 in terms of the four themes.

7.1 Blended Interactions Around Physical Surfaces Were Effective

As an answer to RQ1, a recurring feedback (related to the **Mental Model** theme) that goes across all conditions was that pairs reported feeling like they were together while at the surfaces but not when in-between surfaces. While there were also lots of comments about the perceptual challenges of navigating in the transitions between surfaces, the general feedback was that the sense of co-presence worked as long as the pair remained in the vicinity of the same blended surface. E.g., several participants referred to this aspect as "feeling natural", such as P8B.

P8B: "It felt natural when sitting together at the table."

This is further indicated by an observation of a playful exchange (Figure 10) with bodily communication via the table proxy: P10B is teasing P10A by leaning around the corner of the blended table to cheat by peeking at the cards that were oriented away from them. Thus, the passive haptics from the surface in the local environment




Figure 9: Results from our custom co-presence questionnaire (N = 24). It compares the three conditions regarding the individual user's perception of the shared blended space (1st and 2nd question) and the perception of user movement within it (3rd and 4th question).



Figure 10: Participant P10B playfully cheating in the game by going around the table to look at P10A's cards.

coupled with the participant seeing their partner's avatar touching the same surface seemed to be an effective illusion.

As the transitions between surfaces in the REALIGNMENT and OVERLAY conditions often broke this illusion, participants (not surprisingly) made comparative statements in favor of the PHYSICAL ALIGNMENT condition.

P9B: "[The REALIGNMENT and OVERLAY conditions] felt like being in the same virtual environment, and [the PHYSICAL ALIGNMENT condition] felt more like the physical environment."

Only for the PHYSICAL ALIGNMENT condition, a participant made the following comment.

P11B: "When I took my [headset] off, I was like 'Where are you!' I felt like she was still there."

These quotes are indicators that the sense of co-presence is affected, to some extent, by the partial alignment techniques when comparing to working across physically aligned spaces.

7.2 The User Experience of REALIGNMENT

In the REALIGNMENT condition, users press a button to trigger realignment (either on the hand menu, the whiteboard, or the table). We relate the themes to the questions of their general experience (RQ1) and specifically when users wanted to realign (RQ2).

7.2.1 Asymmetry: Realignment Requires Users to Reorient Themselves. Due to asymmetries in the user experience of realignment, the transition would cause confusion or startle users as they

would suddenly end up in each other's personal space. An illustrative example is seen in Figure 11. Because RealityBlender is an MR experience, the user sees their real space with the remote space as an overlay, and only the remote space and avatar are visually moving during realignment, while the real world remains intact. As a consequence, each participant perceives the other as moving towards them. The asymmetry was revealed in how participants described the intention to realign. One participant expressed jokingly that they could control the partner's movements.

P6A: *"Come over here, get over there."*

But it was only in the interviews that P6A discovered that the partner sees the opposite happen.

7.2.2 Control: A Shared Button May Aid Negotiation of Control. The question of *when* to trigger realignment is highly dependent on the social dynamics. Several pairs started developing social protocols for negotiating the control of the realignment trigger. For instance, the ability for both participants to control the trigger facilitated assisting each other in the transition.

P6A: *"I was holding a card, and switching would be more comfortable. I felt like I could not touch a button because my hand was busy. [...] But I asked P6B to press the button since I could not."*

On the other hand, some participants expressed concerns around the control that the partner had over *their* space as well as the control that they had over their partner's space. This sometimes resulted in disorientation about where their partner was in the shared space. P12B found that the design choice of having individual buttons impeded the sense of shared awareness of action, which made it difficult to decode the partner's intention.

P12B: *"Maybe if it's a shared button, then you can see when someone presses. It will make for a greater shared space. If I know someone is near the button I can prepare and when someone reaches for it I know it's gonna happen."*

7.2.3 Mental Model: Proactive Triggering Works Best. While participants' general comments indicated that the surface-focused collaboration worked, there were several issues with the movements between surfaces. The variety of ways users triggered realignment gives us an impression what works and what does not (RQ2). In most cases, the participants would trigger the REALIGNMENT transition *after* reaching the destination surface.

P3A: *"I didn't realize we were in different locations. To show him the card, we needed to blend to same location."*

This often led to situations where participants looked back at the partner's avatar and triggered the transition on the button. E.g., Figure 11 shows an incident where a participant looked back and then was surprised as the avatar and surfaces rotated towards them. In some cases, we observed that the remote avatar moved *through* the participant, which led to participants' personal space being intruded.

P9B: *"I was not happy that she transitioned through me to get to the whiteboard."*

In a response to these issues, some participants learned reactive or proactive strategies for realigning. E.g., a participant realized how to interact with the system in response to the partner moving.

P9B: *"Her avatar was going away from me. If I press, she would go back to me. It kind of feels like she is coming instead of suddenly being there. It's definitely nice, better than teleporting"*

Some were even more proactive and triggered the REALIGNMENT transition before walking to the destination surface.

