Anchored Multiperspective Visualization for Efficient VR Navigation

Meng-Lin Wu[1]✉ and Voicu Popescu[1]

Purdue University, West Lafayette IN 47907, USA
wu223@purdue.edu

Abstract. This paper presents a novel multiperspective visualization (MPV) approach designed to improve navigation efficiency in Virtual Reality applications. The MPV is continuous and non-redundant, it shows the near part of the scene with a conventional, first-person visualization to anchor the user, and it is controlled with user head translations and rotations reminiscent of natural motion. Three types of anchored MPV are introduced, one that provides a lateral disocclusion effect, allowing the user to see around occluders and through side portals, one that provides a vertical disocclusion effect, allowing the user to see over and on top of occluders, and one that provides teleportation, allowing the user to relocate. The VR navigation efficiency benefits of the anchored MPV have been analyzed in a user study. Significant improvements were achieved in the metrics of number of teleportations and total distance traveled. In these metrics, large or greater Cohen's d effect sizes were observed at p-values below 0.05 in a first VR scene, while medium effect sizes at p-values of 0.1 or better were observed in a second VR scene.

Keywords: Virtual reality · Visualization techniques · Rendering

1 Introduction

In Virtual Reality (VR) applications, a head-mounted display (HMD) tracked with six degrees of freedom supports using real walking for natural navigation, where there is an identity mapping between the user's physical and virtual motion. The user selects the desired view intuitively, by walking to translate the viewpoint, and by rotating their head to change view direction. However, real walking navigation presents several challenges. One challenge is the fact that the real world space hosting the VR application is typically smaller and of a different shape compared to the virtual space, which can prevent the user from reaching some desired viewpoints. For example, a desired viewpoint might coincide with real-world furniture, it might be beyond the walls of the real world room, or it might be high up, on a higher level of a multistory virtual world that is hard to reach.

Another challenge is that in complex virtual world scenes occlusions limit how much the user can see from any given viewpoint. Comprehensive exploration requires translating the viewpoint to circumvent occluders and to gain line of

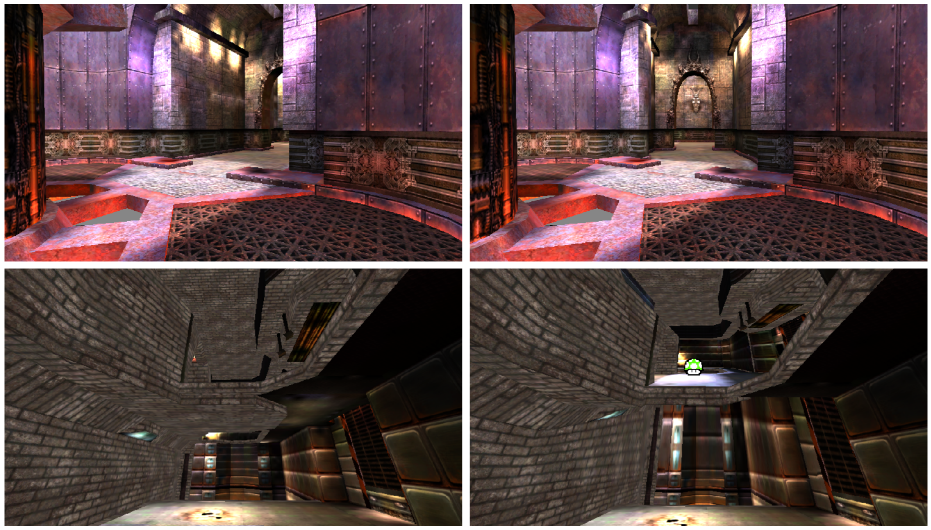

Fig. 1: Top: lateral disocclusion effect. The side corridor is occluded in a conventional visualization (left), and visible in our anchored multiperspective visualization (MPV) (right). The disocclusion effect was deployed by the user with a small left translation of their head. The MPV shows the near part of the scene conventionally, anchoring the user. Bottom: vertical disocclusion effect. A conventional visualization does not show on top of the ledge (left), whereas our anchored MPV does (right). The disocclusion effect was deployed by the user with a small upward translation of their head achieved by getting up on their tiptoes. The MPV shows the ledge and the walls in front of the ledge conventionally, anchoring the user.

sight to all potential regions of interest (ROIs). When a potential ROI turns out to be of no interest, the user has to retrace their path and to explore the next one. Such sequential scene exploration is inefficient. Furthermore, when scene understanding depends on seeing several ROIs simultaneously, or on visualizing dynamic, possibly evading targets, sequential scene exploration is ineffective.

Another reason why real walking might not always be desirable is based on ergonomics considerations. For some applications, the user might prefer not to expend the energy needed to navigate the VR world by always walking and rotating their head in the real world. In other words, for applications where the experience of actual physical locomotion is not essential, users might prefer navigation interface constructs that allow them to see more with less physical effort in a shorter amount of time.

Many approaches have been investigated for overcoming these challenges of using real walking for navigation in VR. One promising approach is based on multiperspective visualization (MPV), which relies on images that integrate samples captured from multiple viewpoints. Consider a virtual scene with two corridors intersecting at a right angle. Using a conventional visualization, a user has to translate the viewpoint up to the intersection to examine the side corridors in search of an ROI. If, on the other hand, an MPV shows not only the main corridor but also the side corridors, the user can examine the side corridors from their current location, which avoids the unnecessary navigation to the intersection when the side corridors turn out to be empty. Similarly, MPV can let the user see distant parts of the scene, without having to move beyond the walls of the real world space hosting the VR application. An MPV can also let the user examine two potential ROIs simultaneously, in parallel, even when no conventional visualization can show both ROIs at the same time.

Harvesting these potential advantages of MPV in the context of VR navigation requires solving two problems: (1) to design an MPV that is effective, i.e. that has the high information payload needed for navigation efficiency, but that remains easy to interpret by the user, and that does not induce user disorientation or motion sickness; (2) to devise navigation interface elements that allow the user to invoke their MPV superpower intuitively, in order to benefit from the additional perspective quickly and to the fullest extent.

In this paper we present anchored multiperspective visualization, a novel multiperspective visualization method designed to improve VR navigation efficiency. Our method was designed based on the following principles: (1) the MPV image should be continuous and non-redundant; (2) the MPV should show the near part of the scene with a conventional first-person visualization controlled through natural motion, anchoring the user; and (3) the MPV effect should be controlled with user motions reminiscent of natural motion, by tethering the secondary perspective selection to the user's head rotations and translations.

