IPT-EGVE Symposium (2007)
B. Fröhlich, R. Blach, and R. van Liere (Editors)

Visual Consistency in Rotational Manipulation Tasks in Sheared-Perceived Virtual Environments

Michal Koutek, René Molenaar, Gerwin de Haan and Frits H. Post

Delft University of Technology, The Netherlands

Abstract

Sharing one large screen projection display, such as the Virtual Workbench, with multiple users can cause distortions in viewing and interaction, as the users perceive a sheared and moving space. Generally, object rotation in non-orthogonal, sheared coordinate spaces is something that should be avoided as it consequently invalidates homogeneous geometric object transformations. Although this problem seems to be artificial, in shared multiple-user VR this issue gets a real, substantial flavor. It is typically a problem for the "secondary" users viewing and interacting with a VR system. Due to rendering the scene from a different viewpoint, they perceive several distortions in the stereo image, one of which is that the VE appears to be sheared. This is also the case in our affordable approach to a multiple user VR Workbench. In this paper we describe technical aspects of our novel viewpoint compensation method to make the object selection, translation, and rotation consistent with the (secondary) user's view of the scene. We focus on description of object rotation inside sheared VE's. We demonstrate that our techniques ensure a sense of interaction consistency despite the view distortions.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computing Methodologies]: Computer GraphicsVirtual Reality;

1. Introduction

Virtual Reality Workbenches have the potential to be used for direct collaboration by a small group of people, and for interactive VR demonstration. To make a Workbench suitable for these scenarios, two major issues need to be resolved: a suitable stereo display and view consistent interaction for the users. When the typical (single-user) Workbench is employed for this situation without any adjustments, the usability is be rather limited.

One solution would be to provide a robust, multi-user stereo display system. Every user would be headtracked and able to correctly interact in 3D with the system. Obviously, this is difficult and expensive to implement. The hardest part for the Workbench system proves to be multi-user stereo display. We have developed an approach and a set of guidelines that make an "ordinary" single-user Workbench suitable for multi-user viewing and interaction in a rather affordable fashion [dHMKP07]. In our approach we can share the screen in a number of ways to provide a good stereo display for certain application/collaboration scenarios [Mol07]. In this paper we use one of the viewing modes, where multiple users share the whole display with only one stereo delivery system.

By applying certain constraints and following our guidelines, we can minimize the adverse effects when one stereo display is being shared. The most challenging part however is to make the interaction of all the users consistent with their visual perception of the 3D scene. The reason is that the main part of the scene is rendered from one viewpoint, while only user-specific parts of the scene can be displayed correctly for their individual viewpoints. This typically leads to a situation where the "secondary" users manipulate objects in a visually distorted VE. In this paper we limit our explanation to the direct and remote (ray-casting) interaction techniques. For the VR effect to work, it is important to retain the consistency between the interaction (scene) space and the visual space. When the visual space of the VE is distorted by shearing, the interaction becomes difficult and often confusing, according to novice users. For example moving an object straight up in the visual/perceived space gets interpreted as moving in a slanted direction relative to the screen plane. Another example is that without compensation, it will not be clear around which axis is an object actually rotating.

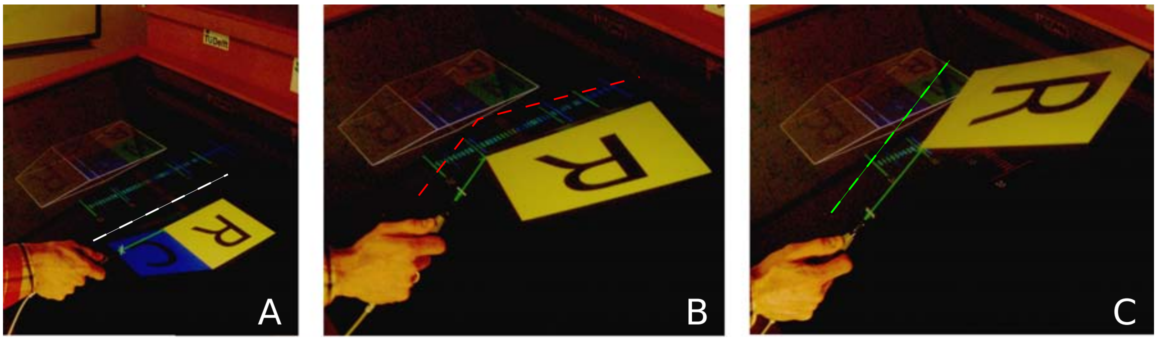

Figure 1: *Motivation toward consistency during object rotations in (visually) sheared virtual environments: (a) selection ray aligned with the object edge, (b) without and (c) with the viewpoint compensation for rotation.*

Without any viewpoint compensation (see Figure 1b), rotations in a sheared-perceived space cause alignment errors and results will be unpredictable for the user. Compensating for the shearing effects while rotating an interaction device will yield an interaction that appears to be consistent from the user's perspective (see Figure 1c).