P12A: *"I took the card, and on my way to the whiteboard, I knew we were gonna end up there, so I pressed and when I arrived she was there."*

Our observation across sessions was that the proactive instances like the above (i.e., triggering before moving) seemed to provide the smoothest experiences with realignment.

7.2.4 Surface Ownership: Aligned Feels Shared, Misaligned Feels Private. Whether the surfaces are aligned or not seemed to affect the perceived ownership of the surfaces. When describing surfaces in a misaligned state, the participants often referred to them as either their own or their partner's surface.

P9B: *"[...] when she walks to her whiteboard she doesn't go to mine. But I needed [...] her to come to me."*

When surfaces in turn were aligned, participants in the focus groups almost exclusively referred to it as "the whiteboard" or "the table". Based on these perceptions, several pairs speculated on how surface ownership could be configured, indicating that there may be a benefit to the separation of private and shared surfaces.

P9B: *"Being able to control it is also nice. If she just walks to her whiteboard to reorganize her notes and I'm still here, I don't want it to suddenly rotate."*

P6B: *"The REALIGNMENT condition would have [supported] to work separately and then work together on the result."*

These suggestions point to an interesting direction for PBR, where the purpose is not merely to solve the alignment problem, but rather to allow for configuring the blended space together.

7.3 The User Experience of OVERLAY

In OVERLAY, users see both partial alignments at once with the remote user's avatar appearing twice. Themes are related to this experience (RQ1) and how users manage their attention within it (RQ3).

7.3.1 Asymmetry and Control: Difficult to Manage Visual Attention. In OVERLAY, several participants appreciated that they could implicitly switch their attention and did not have to explicitly press a button when they wanted to go between the surfaces.

P9B: *"I like how OVERLAY is kind of an automatic switch. It's like it's in your head the switch is instead of software or button. You need to figure out now you switch to this avatar."*

While some preferred mentally switching themselves, the general feedback was that overlaid was confusing to navigate. Participants often took some time after walking to another surface



Figure 11: REALIGNMENT may interfere with personal space: P3B presses the button on their whiteboard. P3A transitions from their table to be close to P3B's face. They both appear to be startled by the intrusion of their personal space.



(a) Recording from P3B's POV



(b) Recording of P3B's physical space.

Figure 12: OVERLAY: P3B stands near his whiteboard and looks at P3A bringing a card to their whiteboard. P3B jumps back and makes a short squeak sound as P3A appears right next to him at the whiteboard.

to reorient themselves, and pairs frequently struggled to achieve mutual gaze with the right avatar replica.

P11B: “I felt I knew when I should look at you and at what place. I knew you were standing at the whiteboard and should look at the whiteboard. I feel like I talk to you and you stand next to me, but you are looking at the one standing away.”

In response to this challenge, several pairs discussed (in the focus group) how the system could aid the user in managing their visual attention. When prompted to suggest how it could be improved, several proposed the idea of using transparency on one layer and then fading in and out layers depending on where the user's attention is.

P5A: “The avatar we are working with is fully opaque and highlighted. The other one can be a bit transparent.”

When prompted to consider which replica layer should be shown, pair 3 were in agreement.

P3B: “Only show the closest avatar.” (**P3A:** “Yeah.”)

7.3.2 Mental Model: “I Think You Could Get Used To It”. A general impression with OVERLAY is that there seemed to be a quick learning curve, with initial confusion and then later learning how to cope with the duplicated overlays. Several had initial experiences that led them to be momentarily startled, such as turning the head and suddenly realizing they were standing too close to the other avatar (Figure 12). Subsequently, we could observe that participants often learned from their experiences and next time were more aware of the mental switch during the transition to another surface.



Figure 13: OVERLAY: P2B instantly shifts attention from the cloned whiteboard surface to their local real surface, as P2A gets up and walks to their whiteboard.

P11B: “[...] I knew when I should look at you at what place. I knew you were standing at the whiteboard and should look at the whiteboard. [...] I think you could get used to it.”

Participant P6A mentioned that OVERLAY helped them understand their partner's physical space better. Participant P12B mentioned that the cartoonish avatars representing the partner made it easier for them to accept that there are two representations.

P12B: “I found it surprisingly good. You would expect to have some spatial awareness that the person can only be at one place. But you become a bit more aware that they are an avatar, but they are also cartoonish. You can just say ‘goodbye’ to one and ‘hello’ to the next.”

The mental distinction between “my surface” and “my partner's” (**Surface Ownership**) seemed to be evident in the OVERLAY, similar to the REALIGNMENT experience. The replicated surfaces and avatar seemed to help them realize how they should navigate by switching focus from the virtual overlay to the surfaces in their local environment (c.f. Figure 13).