We have designed three types of anchored MPV. The first type allows the user to achieve a lateral disocclusion effect (Fig. 1, top) The user cannot see down the right side corridor with a conventional visualization (left). The MPV (right) integrates a secondary perspective into the main user's perspective, allowing the user to see down the right side corridor. The secondary perspective is controlled by the user translating their head to the left as if to look around a corner. The small head translation is amplified and applied to a secondary viewpoint that swings into place to reveal the side corridor. The user view change is used directly, without amplification, to render the nearby geometry, which remains in agreement with the user's proprioception to anchor the user.

The second type of anchored MPV allows the user to achieve a vertical disocclusion effect (Fig. 1, bottom). The user cannot see on top of the ledge in a conventional visualization (left). The MPV (right) integrates an additional perspective, with a high up viewpoint, to reveal the object on the ledge. The secondary perspective is controlled by the user by getting up on their tiptoes as if to examine a tall shelf above eye level. The small vertical user viewpoint translation is amplified and applied to a secondary viewpoint that translates up the necessary amount to see on top of the ledge.

The third type of anchored MPV allows the user to teleport from one location to another. MPV disoccludes parts of the scene not visible from the main user viewpoint, but it does not and should not produce a visualization that shows the entire scene. Consequently, the need to quickly move directly to a distant location of the scene remains even in MPV navigation. We have designed an anchored MPV teleportation method that proceeds in two stages, evocative of how a caterpillar moves (Fig. 2). First, the secondary viewpoint moves forward, getting closer to the far part of the scene, translating from the origin to the destination, while the primary viewpoint doesn't move, remaining at the origin. Second, the primary viewpoint moves forward to the secondary viewpoint, while the secondary viewpoint doesn't move.

We have conducted a user study to detect and quantify any VR navigation efficiency benefits brought by our anchored MPV method. 16 participants were divided evenly in a control group, who used conventional visualization, and an experiment group, who used our anchored MPV. Each participant performed a searching task and a matching task in each of two virtual environments: a single-story area of connected rooms, and a larger room with walkways suspended from the periphery walls, high above the room floor. For the first environment the experiment group participants had available our lateral disocclusion anchored MPV, and for the second environment they had available our vertical disocclusion anchored MPV. In all cases, all participants had the ability to teleport to any scene location to which they had line of sight. The experiment group used our MPV teleportation. The experiment group performed significantly better than the control group in the first virtual environment, achieving improvements in the metrics of distance traveled and number of teleportations. Cohen's d effect size of large and greater was observed with p-values below 0.05. In the more complex second virtual environment, the experiment group achieved improvements of medium Cohen's d effect size at p-values of 0.1 and less. Experiment group participants also reported improvements in spatial awareness and perceived navigation efficiency.

In summary, our paper makes the following contributions: (1) a set of principles for designing VR navigation methods based on multiperspective visualization, (2) three anchored multiperspective visualization based on our design principles, one for lateral disocclusion, one for vertical disocclusion, and one for teleportation, and (3) a user study confirming the potential of our anchored MPV to improve VR navigation efficiency.

2 Prior Work

In VR, a preferred scene navigation modality is actual user locomotion in the physical space, which is translated to matching view changes in the virtual world. However, the physical space typically differs considerably from the virtual space. Due to this mismatch, some virtual viewpoints become inaccessible. In Section 2.1, we discuss this challenge and prior work aimed at alleviating it. Another challenge arises from the reduction in visualization efficiency due to occlusions of ROIs by scene geometry, forcing the user to search for an unobstructed line of sight through extensive viewpoint navigation. In Section 2.2 we review prior

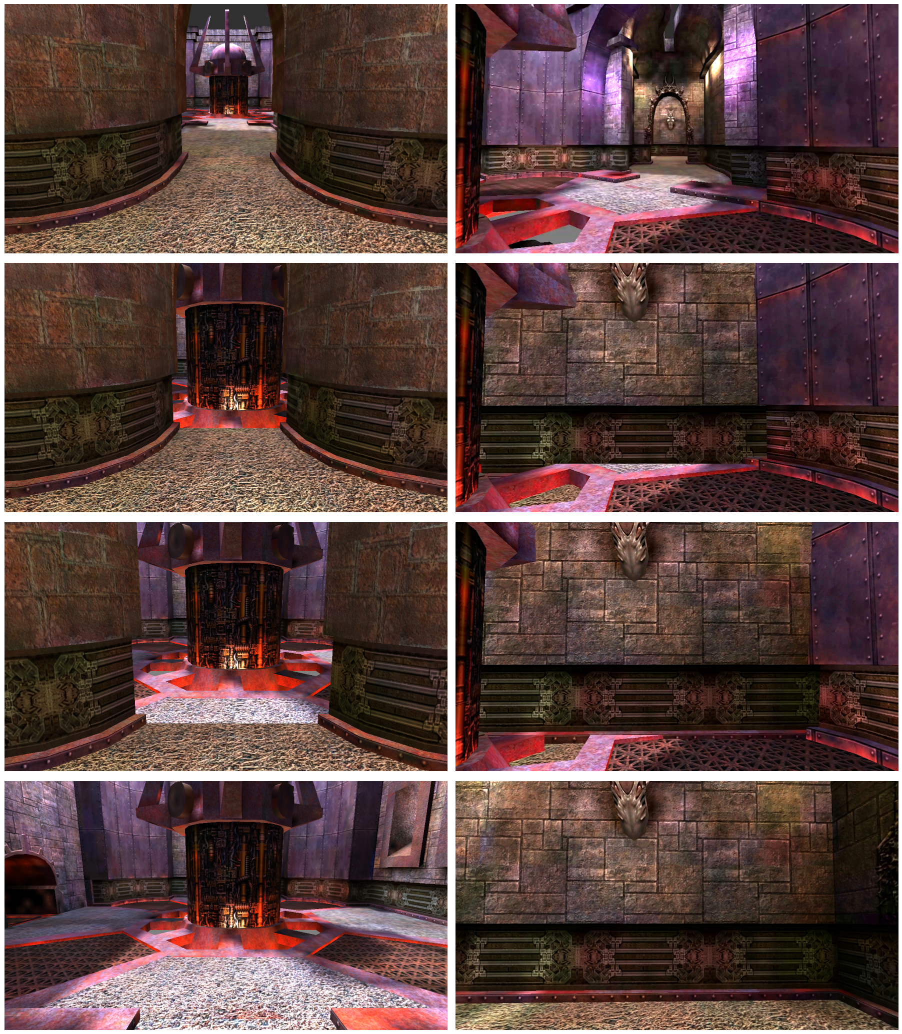

Fig. 2: Two-stage MPV teleportation concept. In the first stage (top two images), the primary perspective stays locked on the origin, anchoring the user, while the secondary perspective translates to the destination. In the second half (bottom two images), the secondary perspective stays fixed, anchoring the user, while the primary perspective moves to assume the secondary perspective.

work for improving VR navigation efficiency using the multiperspective occlusion management approach.