The contribution of this paper is as follows. We provide a technical explanation of our viewpoint compensation method for the scene rendering and for the interaction handling. We analyze the effects of object rotation in sheared spaces and we focus on making the rotation consistent with the view perception of the interacting user.

The remainder of this paper is as follows: First, we describe related and previous work in section 2. Then, in section 3, we describe the issues of a multi-user stereo display on a VR Workbench system and their consequences on the user interaction. In section 4, we present a technical description of our method. We performed experiments to demonstrate the problems and the handling of the rotations in a sheared VE, of which relevant results are presented in section 5.

2. Related Work

In many cases 3D stereo displays are designed for a single (tracked) user. Other users can perceive images, but experience several distortions [ABM*97] [BLCN02] [SS05]. This strongly reduces the immersive experience and precludes accurate interaction. The main design challenge for multi-user 3D stereo displays has been to provide a separate, completely correct image for each eye of each user [BMC04]. Common solutions use optical filtering (e.g. polarization filters), time multiplexing [ABM*97] [MB97] (e.g. shutter glasses), multiple screens, or hybrid combinations of these [FHH*05]. Separate screens and mirrors and physical masks can also be used to create a shared virtual environment for all users [APT*98] [MB04] [BFSE01] [KKYK01]. These techniques can be used to extend existing, single-user systems, or to construct new multi-user systems. Limitations however, can include reduced image brightness, reduced image dimensions, crosstalk (ghosting) and image flicker. Existing VR systems cannot always be easily extended, additional users come at the cost of adding expensive hardware, extra screens and lesser image quality or dimensions. Simon et al. [SS05], [Sim05], [Sim07] presented a different approach on their panoramic screen VR system. Here the scene was rendered from a static, generic viewpoint. Only some elements in the scene (like interaction tools) were displayed correctly for its corresponding user, eliminating distortion like [WHR99] did for false eye separation. Existing systems can be extended for multiple users by tracking the extra users. Their results were our main motivation to extend this work at the following points. Our workbench differs greatly from the panoramic screen used in Simon's study [Sim07] in terms of distortion. In our other publications we therefore analyzed the problems of multiple-user viewing and interaction on the Workbench [dHMKP07] [Mol07] and the source and effects of the distortions in more detail. We discussed possible solutions and introduced our approach to ensure visual and interaction consistency for all interacting users. We used a shared dynamic-camera model. The camera is placed at an average viewpoint, allowing all users to experience stereo and motion parallax effects. To evaluate our method, we performed both a quantitative and a qualitative user study. We evaluated effects of various amounts of view distortion on task performance and accuracy, and the impact of our view compensation approach. For this, we focused on basic object selection and manipulation tasks (mainly translation). In many cases, interaction with view compensation can be as effective as in the single-user case. During this previous study however, we noticed additional distortions for rotations in distorted projections. And this initiated the research presented in this paper.

3. Sharing a Single Stereo Display

In the single user situation, images are rendered only for the first headtracked user. This user has optimal images, rendered according to his point of view. He/she correctly experiences the depth cues (stereo and motion parallax) that are most important to make the VR work. Because of headtracking only this user, other viewers experience adverse side effects that makes interaction much more difficult for them, if not impossible. First of all the stereo fusion may be lost. When the headtracked user is tilting his head, the stereo parallax becomes correct only for him. Other viewers also experience both point of view distortions and motion distortions. Objects appear sheared and at the wrong locations. Consequently, if such a viewer for example wishes to use a ray casting interaction tool from this viewpoint, he does not see the ray come from the stylus. He has to point at a different place in the scene to shoot the ray at the intended object. Motion distortions also influence interaction and are the result of the absence or presence of parallax effects. To support multiple users on one system we use alternative shared camera and headtracking models. The (stereo) camera remains horizontal and is dynamically placed in between the users, a slider widget allows users to take the camera to their viewpoint. All users must be headtracked. This allows us to calculate a proper position of interaction tools, making them appear correct for the corresponding user despite their visual (mis-)perception of the scene. We present a more thorough analysis and description in [dHMKP07].