P11A: “I felt I wanted to follow you walking up and put the card on your whiteboard. But then I remember I have my own whiteboard.”

P2B: “I would only speak to the one at my surface.”

8 DISCUSSION

In the following, we return to our research questions to outline implications for designing Partially Blended Realities.

8.1 The User Experience of Partial Alignment

Starting with **RQ1**, it was clear that participants found the surface-based interaction an effective illusion for creating a blended interaction space. But comparing the two partial alignment conditions to the **PHYSICAL ALIGNMENT** condition, it is also clear that the illusion




Figure 14: Showing the different flows depending on which button was used for triggering realignment. Blended Origin: When triggering *after* moving. Blended Destination: when triggering *before* movement.

can break once it comes to the transitions where one or both users walk between the table and the whiteboard.

The PHYSICAL ALIGNMENT condition represents a global alignment solution. Thus, our study results provide a starting point for discussing trade-offs between PBR and global approaches. With only two surface pairs across spaces with minimal dissimilarities, the benefits of PBR may not be immediately clear; in such simple cases, global solutions that *continuously* adapt avatar poses (such as [5, 14, 43, 44]) are relatively straightforward. However, avatar adaptations may still cause uncanny distortions. Consider a dissimilar setup where the remote user A has the two physical surfaces close to each other, and the local user B has them spaced far apart. Consider then A standing between these two surfaces and pointing back and forth between them. While the avatar adaptation may be correct, B will consequently perceive A's minimal arm movements (between two nearby surfaces) as large uncanny pointing gestures by the avatar (between two surfaces far apart). With PBR, on the other hand, avatars are not distorted. Instead, pointing gestures and gaze may be directed to virtual surfaces that are not currently aligned. Although we found this to cause confusion, the confusion also decreased over time as users learned to mentally comprehend the reason for these disparities.

PBR and global *discrete* solutions, such as Yoon et al. [45], both maintain undistorted space for users in the vicinity of landmarks such as physical surfaces. However, they differ in how avatars transition between landmarks. Yoon et al.'s technique *teleports* the avatar when the remote user reaches the landmark (i.e., the avatar disappears from one location and reappears in another). In contrast, PBR shows the avatar's pose transformed into another reference frame, either through an animated transition (REALIGNMENT) or through mental switching to the other reference frame (OVERLAY).

While this work has allowed us to highlight these differences between PBR and prior approaches, future work is needed to better understand how these techniques compare in terms of co-presence and collaborative performance.

8.2 Control and Scalability of PBR

As we have illustrated in our problem definition, it is an open challenge how to design scalable solutions for blended realities. Thus, the more significant impact of PBR (compared to prior approaches) may be its potential for *scalability*, and we regard this as an interesting avenue for future research. What happens when distributed

teams come together with several local and remote users using more surfaces than one table and one whiteboard? With increased numbers of multi-surface and multi-user environments, the REALIGNMENT and OVERLAY techniques will be increasingly difficult to comprehend and navigate. Thus, to further this research direction, we offer recommendations for how to address their scalability issues based on our findings.

For RQ2 (the realignment trigger), the results encourage the idea of explicit input for REALIGNMENT. As there were several challenges with negotiating the control, we recommend redesigning how the trigger button is manifested in the blended space. The study showed that pairs triggering *after* getting to their destination were often confused as they did not arrive at the same surface (Figure 14 Blended Origin). The pairs who triggered *before* moving to the other surface had a smoother transition experience as they would convene at the same spot (Figure 14 Blended Destination). In scaling the REALIGNMENT technique to more than two surfaces, a technique for transitioning from the origin surface would then need to take into account that there are multiple potential destination surfaces. For RQ3 (mentally shifting between alignments), we found that participants initially struggled to navigate between layers in OVERLAY, but that most participants learned how to cope with this condition within a short period of time. However, for each additional surface, the current OVERLAY technique would spawn one additional avatar, which clearly does not scale well. Thus, it is worth considering either to provide techniques for users to explicitly navigate between layers such as OVRlap [33] or, as suggested by several of our participants, that the system fades between layers based on implicit input.

8.3 Aligned vs. Configurable Spaces

While this work focuses on the problem of aligning distributed spaces, the PBR concept invites thinking about how to go beyond perfectly aligning spaces. In the current version of RealityBlender, the simple alignment rule (i.e., to align surfaces around the center and preserve real-world scale) constrains the alignment of surfaces to merely consider them as pairs that act as physical proxies of the same surface. However, our study revealed the potential for a broader design space of possibilities. Across both REALIGNMENT and OVERLAY, discussions often emerged around the idea of using surfaces that are not aligned as private surfaces. This alludes to a potential expansion of the PBR design space, where the alignment problem is reframed as enabling users to harness the partial blending of spaces for *reconfiguring* the blended space to support different modes of collaboration—akin to prior work on shape-changing collaboration spaces (e.g., [9, 42]).