2.1 VR Navigation Challenges

The most intuitive VR navigation is an identity mapping between physical and virtual motion [25, 20]. One common problem is that the physical space is considerably different than the virtual space. Usually the physical space is more restricted than the virtual world.

To fully explore the virtual world, the real and virtual locomotion must purposefully diverge to allow sufficient virtual motion while limiting physical motion. One approach is teleportation, which allows the user to designate a destination in the virtual world, and then to instantly relocate to that destination [1]. The visualization is discontinuous as the user changes location instantaneously, without any indication of the position of the destination relative to the origin. Therefore, the user needs some time to reorient themselves after arriving at the destination. A technique to reduce this discontinuity is to translate the user from the origin to the destination along a straight path [9]. However, a slow translation might cause nausea, as the user's viewpoint changes without any perceived acceleration, while a fast translation does not resolve the visualization discontinuity issue. In practice, a visual "blink", i.e. fade-out followed by fade-in, is applied as the viewpoint translates, to minimize nausea while providing some visual connection between the origin and the destination [4].

Another approach is artificial, or free, locomotion, where the user relies on input devices such as joysticks or keyboards to navigate the view beyond the tracked head pose. The divergence between virtual and physical motion is thus directly controlled by user input. This method preserves visualization continuity, so spatial awareness is not compromised. However, due to the detachment of the user's virtual movement from their physical movement, the artificial locomotion method induces more motion sickness compared with teleportation-based methods [5]. Specifically, the visual and physical senses of acceleration are out of sync. One technique for alleviating nausea is to limit the user to discrete artificial locomotion steps, which are enacted abruptly to break the sensation of artificial acceleration. However, larger steps impact spatial awareness, while smaller steps incur frequent visual discontinuity [4].

Other approaches hide the mismatch between the physical and the virtual worlds by deviating from the tracking data, for example by making the user cover long straight lines in the virtual world by walking in circles in the physical world [19, 18], by resetting user pose [27], and by modifying input gain [30, 17]. Another approach is to distort the virtual world to pack it tightly in the limited confines of the physical world [24]. An approach that blends physical and artificial locomotion is the treadmill approach, or the smart platform approach, where the user actually walks, but without covering large distances in the physical world [23]. The shortcomings of the approach are confusing motion divergences, tethering the user, and reliance on expensive and bulky hardware.

2.2 Multiperspective Visualization in VR

MPV is a class of visualization techniques that integrate multiple perspectives into the main user perspective. MPV originated in the visual arts, e.g. Picasso's

Cubism, and is applied to achieving comprehensive visualization, such as a ski trail map showing simultaneously trails not all visible from a single viewpoint.

Earlier research work focused on relaxing the single center of projection constraint, but the sampling rays remain linear [6, 16, 31]. More recent work introduced piecewise linear or even curved sampling rays that provide the flexibility needed to go around occluders to reach distant ROIs [11, 3, 15]. While relaxation of the constraints opened up more degrees of freedom in the camera model used to render the visualization, the camera model generalization also created the need for automatic and interactive constructors that provide the application with the desired disocclusion effect [28].

In VR occlusion management, MPV is also found to be an effective technique [29], while conventional desktop occlusion management techniques such as transparency and explosion visualizations face various challenges. Transparency techniques introduce visual clutter that scale with scene complexity [7], whereas the MPV approach does not introduce additional geometry and does not violate pictorial depth cues. Explosion techniques disturb scene geometry [10], which impact the user's spatial awareness in VR, whereas the MPV approach does not disturb surface connectivity.

Portal-based visualization is a technique closely related to MPV. It composites additional views of the scene within the main view in a picture-in-picture fashion [8]. In VR applications, it supports teleportation navigation where the user teleports to destinations revealed through the portal [14]. However, the teleportation destination, as viewed through the portal, is beyond the vista space [13]. The user is therefore unable to trace the path of teleportation.

Our MPV method increases scene exploration efficiency by giving the user a preview of ROIs that are occluded in a conventional visualization. Compared with portal-based visualization, our MPV incorporates disoccluded ROIs into the vista space, avoiding disorientation due to untraceable teleportation [26]. Compared with the prior work in MPV navigation [29], our MPV supports both lateral and vertical disocclusion. It provides more versatile and at the same time more intuitive ways of controlling the additional perspectives, and it allows the user to assume seamlessly any additional perspective revealed by MPV using teleportation. The user's spatial awareness is further increased by visually anchoring the user in the scene as the additional perspectives are deployed and retracted, and during teleportation.

3 Anchored Multiperspective Visualization

In this section, we first discuss in more detail our three principles for the design of effective multiperspective visualization for VR navigation, and then we present our three methods for anchored MPV.

3.1 Design Principles for Effective Multiperspective Visualization in VR

(1) *MPV image continuity and non-redundancy*

This first principle encapsulates general concerns for achieving effective MPV, irrespective of the VR context. An MPV has to be continuous, i.e. points that are close in 3D should project to nearby image locations. This concern disqualifies MPVs obtained through a discontinuous collage of individual perspectives, or through parallel visualization with multiple disconnected rectangular images. An MPV also has to be non-redundant, i.e. it should not show a part of the scene multiple times. Continuity and non-redundancy are necessary conditions for obtaining an MPV that can be parsed by the user without the disadvantage of a significant cognitive load, which is particularly important in the VR navigation context.