3.1. Sheared Coordinate Spaces

This display setup approach suffers from a visual distortion, that appears to be equivalent to a rather simple shearing distortion. Our claim is that this shearing distortion is measurable, correctable, and preventable. We will now describe the cause of the shearing distortion, how we can measure it and how we can correct for the unpleasant effects.

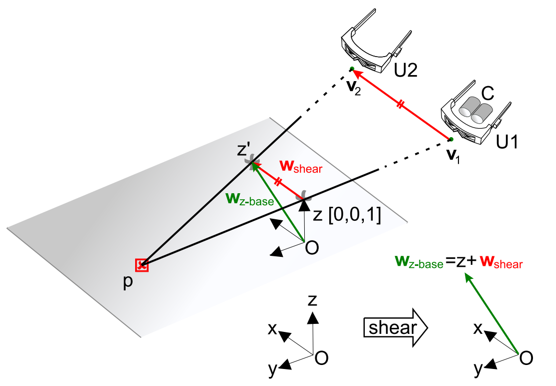

Figure 2: *Construction of the sheared coordinate system: The axis vectors X and Y are not distorted. The Z axis suffers from shear distortion. Point Z is perceived at Z'.*

Figure 2 illustrates this situation. User $U1$ has a correct view and observes the Z axis as being perpendicular to the ground plane, pointing upwards. His perception of the scene space can be described by an orthonormal coordinate system matrix, which is in this case equal to the identity matrix, as all the axis base vectors are orthogonal and have a unit length. The view of the second user $U2$ however, suffers from a shearing distortion. $U2$ experiences the Z axis as pointing into a sheared direction. The X and Y base vectors remain unaffected. The vector $\mathbf{w}_{shear}$ defines the magnitude and direction of the shearing. It can be calculated from any point above the ground plane. Here, we use the base vector $\mathbf{z}[0,0,1]$ for convenience. As shown in Figure 2, its planar projection from viewpoint $\mathbf{v}1$ is the point $\mathbf{p}$. We calculate $\mathbf{w}_{shear}$, which is parallel to the vector $\mathbf{v}1\mathbf{v}2$, see equation 1. Then, we calculate position of the point $\mathbf{z}'$ in the perceived space of the user $U2$, see equation 2.

$$\mathbf{w}_{shear} = \frac{|\mathbf{z}-\mathbf{p}|}{|\mathbf{v}1-\mathbf{p}|}(\mathbf{v}1 - v2) \quad (1)$$

$$\mathbf{z}' = \mathbf{z} + \mathbf{w}_{shear} \quad (2)$$

We use the point $\mathbf{z}'$ to calculate the vector $\mathbf{w}_{z-base}$, which is the z base vector of the coordinate system of the second user, given by: $\mathbf{w}_{z-base} = \mathbf{z} + \mathbf{w}_{shear}$. The amount of distortion for the whole scene can be encapsulated in the shearing matrix $\mathbf{M}_{shear}$, as shown in equation 3. Basically, we insert the vector $\mathbf{w}_{z-base}(s_{xz}, s_{yz}, s_{zz})$ at the place of the original z-base vector.

$$\mathbf{M}_{shearZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ s_{xz} & s_{yz} & s_{zz} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (3)$$

It is clear that in this case, the shearing distortion affects the rendering of the individual geometry vertices in the scene. The amount of shearing for each vertex is a function of the camera position $\mathbf{v}1$, the observer position $\mathbf{v}2$ and the z coordinates of the vertex. In general: the higher a point above the table or the larger the distance between $\mathbf{v}1$ and $\mathbf{v}2$, the larger the shearing distortion.