For this paper, we focused on the subset of blending possibilities that *align* surfaces; what in Marquardt et al.'s multi-surface design space is regarded as *stacked* [23]. However, there is potentially a rich design space for how surfaces (and people) may be arranged, following the broader set of spatial dimensions laid out by Marquardt et al. [23]. RealityBlender could be extended to allow for more complex configurability, which has great potential for improving its scalability beyond merely pairs of users, surfaces, and spaces. For instance, the system could enable users to rescale, move, and rotate remote virtual surfaces in relation to local surfaces to allow

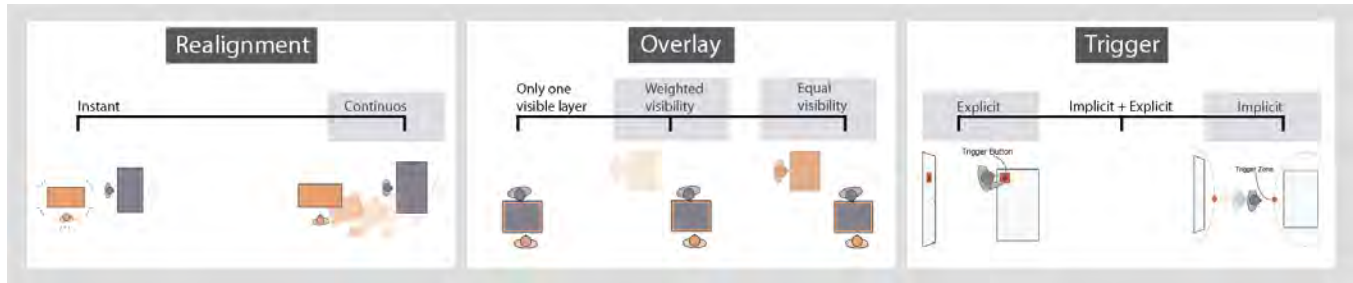


Figure 15: PBR Design Space: Executions (REALIGNMENT and OVERLAY) and trigger mechanisms (explicit, implicit, or hybrid) for partial alignment techniques.

for more expressive forms of blended realities that are not constrained to simply aligning surfaces. However, this direction poses new challenging questions, such as: What happens to the avatars? Should spatial relations in remote spaces always be preserved?

8.4 Limitations and Future Work

Throughout the paper, we have consistently focused on a single distributed spatial setup with two tables and a whiteboard with dissimilar relative rotations. It remains to be technically evaluated to which degree this concept scales, i.e., how many surfaces the system can handle before it starts affecting the headset performance. But more importantly, our user study results reflect only the particular spatial experience and asymmetries, which are incurred by this setup. The individual rooms were arranged with walls (thick borders in Figure 7) to cause intentional physical constraints on the user’s local movement; the constraints made it impossible for participants to physically walk to the virtual surfaces when they were misaligned with the local surfaces. This was to engineer a “right” and a “wrong” way to walk around, because we wanted to control that the dissimilarity caused participants to realign to better utilize their local space. However, this setup was simplified, controlled, and far from the complex dissimilarities that are likely to occur naturally in real-world office spaces, and conducting the same study in another experimental setup may reveal other trade-offs. Hence, more systematic exploration of different blended setups may provide a more generalizable and more ecologically valid account of the pros and cons.

Moreover, the study task was specifically facilitated (with instructions from the facilitator) to require that visual cards were on one surface and had to be moved to another. This was decided to induce frequent proxemic transitions where users move between their physical surfaces individually or in unison. The unnatural frequency of transitions between surfaces certainly reduced the realism of the social dynamics we observed. Moreover, the transitions with REALIGNMENT and OVERLAY had a significant impact on the participants’ overall impression of the conditions and this might not have been so strong if there had been fewer transitions. However, the high frequency was a deliberate choice to make sure that they occurred, because, in more realistic collaborative tasks, such transitions may otherwise be rare.

While this work focused on qualitatively assessing the effects of transitions on the user experience, there are many interesting questions for which future work could analyse the problem more

quantitatively. E.g., conditions could be compared in terms of the time spent on tasks, the amount of movement in space, and the time the pair spent together and separated. Navigation performance could also be measured across techniques by estimating how long it takes for participants to get back on task after a transition. When quantifying user performance, it then becomes especially relevant to account for learning effects across conditions. While we counter-balanced the conditions, we did not find any indication that users improved across conditions in navigating PBR as they got more familiar with the local and remote spaces. We believe this is due to the users being unfamiliar with both local and remote spaces and anticipate that if users work in their familiar environments, the resulting user experience would be less confusing with better performance in navigating PBR. Future work could study this in a more realistic setting where real colleagues work across environments that are familiar to them.

Despite these limitations, our results revealed several important implications for designing PBR. Next, we synthesize these implications into a design space for PBR.