(2) *Primary perspective MPV anchoring*

This second principle ensures that, as the conventional visualization morphs into an MPV, and as the MPV parameters are changed interactively, there is always a significant part of the image that is unaffected by the MPV effect, and that the unaffected part of the image corresponds to the space surrounding the user. The user's visual system relies on this primary visualization of nearby geometry, in sync with their own primary perspective, to remain in agreement with the motion perceived by the user, and to dissociate from the distant parts of the scene that move incongruently with the perceived motion, both of which contribute to preventing disorientation and motion sickness.

(3) *Natural secondary perspective navigation*

An MPV has a significantly higher number of degrees of freedom than a conventional visualization, with each additional perspective introducing six more extrinsic parameters. This third principle prohibits complex navigation interfaces that ask the user to manipulate a high number of degrees of freedom individually, and mandates allowing the user to control the secondary perspectives with natural motions similar to the ones used to control the main user perspective in conventional VR.

3.2 Anchored MPV for Lateral Disocclusion

A frequently needed disocclusion effect is to see around an occluder, e.g. to see around a tree, or to see through a side opening, e.g. through a window in a house facade. Such a lateral disocclusion effect can be provided by integrating a secondary perspective from a viewpoint that has line of sight around the tree or through the window. We provide a lateral disocclusion effect as follows.

Given a virtual scene, we define a set of vertical rectangles in the scene to serve as portals to guide the lateral disocclusion effect. The portals are defined where the user is likely to benefit from disocclusion, e.g. at the doorways that connect various sections of the scene, or between an occluder and nearby walls, such as a column in a middle of a room. The portals are not rendered, but when a user exploring the scene with a VR HMD sees the geometry spanned by a portal, the geometry changes color to indicate the availability of a disocclusion effect. The user activates the disocclusion effect with a controller button. The activation itself does not change the visualization to an MPV. After activation,

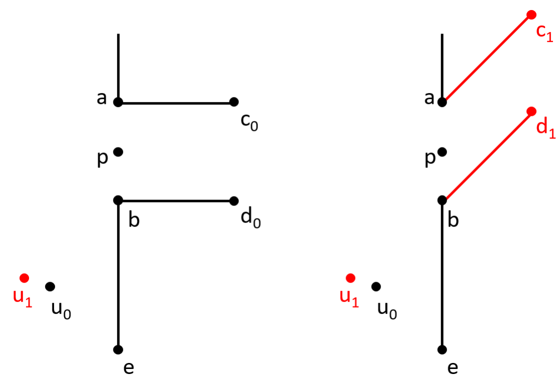

Fig. 3: Lateral disocclusion through anchored MPV. A conventional visualization from viewpoint u_0 does not show the side corridor (left). A small left translation of the head from u_0 to u_1 deploys a disocclusion effect that shows inside the side corridor (right).

lateral translations of the user's head as recorded by the HMD will deploy a secondary perspective that sees through the portal.

In Fig. 3 the user's initial viewpoint is u_0, looking at portal ab. In the conventional visualization (left), the user cannot see deep inside the portal from u_0. Once the user activates the disocclusion effect of the portal, a subsequent left translation of the user viewpoint rotates the geometry behind the portal plane about the pivot point p, which is the center of the portal rectangle. The rotation gives the user line of sight perpendicularly through the portal, e.g. swinging the side corridor wall vertices c_0, d_0 to c_1, d_1, respectively (right). The rotation angle is proportional to the user's lateral head translation, and the gain is tuned such that a small amount of translation u_1u_0 (e.g. 20cm) is sufficient to see down the portal. The small translation wouldn't have been sufficient to see down the portal with a conventional visualization, i.e. the user would see only marginally more inside the portal from u_1 as compared to u_0 (left). The small translation is amplified by our lateral MPV disocclusion effect to introduce the necessary second perspective on the geometry beyond the portal plane.

The rotation angle is capped to the value needed to the see down the portal. Any geometry vertices or fragments that cross the portal plane when rotated are discarded (i.e. they are not drawn), which does not create artifacts as this geometry wouldn't have been visible anyway due to the side corridor walls. Vertex projection is continuous and non-redundant, which enforces the first design principle. Nearby geometry, i.e. the part of the eb wall seen by the user, is drawn conventionally, from the primary perspective, which anchors the user, enforcing the second principle. The additional perspective is deployed by the user translating their head to the left, and by slightly panning the view to the right in order to keep the portal in the center of the image, which is the natural motion the user would make if they were close to point b and wanted to look inside the portal, so the interface is in agreement with the third principle.

3.3 Anchored MPV for Vertical Disocclusion

In addition to lateral disocclusion, an explorer of a 3D virtual scene might also want to be able to see on top of horizontal surfaces that are suspended above the user's eye sight. Given a virtual scene, we define a set of ledge edges. Like in the case of portals, when a VR explorer has a predefined ledge edge into view, a highlight alerts them to the availability of a vertical disocclusion effect, and the user can activate the effect with a controller button.

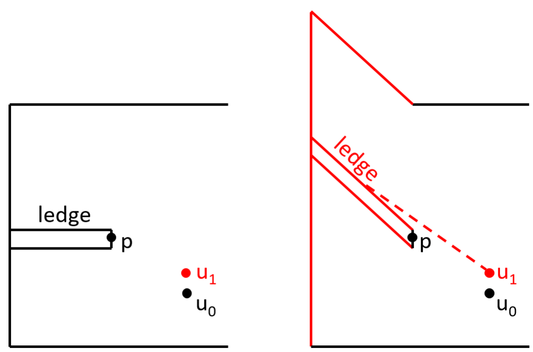

Fig. 4: Vertical disocclusion through anchored MPV. A conventional visualization does not show on top of the ledge (left). The user tiptoes to generate a small upward head translation from u_0 to u_1 that deploys a disocclusion effect that shows on top of the ledge (right).

Fig. 4 shows the conventional visualization of a scene with a ledge (left) and the same scene with our vertical disocclusion MPV effect (right). The initial user viewpoint u_0 is too low to see on the ledge. A small vertical translation of the user viewpoint to u_1 brings in an additional perspective that sees on the ledge. The disocclusion effect is implemented as a rotation of the geometry beyond and above the ledge, such that the new viewpoint u_1 is above the ledge plane, disoccluding the ledge. The viewpoint u_1 wouldn't have been high enough to disocclude the ledge in a conventional visualization (left). The small translation u_0u_1 is amplified to achieve the vertical disocclusion effect.