3.2. Viewpoint Compensation of the Geometry

In our approach, the correction matrix $\mathbf{M}_{warp}$ for counteracting the distortion, is obtained by $\mathbf{M}_{warp} = \mathbf{M}^{-1}_{shear}$, see Equation 4.

$$\mathbf{M}^{-1}_{shearZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \frac{-s_{xz}}{s_{zz}} & \frac{-s_{yz}}{s_{zz}} & \frac{1}{s_{zz}} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4)$$

This matrix can be used to *pre-distort* or to *pre-warp* the geometry in the scene space. We implemented the warping function as a dynamically updating scene graph node, which can be used to easily correct for distortion effects on various parts of the scene. We use the *warping matrix* to properly transform those elements in the scene that are associated with a specific user. This warping matrix allows us to transform points and geometry between the perceived space and the scene space.

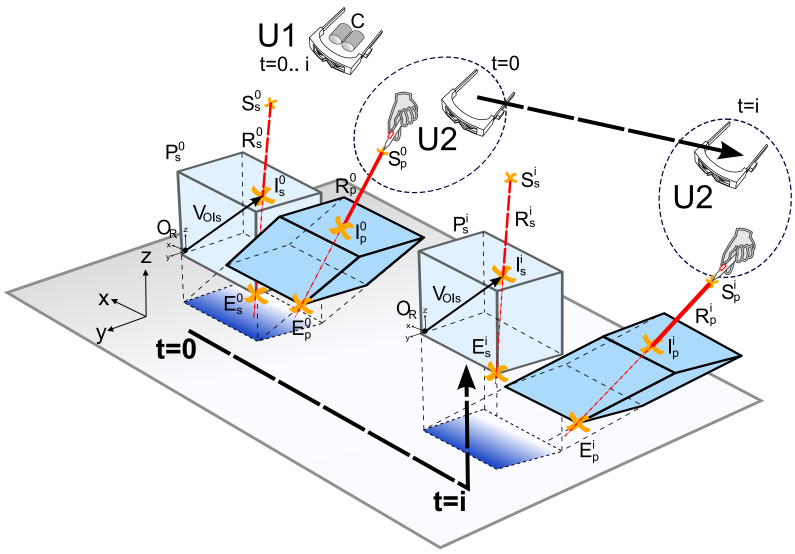

Figure 3: *Object translation in perceived space. U1 has the camera and is not moving; U2 has moved during manipulation time t_i causing the shearing to increase. Notation: P_s^0 - initial object pose, P_s^i - object pose at time i, $S_s^{0,i}$ - stylus position in scene space, $S_p^{0,i}$ - stylus position in perceived space, $R_s^{0,i}$ - ray in scene space, I_s - interaction point in scene space, E_s - exit point, $\mathbf{v}_{OIs}$ - vector: object origin to interaction point.*

4. View-Consistent Interaction

Having only a proper visual perception is not yet sufficient for consistent interaction from a different viewpoint. The problem is that user interactions are measured in the tracker space, which was during the workbench calibration aligned with the scene space. Due to the sheared perception of the scene (the perceived space), interactions do not match with the perceived space. Therefore the interaction handling has to be compensated for the distortion. Interaction techniques typically use transformation matrices and quaternion algebra, but assume orthonormality of the input data. As soon as we would apply distortion corrections (M_{warp}) on the stylus tracker data, orthonormality would be lost. As a result, shearing aspects would interfere with common scene graph calculations and would introduce unexpected deformation of objects in the scene. Therefore this has to be solved differently. We will describe now the adjustments needed to use the basic interaction techniques under shear-distorted viewing conditions.

4.1. Object Selection

For object selection, we use both direct and remote (ray-based) techniques. The *direct selection* technique uses a single 3D point provided by the stylus device. On account of possible viewing distortions, we need to multiply that point in perceived space S_p by the $\mathbf{M}_{shear}$ matrix to obtain the corresponding point in scene space S_s. The *ray-based selection* technique uses both a point and a direction vector provided by the stylus, both of which have to be transformed from tracking space into scene space, see Figure 4.

The (unit) direction vector $\mathbf{W}_S$ can be extracted from the orientation quaternion of the stylus $\mathbf{Q}_S$, see also Figure 8. To avoid distortion effects that could occur by warping the quaternion directly, we divide the calculation of the ray in a transformation of two points.