9 EXPANDING THE DESIGN SPACE OF PBR

Incorporating the implications from our discussion, we build upon our two initial techniques to expand their design space (Figure 15).

9.1 Redesigning the REALIGNMENT Trigger for Shared Interaction

As we have discussed, the pairs that could navigate in REALIGNMENT most easily were the ones that triggered realignment prior to moving to the destination surface. Thus, we recommend that PBR should only anchor realignment trigger buttons at the *origin surface* rather than at the destination surface of a proxemic transition. This way, users can better get used to the new blended state before moving towards the newly aligned surface (Figure 14). Moreover, as suggested by P12B, we recommend making it a *shared* trigger button. This further aligns with the placement choice at the origin surface as it will be aligned such that both users can have the button at the same spatial reference point (Figure 16 Shared Button).

9.2 Enhancing the Scalability of OVERLAY

In response to the confusion and scalability challenges with OVERLAY, we recommend considering a redesign of the technique that incorporates a weighted fading of the different partial alignment



Figure 16: Left: A design reiteration of REALIGNMENT that proposes a shared button. Right: A comparison between OVERLAY with equal visibility and weighted visibility.

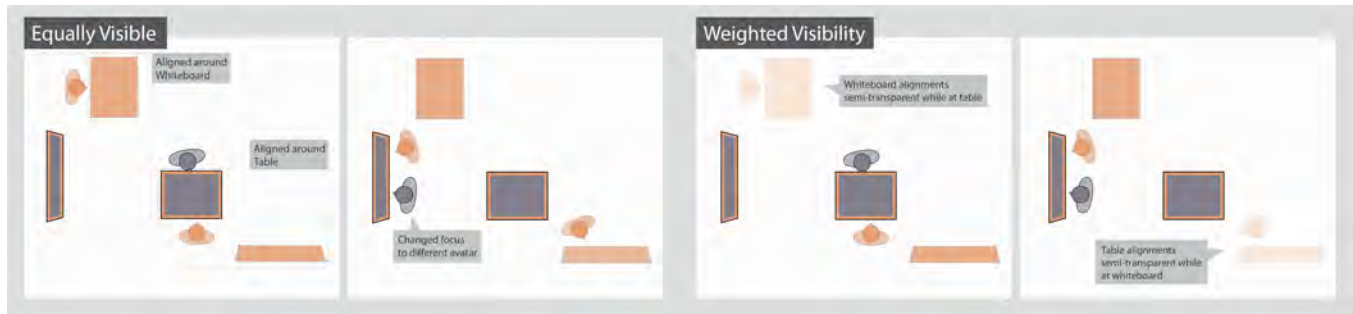


Figure 17: OVERLAY techniques. Equally visible overlays vs. weighted visibility with fading based on user input.




Figure 18: Explicit vs. implicit input. Left: Our explicit techniques rely on buttons to trigger transitions. Right: Implicit techniques rely on proximity thresholds that can trigger either when users enter or leave the area near a surface.

layers. As an alternative to rendering remote replica layers with equal visibility, OVERLAY techniques can incorporate a weighing mechanism of the different layers (Figure 17). This is illustrated with RealityBlender in Figure 16 (Equally Visible vs. Weighted Visibility). In the other end of the continuum of the design space, one layer is fully opaque and the other fully transparent. This experience is then similar to REALIGNMENT techniques, given that only one blended state is visible at once.

9.3 Explicit vs. Implicit User Input

For each of the above techniques, there are several design possibilities for how to trigger interface transitions (e.g., realignment or fading). We consider design aspects of explicit vs. implicit input, and how to design these for scalability. To allow for the techniques

to scale to more complex multi-surface environments (beyond one table and one whiteboard), we propose that the space is divided into discrete zones for triggering (akin to prior work on DISCRETE ALIGNMENT [45]). Figure 18 illustrates how these trigger points could be placed. However, in the case that the user can walk to multiple different surfaces, only the implicit trigger point at the destination surfaces will allow for the system to disambiguate how the interface should transition (e.g., to which destination surface should it realign, or which layer should be faded in). For the explicit trigger point, there would need to be several buttons at the origin surface for selecting between multiple destination surfaces.

In our design proposal for the fading trigger in OVERLAY, the fading is based on which local surface the user is in closest proximity to (Figure 17). Hence, when users collaborate around the whiteboard, the avatar near the whiteboard will be opaque and the other one faded (and vice versa for the table). While surface proximity is a good indicator of where visual attention may be for PBR, it is also a crude generalization and not always the case (e.g., you may be close to the table but pointing to something on the whiteboard). Hence, future work could investigate what might be the best input methods and modalities for interacting with fading.