Like in the case of lateral disocclusion, the vertex projection is continuous and non-redundant, and the user's view of the floor and of the ledge edge doesn't change, anchoring the user. Finally, the vertical up and down translation is controlled by the user's tracked HMD, who gets up on their tiptoes up to see atop the ledge, and back down to revert to a conventional visualization.

3.4 Anchored MPV for Teleportation

We allow the user to teleport between an origin and a destination viewpoint, as needed to change floors, rooms, and, in general, to overcome the constraints of the real world and of the tracking system. Teleportation is a rapid transition between the two viewpoints, which can induce user motion sickness because the user moves, i.e. ”flies”, without actually engaging in locomotion. Prior work suggests that the safest teleportation is a very abrupt one, but that is also the teleportation that disorients the user the most.

We have developed a teleportation that aims to alleviate these disadvantages. The user selects the destination with the cursor, which is always placed at the intersection between the view direction and the scene geometry. Therefore, the destination is selected with the HMD by changing view direction. The destination viewpoint is the point on the vertical through the cursor that is at the user’s height above the ground. If this initial destination viewpoint is too close to a wall, the destination viewpoint is pulled back away from the wall to provide a meaningful view once teleportation is complete. The user triggers teleportation with a controller button.

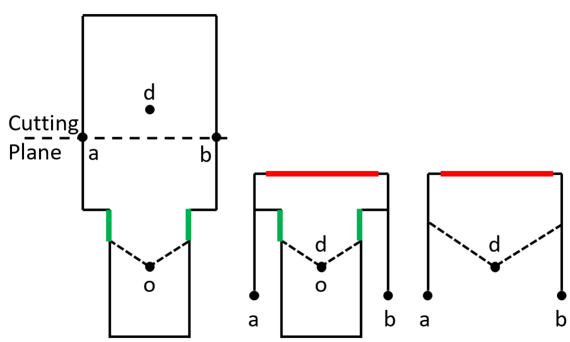

Fig. 5: Anchored MPV teleportation. The viewpoint moves from o to d in two phases: first the part of the scene beyond the cutting plane is brought down, anchoring the user with the near part of the scene (green, from left to middle), and then the cutting plane is brought down, anchoring the user with the far part of the scene (red, from middle to right)

Our anchored teleportation method is illustrated in Fig. 5. The origin and destination viewpoints are o and d. During the first phase (i.e. transition from left to middle in Fig. 5), the far perspective is brought closer, by translating the scene geometry beyond the cutting plane with vector od. The cutting plane is positioned between the two viewpoints, splitting the distance between them at a fixed ratio. This first phase brings the far part of the scene closer, to be as close as it will be when the viewpoint is at d, but without pushing away the near part of the scene. The visualization of the near part of the scene (green) remains unchanged, which anchors the first phase of the teleportation. During the second phase, the cut plane is translated along the vector *do*, which makes the nearby geometry disappear from view. The visualization of the far part of the scene (red) remains unchanged in this second phase, which maintains uninterrupted anchoring.

Fig. 5 illustrates teleportation along a straight line, but the same procedure is followed if the user chooses to teleport using a lateral or vertical disocclusion MPV. The only difference is that once the second phase is complete, any residual distortion of geometry is gradually eliminated. Geometry is distorted only close to the portal or ledge planes, and typically the user desires to teleport deeply through the portal or beyond the ledge, so the geometry is undistorted off screen. Fig. 2, left, shows frames from a "straight line" teleportation. Fig. 2, right, illustrates teleportation into a side corridor, starting from the lateral disocclusion effect.

4 User Study

In order to evaluate the effect of our multiperspective visualization technique on VR navigation efficiency, we conducted a randomized user study, with the approval of our Institutional Review Board. Each participant performed three tasks in VR, and their actions were logged for subsequent analysis. After task completion, the users responded to a questionnaire that provided their subjective evaluation of task performance.

The participant wore a VR HMD (i.e. Windows Mixed Reality headset [12]). The VR HMD performs six degree of freedom SLAM-based tracking of the pose of the participant's head. In addition, the participant used a motion-tracked hand-held controller to enable interactions with the virtual world through actions such as pointing and clicking buttons. The HMD displays stereoscopic image pairs rendered from viewpoints offset by the interpupillary distance to provide stereoscopic depth perception.

4.1 Participants

A total of 16 participants (14 male) completed our study. The participants were graduate students of ages between 23 and 37. They were randomly assigned to experiment and control groups of 8 participants each. Participants in the control group performed all tasks using only conventional, single-perspective visualization, while the participants in the experiment group performed all tasks using our anchored MPV.

The between-group design was chosen over the within-subject design to avoid any learning effect when repeating tasks from one condition to the next. Another reason is to mitigate the fatigue factor that can affect performance in extended task performance in VR environments–with the between-group design a participant's involvement time is reduced in half. However, the effect of prolonged usage of MPV navigation is an important avenue for future research, which is discussed further in Section 5.

4.2 Tasks and Evaluation

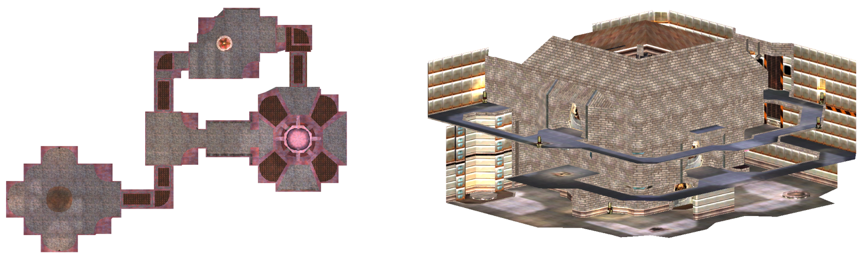

Fig. 6: VR scenes used in the user study. Left: Scene 1 is a set of rooms connected by corridors. Right: Scene 2 is a 3-story building with multi-level walkways.

The participants performed tasks that required extensive virtual locomotion in two VR scenes using environments adapted from the Quake 3 Arena [22]. The first VR scene (Scene 1) is a single-story indoors area consisting of a set of rooms connected by corridors (Fig. 6, left). Two of the rooms contain cylindrical pillars in the center, which partially occlude the interior of the rooms regardless of the participant's current viewpoint. The second VR scene (Scene 2) is a large 3-story building where the center is an inaccessible tower, while the levels are connected by walkways attached to the perimeter walls (Fig. 6, right).