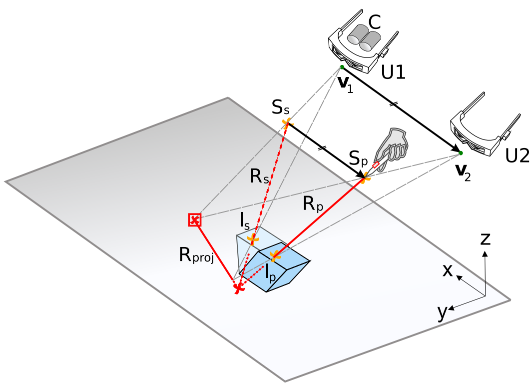

Figure 4: *Viewpoint compensation for ray-casting selection: The camera is at $U1$. S, R, I stand for stylus, ray, and intersection point; subscripts $-s, p, proj-$ stand for in scene, perceived and projected. U2 holds a stylus and points a ray at the cube. U2 should be seeing R_p as the ray coming from S_p in his perception of the scene. The ray's projection is R_{proj}. We correct for the viewing distortion by transforming the ray to R_s at another position S_s in the scene space.*

For this, we calculate the position of an arbitrary second point along the stylus ray. The two points, the stylus tip and

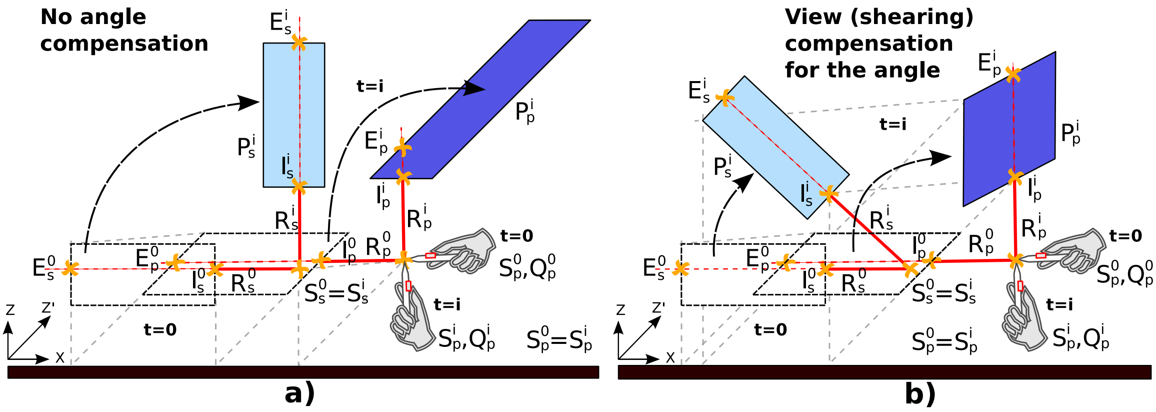

Figure 5: *Object rotation in sheared space: (a) The rotation angle applied on the object is not compensated for the shearing distortion. See the "'flip-flop"' effect. (b) The rotation angle is compensated and the stylus-object intersection line remains visually consistent in the perceived space.*

the second point, will be converted into the scene space. Based on these two points, the correct selection ray can be reconstructed in the scene space and the selection query on the scene graph can be performed. The selection routine returns: which object is selected and its interaction (intersection) point I_s in the scene space.

4.2. Object Manipulation: Translation

For basic object manipulation we use both direct and remote (ray-based) techniques. Once the object has been selected, both techniques use the same algorithm for manipulations. An important adjustment here is the separate treatment of translation and rotation transformations.

The reason for this is that the translation matrices remain pure translation matrices despite of shear distortion. For rotations this is not the case. The pose of an object P_s^i consists of a position (x,y,z) and orientation (quaternion $\mathbf{Q}_{obj}$). After object selection, we obtain the so-called *interaction point* I_s^0, see Figure 3. We have to compute its equivalent in the perceived space by $I_p^0 = I_s^0 \cdot \mathbf{M}_{shear}$. The stylus position S_p^0 comes from the tracker device. We store the distance vector R_p^0 between I_p^0 and S_p^0 and the vector $\mathbf{v}_{OIs}$ between the object origin and I_s^0. For direct manipulation, distance R_p^0 is zero, while for ray-based manipulation this corresponds to the length of the selection ray. During manipulation, we combine the stored $\mathbf{v}_{OIs}$ and R_p^i with the new stylus position S_p^i to translate the object to a new position P_s^i in the scene space, as show bellow.

$$P_s^i(x,y,z) = (S_p^i + R_p^0) \cdot \mathbf{M}_{warp} - \mathbf{v}_{OIs} \qquad (5)$$