10 CONCLUSION

As the world changes to new hybrid forms of work, where people are working from anywhere, collaborative MR solutions need to be adaptable to different environments. We have introduced Partially Blended Realities (PBR) – a novel class of solutions to partial alignment of dissimilar spaces for distributed MR meetings. We have

developed two types of partial alignment techniques and built a prototype system, RealityBlender, to implement and demonstrate two solutions to PBR; REALIGNMENT and OVERLAY. In a user study, we have studied the user experience of the two techniques in a remote collaboration task. Incorporating findings from the user study, we have expanded on the design space of PBR. We have discussed reiterated design proposals of our two techniques, along dimensions such as how explicit control is spatially configured, how to design for weighted visibility of overlays, and trade-offs in explicit vs. implicit triggering of transitions in the partial alignments. The main takeaway from this work is that blending physical surfaces as in PBR is an effective solution for enabling MR meetings across remote and dissimilar spaces, pointing to a rich design space for supporting users in blending collaboration spaces for the future of hybrid work.

ACKNOWLEDGMENTS

This work was supported by the European Research Council (ERC) under the European Union Horizon 2020 research and innovation programme (grants no. 740548 CIO, and no. 101021229 GEMINI) and by Aarhus University Research Foundation (grant no. AUFF-F-2021-7-2).