Each participant was required to perform 3 tasks: two search tasks, each in a different VR scene, and one pair matching task. Before beginning each task, the participant completed a short warm-up exercise which is similar to the actual task, but differs in content and is much shorter. This warm-up period ensured that the participant correctly operated the HMD and the hand-held controller, and that the test procedures were clear. Participants received no other training beyond this warm-up period.

The participant's performance was evaluated from recorded data using objective metrics, which were unknown to the participants. During the performance of each task, the participant's tracked physical HMD pose was logged, along with any interaction events such as initiating teleportation, acquiring a target object, or matching a pair of target objects. The logs were processed to extract metrics of interest. Additionally, each participant was asked to respond to a questionnaire after completion of the tasks. The questionnaire recorded subjective evaluation of performance and comfort.

4.3 User Interface for Locomotion and Anchored MPV

In both Scene 1 and Scene 2, the user was free to employ real walking by moving within a $2m$ by $2m$ physical space. To navigate beyond the limits of the

Fig. 7: The user selects the teleportation destination using a hand-held controller. Left: In both VR scenes, the user teleports by pointing a "laser" beam at the destination and clicking using the controller. Right: In the second VR scene, the walkway edges that are eligible as teleportation destinations are highlighted in blue when swiped with the laser beam.

physical space, the user employed teleportation. The user pointed the hand-held controller at the intended destination, and then clicked the controller button to initiate the teleportation. (Fig. 7, left). In Scene 2, the user could additionally access higher floor walkways by designating the walkway edges as a teleportation destination. As the user points the hand-held controller at a walkway edge, the edge is highlighted to indicate that the edge is a valid destination (Fig. 7, right). The user can also select any visible lower floor surface. As the controller button is clicked, the virtual viewpoint is teleported to the destination. The teleportation is not completed instantly, but the virtual viewpoint is translated to the destination at a fast $40m/s$ straight line velocity. During the translation, participants in the control group experienced the conventional "blink" visual effect, whereas participants in the experiment group experienced the anchored MPV teleportation described in Section 3.4.

Our lateral disocclusion anchored MPV (Fig. 1, top) was accessible to experiment group participants performing tasks in Scene 1. The anchored MPV is activated automatically as the user gazes at virtual portals defined by corridor archways or connecting cylindrical pillars to side walls. The available archways and pillars in Scene 1 are predetermined prior to the study. Once the MPV is activated, the user moves laterally in small amounts to rotate the secondary perspective horizontally.

Our vertical disocclusion anchored MPV (Fig. 1, bottom) was accessible to experiment group participants performing tasks in Scene 2. The vertical disocclusion anchored MPV is manually activated. First, the user designates a walkway edge on a different floor by pointing with the hand-held controller and holding down the button. Then, with the button held, the user tiptoes or crouches a small amount to raise or lower the secondary viewpoint. Due to the need to hold down the button for activating the anchored MPV, the user of the MPV is required to double-click the controller button in order to initiate teleportation in Scene 2. Any new area revealed by the anchored MPV is also a valid destination for teleportation.

Search Tasks (Search 1 and Search 2) The search tasks Search 1 and Search 2 were performed in Scene 1 and Scene 2 respectively. In both search tasks, the participant was asked to find target objects in the form of gold coins placed in the VR scene. There were 6 coins that appeared one after another. Their locations were unknown to the participant, but were identical for all participants. A coin disappeared once it was found and collected by the participant by getting within $1m$ of the coin, either through real walking or through teleportation. An audio cue was triggered at the collection of a coin to notify the participant. The task was complete when all coins are collected.

Pair Matching Tasks (Match) The pair matching task (Match) was performed in Scene 1. In this task, the participant was asked to identify target objects in the form of colorful mushrooms, of identical color pattern. There were 8 objects in 4 distinct color patterns, whose placement was identical for all participants. The participant first pointed to one visible target object and selected it by clicking a button. When the participant clicked on the second object of identical pattern, the pair was considered matched and it disappeared from the scene, with an audio cue, The task was complete when no targets remained.

Post-Performance Questionnaire Each participant was asked to respond to a questionnaire after completing the three tasks. The purpose of the questionnaire was to evaluate the participant's perception of their own performance and of of our MPV technique. The participant responded to each of three statements "You always felt present in the virtual world", "Your spatial awareness was maintained while moving around", and "You could reach any intended destination efficiently", by choosing an answer on a 1 to 5 scale, with 1 meaning "Not at all" and 5 points meaning "Very much".

4.4 Results and Discussion

We analyzed each participant's recorded logs to measure performance along several metrics. The most relevant metrics for this study are the metrics for navigation efficiency: 1) the number of times the participant initiated teleportation, 2) the accumulated teleportation distance, and 3) the time required to complete a task. These metrics were not revealed to the participants.

These metrics are evaluated for their effect sizes using Cohen's d [2], where qualifiers "small", "medium", and "large" are applied to cases where $d > 0.2$, $d > 0.5$, and $d > 0.8$, respectively. The qualifiers are extended to include "very small", "very large", and "huge" for $d < 0.01$, $d > 1.2$, and $d > 2.0$, respectively [21]. Statistical significance is evaluated using the two-sample t-test. We report the p-value to identify measurements that were due to chance.

Finally, we report the aggregate results of the questionnaire responses for discussion of subjective metrics which cannot be extracted from the task logs.

Number of Teleportations Each teleportation incurs a discontinuity in the user's mental localization in the virtual world, after which the user must re-orient themselves. Therefore, it is desirable to minimize the number of teleportations. Furthermore, the ill-effects of teleportation can be reduced by improving visualization during teleportation.

We expect our anchored MPV to reduce the number of teleportations for two reasons. First, our MPV disoccludes ROIs and allows the user to plan their path more efficiently. Therefore, the user is able to avoid teleportating to destination only to find that it is not of interest. The second reason is that our anchored MPV extends the set of possible teleportation destinations to those areas newly disoccluded by the visualization. As the user is not limited to teleporting only to destinations with a line of sight, they are able to traverse at once a path that might require multiple teleportations using conventional visualization. Table 1 shows the result for the metric of number of teleportations. In Search 1 and Match tasks, which are both performed in the single-floor Scene 1, our anchored MPV significantly reduced the number of teleportations required to complete the tasks, with the p-value well below 0.05. The effect sizes are very large and large for tasks Search 1 and Match respectively. In the Search 2 task, which is performed in the more complex 3-story Scene 2, the results are positive with a medium effect, although with a higher p-value of 0.06.