In common ray-based manipulation, the length R_s of the ray remains constant in the scene. In our situation, the visible ray R_p needs to maintain constant length in the perceived space, that is $|R_p^0| = |R_p^i|$, while R_s will be changing accordingly. First, the distortion within the perceived space causes the R_s vector to change its length depending on its direction. Next, head movements of either user cause the amount of distortion in $\mathbf{M}_{shear}$ to change constantly. To compensate for this, we recalculate the correct length R_s^i of the ray in the scene by taking the distance between the current stylus position S_s^i and the unwarped interaction point I_s^i. An example of object manipulation is shown in Figure 3. Note that the user $U2$ has moved his head and the stylus. The shearing has increased between times $t_{0..i}$. To compensate for an induced rotation from our view compensation method, the user had to extra contra-rotate the stylus to keep the object initial orientation at the begin of the manipulation in the scene space. This effect will become more clear after explanation of the rotation in sheared spaces, see also Figures 6 and 7.

4.3. Object Manipulation: Rotation

To calculate a new orientation of the object during manipulation, we first store the quaternion difference $\mathbf{Q}_{delta}$ between the initial object orientation $\mathbf{Q}_{obj}^0$ and the stylus orientation $\mathbf{Q}_S^0$ in the scene. Then, at every manipulation step the current quaternion $\mathbf{Q}_P^i$ in the perceived space is warped to match the corresponding rotation in the scene space $\mathbf{Q}_S^i$ that is used for the actual interaction. To calculate a new object rotation around the point I_S^i, $\mathbf{Q}_S^i$ is multiplied with $\mathbf{Q}_{delta}$, see Figures 6 and 7. Note that the increased shearing in Figures 7b,c) compared to Figures 6b,c) has to be compensated with translation and extra rotation of the manipulated object to ensure a consistent appearance of the manipulated object during manipulation with the ray of the stylus.

Figure 5a) shows the situation in which the rotation of the stylus Q_P^i from the perceived space is directly applied to the object in the scene space to obtain a new object pose P_S^i. A 90 degrees turn with the stylus (in perceived space) will cause a 90 degrees rotation of the object in the scene space. However, due to the shearing of the perceived space, this rotation will introduce unexpected side effects. The angle between

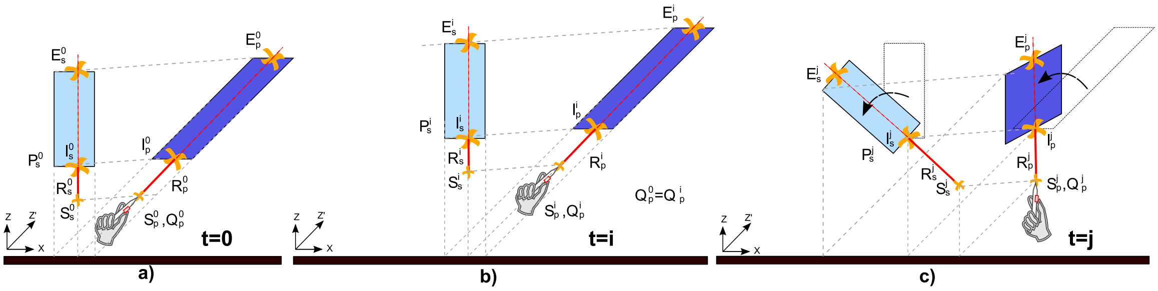

Figure 6: *Object manipulation under constant shearing: a) initial object selection, b) translation c) rotation.*

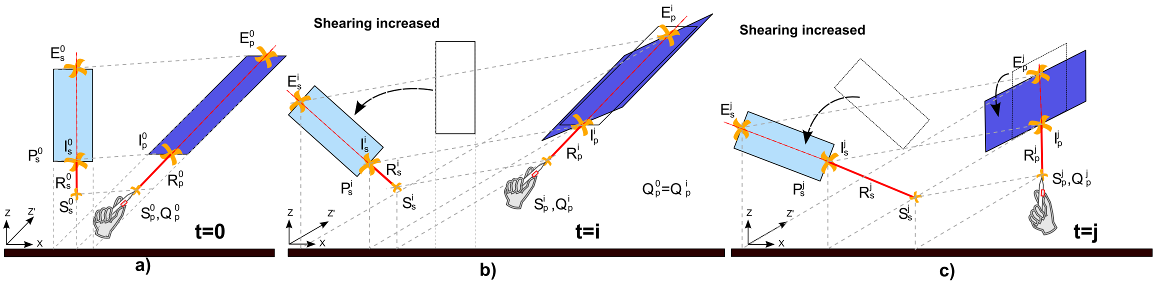

Figure 7: *Object manipulation when shearing has increased: a) initial object selection, b) translation c) rotation.*