REFERENCES

- [1] Jakob E. Bardram and Claus Bossen. 2005. Mobility Work: The Spatial Dimension of Collaboration at a Hospital. *Computer Supported Cooperative Work (CSCW)* 14, 2 (2005), 131–160. <https://doi.org/10.1007/s10606-005-0989-y>
- [2] Mark Billinghurst and Hirokazu Kato. 1999. Real World Teleconferencing. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania) (CHI EA '99). Association for Computing Machinery, New York, NY, USA, 194–195. <https://doi.org/10.1145/632716.632838>
- [3] Doug A. Bowman, David Koller, and Larry F. Hodges. 1997. Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. In *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*. IEEE, New York, NY, USA, 45–52. <https://doi.org/10.1109/VRAIS.1997.583043>
- [4] Michael Broughton, Jeni Paay, Jesper Kjeldskov, Kenton O'Hara, Jane Li, Matthew Phillips, and Markus Rittenbruch. 2009. Being Here: Designing for Distributed Hands-on Collaboration in Blended Interaction Spaces. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7* (Melbourne, Australia) (OZCHI '09). Association for Computing Machinery, New York, NY, USA, 73–80. <https://doi.org/10.1145/1738826.1738839>
- [5] Ben J. Congdon, Tuanfeng Wang, and Anthony Steed. 2018. Merging Environments for Shared Spaces in Mixed Reality. In *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology* (Tokyo, Japan) (VRST '18). Association for Computing Machinery, New York, NY, USA, Article 11, 8 pages. <https://doi.org/10.1145/3281505.3281544>
- [6] Scott Elrod, Richard Bruce, Rich Gold, David Goldberg, Frank Halasz, William Janssen, David Lee, Kim McCall, Elin Pedersen, Ken Pier, John Tang, and Brent Welch. 1992. Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations, and Remote Collaboration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Monterey, California, USA) (CHI '92). Association for Computing Machinery, New York, NY, USA, 599–607. <https://doi.org/10.1145/142750.143052>
- [7] Jens Emil Grønbaek, Mille Skovhus Knudsen, Kenton O'Hara, Peter Gall Krogh, Jo Vermeulen, and Marianne Graves Petersen. 2020. Proxemics Beyond Proximity: Designing for Flexible Social Interaction Through Cross-Device Interaction. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376379>
- [8] Jens Emil Grønbaek, Henrik Korsgaard, Marianne Graves Petersen, Morten Henriksen Birk, and Peter Gall Krogh. 2017. Proxemic Transitions: Designing Shape-Changing Furniture for Informal Meetings. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 7029–7041. <https://doi.org/10.1145/3025453.3025487>
- [9] Jens Emil Grønbaek, Majken Kirkegaard Rasmussen, Kim Halskov, and Marianne Graves Petersen. 2020. KirigamiTable: Designing for Proxemic Transitions with a Shape-Changing Tabletop. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3313831.3376834>
- [10] Edward T. Hall. 1966. *The Hidden Dimension*. Doubleday, New York, NY, USA.
- [11] Anuruddha Hettiarachchi and Daniel Wigdor. 2016. Annexing Reality: Enabling Opportunistic Use of Everyday Objects as Tangible Proxies in Augmented Reality. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1957–1967. <https://doi.org/10.1145/2858036.2858134>
- [12] Leila Homaieian, Nippun Goyal, James R. Wallace, and Stacey D. Scott. 2018. Group vs Individual: Impact of TOUCH and TILT Cross-Device Interactions on Mixed-Focus Collaboration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173647>
- [13] Brent Edward Insko. 2001. *Passive haptics significantly enhances virtual environments*. The University of North Carolina at Chapel Hill, North Carolina, USA.
- [14] Dongsik Jo, Ki-Hong Kim, and Gerard Jounghyun Kim. 2015. SpaceTime: adaptive control of the teleported avatar for improved AR tele-conference experience. *Computer Animation and Virtual Worlds* 26, 3-4 (2015), 259–269.
- [15] Janet G. Johnson, Danilo Gasques, Tommy Sharkey, Evan Schmitz, and Nadir Weibel. 2021. Do You Really Need to Know Where “That” Is? Enhancing Support for Referencing in Collaborative Mixed Reality Environments. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 514, 14 pages. <https://doi.org/10.1145/3411764.3445246>
- [16] Brennan Jones, Yaying Zhang, Priscilla N. Y. Wong, and Sean Rintel. 2020. VROOM: Virtual Robot Overlay for Online Meetings. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3334480.3382820>
- [17] Adam Kendon. 2010. Spacing and Orientation in Co-present Interaction. In *Development of Multimodal Interfaces: Active Listening and Synchrony*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1–15.
- [18] Jesper Kjeldskov, Jacob H. Smedegård, Thomas S. Nielsen, Mikael B. Skov, and Jeni Paay. 2014. EyeGaze: Enabling Eye Contact over Video. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces* (Como, Italy) (AVI '14). Association for Computing Machinery, New York, NY, USA, 105–112. <https://doi.org/10.1145/2598153.2598165>
- [19] Peter Gall Krogh, Marianne Graves Petersen, Kenton O'Hara, and Jens Emil Grønbaek. 2017. Sensitizing Concepts for Socio-Spatial Literacy in HCI. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 6449–6460. <https://doi.org/10.1145/3025453.3025756>
- [20] Hideaki Kuzuoka. 1992. Spatial Workspace Collaboration: A SharedView Video Support System for Remote Collaboration Capability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Monterey, California, USA) (CHI '92). Association for Computing Machinery, New York, NY, USA, 533–540. <https://doi.org/10.1145/142750.142980>
- [21] Nicolas H. Lehment, Daniel Merget, and Gerhard Rigoll. 2014. Creating automatically aligned consensus realities for AR videoconferencing. In *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 201–206. <https://doi.org/10.1109/ISMAR.2014.6948428>
- [22] Paul Luff and Christian Heath. 1998. Mobility in Collaboration. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (CSCW '98). Association for Computing Machinery, New York, NY, USA, 305–314. <https://doi.org/10.1145/289444.289505>
- [23] Nicolai Marquardt, Nathalie Henry Riche, Christian Holz, Hugo Romat, Michel Pahud, Frederik Brudy, David Ledo, Chunjong Park, Molly Jane Nicholas, Teddy Seyed, Eyal Ofek, Bongshin Lee, William A.S. Buxton, and Ken Hinckley. 2021. AirConstellations: In-Air Device Formations for Cross-Device Interaction via Multiple Spatially-Aware Armatures. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '21). Association for Computing Machinery, New York, NY, USA, 1252–1268. <https://doi.org/10.1145/3472749.3474820>
- [24] Paul Marshall, Yvonne Rogers, and Nadia Pantidi. 2011. Using F-formations to analyse spatial patterns of interaction in physical environments. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. ACM, Association for Computing Machinery, New York, NY, USA, 445–454.
- [25] Jens Müller, Roman Rädle, and Harald Reiterer. 2017. Remote Collaboration With Mixed Reality Displays: How Shared Virtual Landmarks Facilitate Spatial Referencing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 6481–6486. <https://doi.org/10.1145/3025453.3025717>
- [26] Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yuri Degtyarev, David Kim, Philip L. Davidson, Sameh Khamis, Ming-song Dou, Vladimir Tankovich, Charles Loop, Qin Cai, Philip A. Chou, Sarah