Table 1: Average number of teleportations per participant.

task	control	experiment	diff.	p	d	effect size
Search 1	83.4 ± 48.9	21.6 ± 18.5	61.8	< 0.01	1.7	very large
Match	51.8 ± 24.8	27.6 ± 22.7	24.1	0.03	1.0	large
Search 2	98.3 ± 71.5	50.3 ± 31.5	48.0	0.06	0.9	medium

Accumulated Teleportation Distance We analyze the accumulated teleportation distance because the majority of locomotion that is conducted in our VR scenes is through teleportation, therefore it is representative of the total amount of virtual viewpoint travel. We expect users of our anchored MPV to accumulate less traveled distance. This is due to the increased path planning efficiency as discussed in 4.4. Table 2 lists the average per participant distance traveled for each task. Significant improvements of huge and large Cohen d's effect sizes were observed for both Search 1 and Match tasks performed in Scene 1. This is in line with our expectation that the user is able to explore maps effectively with less required virtual locomotion. The result for task Search 2 performed in the

Table 2: Average distance traveled per participant, in meters.

task	control	experiment	diff.	p	d	effect size
Search 1	464 ± 135	222 ± 90	242	< 0.01	2.1	huge
Match	326 ± 103	228 ± 86	99	0.03	1.0	large
Search 2	677 ± 306	527 ± 236	150	0.1	0.5	medium

more complex Scene 2 also shows an improvement of medium effect size. The statistical significance at $p = 0.1$ is not as strong as with tasks performed in Scene 1. However, the positive effect size suggests a more significant result is possible with more user study participants.

Table 3: Task completion time, in seconds.

task	control	experiment	diff.	p	d	effect size
Search 1	113 ± 45	56 ± 29	57	< 0.01	1.5	very large
Match	89 ± 25	78 ± 25	11	0.2	0.4	small
Search 2	157 ± 71	175 ± 88	(18)	0.7	0.2	small

Task Completion Time Table 3 reports the time our participants took to complete the three tasks. The experiment group had a statistically significant advantage for the first task. For the second task, the experiment group was faster, but the advantage was not statistically significant. For the third task, the average completion time for the experiment group was longer than for the control group. From our observation of the participants during the experiment, we explain this based on a longer time the participant needed to engage the MPV interface. As shown in Table 1, the number of jumps is significantly lower for the experiment group even for the third task, indicating that the MPV is effective, except that using it takes longer. A direction of immediate future work is to improve the time effectiveness in which the vertical disocclusion MPV is used, by suggesting to the user the availability of the effect in more salient way during training (the blue highlight is sometimes easy to miss), and then by suggesting the tiptoeing mechanism that actually implements the MPV.

Questionnaire Responses The final metric we used to compare the two participant groups was a compilation of the self-evaluation questionnaire responses (Table 4). The experiment group self-reported higher spatial awareness and higher navigation efficiency, while they reported a lower sense of actual presence in the virtual environment. It comes as no surprise that the experiment group had better spatial awareness than the control group as the MPV essentially provides a preview of the scene, without as much disorienting backtracking as required by the sequential exploration with a conventional visualization. Furthermore, our MPV was designed to anchor the user at all times, so the additional information presented did not come at the cost of confusing the user. Similarly, the improvement of navigation efficiency is a reasonable hypothesis based on the same arguments. We explain the decreased sense of presence by the fact that, like any MPV, our method upgrades the user from an uninterrupted immersive first-person view of the scene to an occasional second or even third-person monitoring of the scene. The user shifts fluidly viewpoint and their association with a specific location in the scene, and the consequent decrease in sense of presence is a reasonable trade-off towards gaining navigation efficiency.

Table 4: Overall subjective evaluation.

	control	experiment	diff.	d	effect size
Spatial awareness	3.3 ± 1.3	4.3 ± 1.0	1.0	0.7	medium
Efficiency	4.1 ± 1.5	4.6 ± 0.5	0.5	0.4	small
Presence	4.0 ± 1.1	3.5 ± 1.1	(0.5)	0.4	small

5 Conclusions and Future Work

We have presented a novel method for multiperspective visualization for Virtual Reality that promises to improve navigation efficiency. Our method visualizes the scene with images that show more than what is visible from a single viewpoint, by integrating additional perspective continuously and non-redundantly. Another goal of our MPV is to always anchor the user by showing part of the scene geometry conventionally, from the user's first-person view. Finally, MPV navigation should remain as intuitive as possible, by allowing the user to control the additional perspectives through the tracked HMD. We describe anchored MPV techniques for lateral disocclusion, for vertical disocclusion, and for teleportation. The MPV benefits have been confirmed in a user study.

Additional user studies should be conducted to explore in depth the subjective effects of MPV navigation. Our post-performance questionnaire provided preliminary insights to the users' perceptions, but the study is not tailored to measure subjective effects such as visual quality, cognitive effort, and user comfort. Especially, tasks which require prolonged usage of MPV navigation should be designed to examine any cumulative effect of simulator sickness, even though no test participant expressed discomfort at any point during the study. At the same time, these longer tasks allow examination of any training effects as participants become familiar with the user interface.

One direction of future work is to find automatically the places in the scene where disocclusion effects are useful. This is in view of the limitation that the virtual portals needed to be manually marked in Scene 1 of our user study. One option is to preprocess the scene, and another option is to decide on the potential for disocclusion on the fly, based on the current frame.

Another direction of future work is to investigate the anchored MPV benefits in the context of dynamic, and even evading targets, which place even more stringent requirements on the quality of the disocclusion effect and on the intuitiveness of the interface for deploying it. These requirements are particularly relevant to gaming applications, where targets could follow complex strategies, or they could be other humans. Furthermore, there is interplay between leveraging visualization to facilitate navigation, and designing visualization for game mechanics. It is worth studying how to optimize for both sets of goals in the design space for VR visualization. As VR interfaces strive to become mainstream and to move beyond entertainment and into day to day use, a scenario that requires special attention is the sit at desk scenario, for which multiperspective visualization might be particularly well suited.