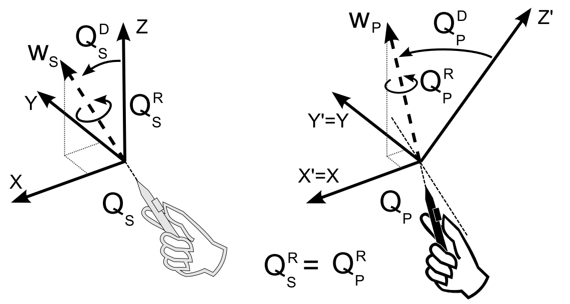

Figure 8: *Rotation quaternion of the stylus: (right) measured and seen in the perceived space $\mathbf{Q}_P$; (left) transformed into the scene space $\mathbf{Q}_S$, while compensating for the shearing distortion; $\mathbf{Q}^R_{P,S}$ is rotation around the stylus axis $W_{P,S}$. Hand with stylus exists only in the perceived (real) space.*

the stylus ray and the object do not remain visually constant. And the rotation axis will be variant. We call this a "flip-flop" or "bobbing" effect. Figure 5b) explains our shearing compensation for the rotation. Basically, we ensure that the ray-object intersection line with the entering point I_P and the exit point E_P remains visually persistent with any position and orientation of the stylus in the perceived space during manipulation. The effect is that a certain rotation angle in the perceived space will be transformed into a corresponding angle in the scene space depending on the actual amount of shearing distortion. In this example, a turn of 90 degrees in perceived space would result in one of 45 degrees in the scene space. Compare Figures 5a) and b) to see what happens with the exit point E_P relative to the object geometry. There will be no extra "flip-flop" rotation behavior around the intersection point I_P. Only the shearing during the object rotation or a change of $U1$ or $U2$ viewpoints will affect the location of the object corners in the perceived space.

The conversion of $\mathbf{Q}_P$ into $\mathbf{Q}_S$ is done as follows, see Figure 8. Stylus orientation in the perceived space $\mathbf{Q}_P$ measured by the tracking system can be written as combination of two rotations: $\mathbf{Q}_P = \mathbf{Q}^R_P \bullet \mathbf{Q}^D_P$, see Figure 8 right). In a similar way the stylus orientation quaternion in the scene space $\mathbf{Q}_S$ can be written as: $\mathbf{Q}_S = \mathbf{Q}^R_S \bullet \mathbf{Q}^D_S$, , see Figure 8 left). We measure $\mathbf{Q}_P$ and we want to obtain $\mathbf{Q}_S$. The visual (perceived) direction vector of the stylus W_S is calculated from the measured W_P by warping this vector into the scene space (using $\mathbf{M}_{warp}$ matrix). The rotation around the stylus axis (W_P or W_S, resp.) is invariant to the existing form of shearing and therefore $\mathbf{Q}^R_P = \mathbf{Q}^R_S$. The rotation $\mathbf{Q}^D_P$ is calculated between vectors Z' and W_P. Note that Z' has to be updated every frame according to the actual shearing caused by changing positions of $U1$ and $U2$. The final solution for $\mathbf{Q}_S$ is outlined as:

$$\mathbf{Q}_S = \mathbf{Q}_P \bullet (\mathbf{Q}^D_P)^{-1} \bullet \mathbf{Q}^D_S \quad (6)$$