- Mennicken, Julien Valentin, Vivek Pradeep, Shenlong Wang, Sing Bing Kang, Pushmeet Kohli, Yuliya Lutchyn, Cem Keskin, and Shahram Izadi. 2016. Holoportation: Virtual 3D Teleportation in Real-Time. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 741–754. <https://doi.org/10.1145/2984511.2984517>
- [27] Kenton O'hara, Jesper Kjeldskov, and Jeni Paay. 2011. Blended Interaction Spaces for Distributed Team Collaboration. *ACM Trans. Comput.-Hum. Interact.* 18, 1, Article 3 (May 2011), 28 pages. <https://doi.org/10.1145/1959022.1959025>
- [28] Mark Perry, Kenton O'Hara, Abigail Sellen, Barry Brown, and Richard Harper. 2001. Dealing with mobility: understanding access anytime, anywhere. *ACM Transactions on Computer-Human Interaction (TOCHI)* 8, 4 (2001), 323–347.
- [29] Mark J Perry and Kenton O'Hara. 2003. Display-Based Activity in the Workplace.. In *INTERACT*. Citeseer, Citeseer, Pennsylvania, USA.
- [30] Thammathip Piumsomboon, Gun A. Lee, Jonathon D. Hart, Barrett Ens, Robert W. Lindeman, Bruce H. Thomas, and Mark Billinghurst. 2018. Mini-Me: An Adaptive Avatar for Mixed Reality Remote Collaboration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173620>
- [31] Holger Regenbrecht, Michael Haller, Jörg Hauber, and Mark Billinghurst. 2006. Carpeno: interfacing remote collaborative virtual environments with table-top interaction. *Virtual Reality* 10 (2006), 95–107.
- [32] Yvonne Rogers and Siân Lindley. 2004. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers* 16, 6 (2004), 1133–1152.
- [33] Jonas Schjerlund, Kasper Hornbæk, and Joanna Bergström. 2022. OVRlap: Perceiving Multiple Locations Simultaneously to Improve Interaction in VR. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 355, 13 pages. <https://doi.org/10.1145/3491102.3501873>
- [34] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA) (CSCW '04). Association for Computing Machinery, New York, NY, USA, 294–303. <https://doi.org/10.1145/1031607.1031655>
- [35] Stacey D. Scott, Karen D. Grant, and Regan L. Mandryk. 2003. System Guidelines for Co-located, Collaborative Work on a Tabletop Display. In *ECSCW 2003*, Kari Kuutti, Eija Helena Karsten, Geraldine Fitzpatrick, Paul Dourish, and Kjeld Schmidt (Eds.). Springer Netherlands, Dordrecht, 159–178.
- [36] Keisuke Shiro, Atsushi Okada, Takashi Miyaki, and Jun Rekimoto. 2018. OmniGaze: A Display-Covered Omnidirectional Camera for Conveying Remote User's Presence. In *Proceedings of the 6th International Conference on Human-Agent Interaction* (Southampton, United Kingdom) (HAI '18). Association for Computing Machinery, New York, NY, USA, 176–183. <https://doi.org/10.1145/3284432.3284439>
- [37] Adalberto L. Simeone, Eduardo Velloso, and Hans Gellersen. 2015. Substitutional Reality: Using the Physical Environment to Design Virtual Reality Experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 3307–3316. <https://doi.org/10.1145/2702123.2702389>
- [38] Mauricio Sousa, Rafael Kufner dos Anjos, Daniel Mendes, Mark Billinghurst, and Joaquim Jorge. 2019. Warping Deixis: Distorting Gestures to Enhance Collaboration. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300838>
- [39] Misha Sra, Aske Mottelson, and Pattie Maes. 2018. Your Place and Mine: Designing a Shared VR Experience for Remotely Located Users. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) (DIS '18). Association for Computing Machinery, New York, NY, USA, 85–97. <https://doi.org/10.1145/3196709.3196788>
- [40] Richard Stoakley, Matthew J. Conway, and Randy Pausch. 1995. Virtual Reality on a WIM: Interactive Worlds in Miniature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 265–272. <https://doi.org/10.1145/223904.223938>
- [41] Ryo Suzuki, Hooman Hedayati, Clement Zheng, James L. Bohn, Daniel Szafir, Ellen Yi-Luen Do, Mark D. Gross, and Daniel Leithinger. 2020. RoomShift: Room-Scale Dynamic Haptics for VR with Furniture-Moving Swarm Robots. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3313831.3376523>
- [42] Kazuki Takashima, Takafumi Oyama, Yusuke Asari, Ehud Sharlin, Saul Greenberg, and Yoshifumi Kitamura. 2016. Study and Design of a Shape-Shifting Wall Display. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) (DIS '16). Association for Computing Machinery, New York, NY, USA, 796–806. <https://doi.org/10.1145/2901790.2901892>
- [43] Xuanyu Wang, Hui Ye, Christian Sandor, Weizhan Zhang, and Hongbo Fu. 2022. Predict-and-Drive: Avatar Motion Adaption in Room-Scale Augmented Reality Telepresence With Heterogeneous Spaces. *IEEE Transactions on Visualization and Computer Graphics* (2022), 1–10. <https://doi.org/10.1109/TVCG.2022.3203109>
- [44] Leonard Yoon, Dongseok Yang, Choongho Chung, and Sung-Hee Lee. 2021. A Full Body Avatar-Based Telepresence System for Dissimilar Spaces. *arXiv preprint arXiv:2103.04380* (2021).
- [45] Leonard Yoon, Dongseok Yang, Jaehyun Kim, ChoongHo Chung, and Sung-Hee Lee. 2022. Placement Retargeting of Virtual Avatars to Dissimilar Indoor Environments. *IEEE Transactions on Visualization and Computer Graphics* 28, 3 (2022), 1619–1633. <https://doi.org/10.1109/TVCG.2020.3018458>