References

1. Bozgeyikli, E., Raij, A., Katkoori, S., Dubey, R.: Point & teleport locomotion technique for virtual reality. In: Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play. pp. 205–216. CHI PLAY '16, ACM, New York, NY, USA (2016)
2. Cohen, J.: Statistical power analysis for the behavioral sciences: Jacob Cohen. Psychology Press (2009)
3. Cui, J., Rosen, P., Popescu, V., Hoffmann, C.: A curved ray camera for handling occlusions through continuous multiperspective visualization. IEEE Transactions on Visualization and Computer Graphics **16**(6), 1235–1242 (Nov 2010)
4. Davis, S., Nesbitt, K., Nalivaiko, E.: A systematic review of cybersickness. In: Proceedings of the 2014 Conference on Interactive Entertainment. pp. 8:1–8:9. IE2014, ACM, New York, NY, USA (2014)
5. Habgood, J., Moore, D., Wilson, D., Alapont, S.: Rapid, continuous movement between nodes as an accessible virtual reality locomotion technique. In: Proceedings of the 25th IEEE Conference on Virtual Reality and 3D User Interfaces. VR '18, Reutlingen, Germany (2018)
6. Hartley, R.I., Gupta, R.: Linear pushbroom cameras, pp. 555–566. Springer Berlin Heidelberg, Berlin, Heidelberg (1994)
7. Kameda, Y., Takemasa, T., Ohta, Y.: Outdoor see-through vision utilizing surveillance cameras. In: Third IEEE and ACM International Symposium on Mixed and Augmented Reality. pp. 151–160 (Nov 2004)
8. Kunert, A., Kulik, A., Beck, S., Froehlich, B.: Photoportals: Shared references in space and time. In: Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing. pp. 1388–1399. CSCW '14, ACM, New York, NY, USA (2014)
9. Laviola, J.J.: A discussion of cybersickness in virtual environments. SIGCHI Bulletin (2000)
10. Li, W., Agrawala, M., Curless, B., Salesin, D.: Automated generation of interactive 3d exploded view diagrams. ACM Trans. Graph. **27**(3), 101:1–101:7 (Aug 2008)
11. Mei, C., Popescu, V., Sacks, E.: The Occlusion Camera. Computer Graphics Forum (2005)
12. Microsoft: Windows mixed reality http://www.microsoft.com/
13. Montello, D.R.: Scale and multiple psychologies of space. In: Frank, A.U., Campari, I. (eds.) Spatial Information Theory A Theoretical Basis for GIS. pp. 312–321. Springer Berlin Heidelberg, Berlin, Heidelberg (1993)
14. Neat Corporation: Budget cuts (2018), http://neatcorporation.com/
15. Popescu, V., Rosen, P., Adamo-Villani, N.: The graph camera. ACM Trans. Graph. **28**(5), 158:1–158:8 (Dec 2009)
16. Rademacher, P., Bishop, G.: Multiple-center-of-projection images. In: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques. pp. 199–206. SIGGRAPH '98, ACM, New York, NY, USA (1998)
17. Ragan, E.D., Scerbo, S., Bacim, F., Bowman, D.A.: Amplified head rotation in virtual reality and the effects on 3d search, training transfer, and spatial orientation. IEEE Transactions on Visualization and Computer Graphics **23**(8), 1880–1895 (Aug 2017)
18. Razzaque, S., Kohn, Z., Whitton, M.C.: Redirected Walking. In: Eurographics 2001 - Short Presentations. Eurographics Association (2001)

19. Razzaque, S., Swapp, D., Slater, M., Whitton, M.C., Steed, A.: Redirected walking in place. In: Proceedings of the Workshop on Virtual Environments 2002. pp. 123–130. EGVE '02, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2002)
20. Ruddle, R.A., Lessels, S.: The benefits of using a walking interface to navigate virtual environments. ACM Trans. Comput.-Hum. Interact. **16**(1), 5:1–5:18 (Apr 2009)
21. Sawilowsky, S.S.: New effect size rules of thumb. Journal of Modern Applied Statistical Methods **8**(2), 597–599 (2009)
22. id Software: Quake 3 arena (1999)
23. Souman, J.L., Giordano, P.R., Schwaiger, M., Frissen, I., Thümmel, T., Ulbrich, H., Luca, A.D., Bülthoff, H.H., Ernst, M.O.: Cyberwalk: Enabling unconstrained omnidirectional walking through virtual environments. ACM Trans. Appl. Percept. **8**(4), 25:1–25:22 (Dec 2008)
24. Sun, Q., Wei, L.Y., Kaufman, A.: Mapping virtual and physical reality. ACM Trans. Graph. **35**(4), 64:1–64:12 (Jul 2016)
25. Usoh, M., Arthur, K., Whitton, M.C., Bastos, R., Steed, A., Slater, M., Brooks, Jr., F.P.: Walking > walking-in-place > flying, in virtual environments. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques. pp. 359–364. SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA (1999)
26. Weissker, T., Kunert, A., Froehlich, B., Kulik, A.: Spatial updating and simulator sickness during steering and jumping in immersive virtual environments. In: Proceedings of the 25th IEEE Conference on Virtual Reality and 3D User Interfaces. VR '18, Reutlingen, Germany (2018)
27. Williams, B., Narasimham, G., Rump, B., McNamara, T.P., Carr, T.H., Rieser, J., Bodenheimer, B.: Exploring large virtual environments with an hmd when physical space is limited. In: Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization. pp. 41–48. APGV '07, ACM, New York, NY, USA (2007)
28. Wu, M.L., Popescu, V.: Multiperspective focus+context visualization. IEEE Transactions on Visualization and Computer Graphics **22**(5), 1555–1567 (May 2016)
29. Wu, M.L., Popescu, V.: Efficient vr and ar navigation through multiperspective occlusion management. IEEE Transactions on Visualization and Computer Graphics **PP**(99), 1–1 (2017)
30. Xie, X., Lin, Q., Wu, H., Narasimham, G., McNamara, T.P., Rieser, J., Bodenheimer, B.: A system for exploring large virtual environments that combines scaled translational gain and interventions. In: Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization. pp. 65–72. APGV '10, ACM, New York, NY, USA (2010)
31. Yu, J., McMillan, L.: General Linear Cameras, pp. 14–27. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)