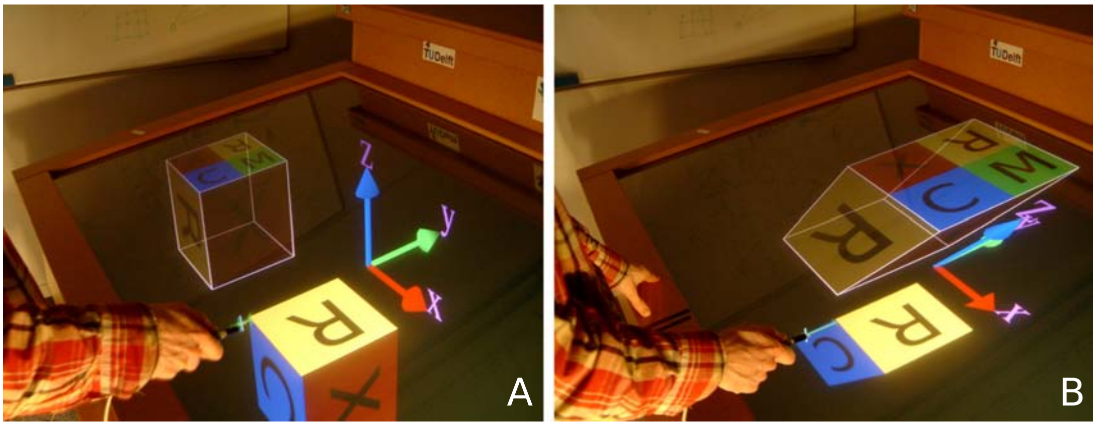

Figure 9: *Initial state of the docking task - cube selection; (a) view without sheared distortion, (b) sheared-distorted view, proportional to 1 meter headtracker offset from the correct viewpoint; viewpoint compensantion applied.*

5. Results

To show the issues related to the rotation in sheared perceived VE's, we have prepared a simple docking task, see Figure 9. The scene is rendered from a viewpoint that is 1 meter from the interacting user, see Figure 9b). The goal of the docking task is to orient and place the cube on the right side, into the transparent cube on the left. For the sake of clarity, the user is selecting the object approximately at one corner with the ray visually aligned with one edge of the object. In this situation, the ray points in the Y direction. This means, that for this task, the final orientation of the stylus will be pointing upwards in the perceived direction of the Z-axis. Figure 9a) shows the situation when the user has a correct view and this task will be rather simple without any conflicting cues. The situation in which the user does not have a correct view and the viewpoint compensation is applied, is shown in Figure 9b). The user interaction is made consistent with the sheared perception of the world. His interaction device will remain co-located with the selection/manipulation ray and the interaction icons. Basically, the user is able to select the objects with the same comfort as in the orthogonal-perceived space. The object translations are made consistent with the user's view on the world, therefore going up along the sheared Z-axis will correspond to a perpendicular Z translation in the scene space. We should note that strictly horizontal manipulations are not affected by the shearing effect.

As the user proceeds with this task, larger rotations are applied on the object and certain problems become more obvious. Figure 10 shows what happens in a strongly sheared-perceived view, when the rotation is not corrected for the shearing effects. It is very clear that the visual alignment between the stylus ray and the cube edge is lost during the upwards rotation. From our experiments we know that this alignment error occurs only when the rotation is applied around X axis(pitch) or Y axis(roll). Due to the logic of the Euler angles and the specifics of this visual shearing the heading rotation (around scene Z axis) is not affected. Therefore we have chosen a rotational task that explicitly requires the "critical" rotations. The observed "flip-flop" rotational behavior around the intersection point I_p is strongly inconsistent with the users' expectations and rather contra-intuitive, having several conflicting cues. Figure 10b) shows actually how difficult it is for the user to perform the final docking step. He has to align the cube edge with the scene Z-axis going perpendicularly up, while the other cues show that this edge and actually also the Z-axis are sheared in a different direction.

We will now compare the same task with our angular shearing correction, see Figure 11. The results are convincing. Absolutely no 'flip-flop" rotational behavior around the intersection point I_p is observed. The edge remains aligned with the ray all the way during the manipulation task. The final docking alignment is much simpler because the cues provided are consistent with each other. The world Z-axis looks and "feels" sheared in the same way during interaction. This seems to be less confusing than without our correction method. There is however one aspect that we cannot correct completely. This is the inherent shearing in the perception of the virtual world. Rotation in sheared coordinate spaces will always produce strange visual effects. But the good news is that whatever visual shearing is observed during manipulation (both translation and rotation), the object in the scene space will never become sheared. The user has to understand that compensation effects will appear in the form of additional translation and rotation of the manipulated object to keep this object consistent with the sheared view.

The photo series in Figure 12 show the viewpoint compensation in action. The left user is manipulating a cube in the scene, which is rendered from a viewpoint of the right user. Actions of the left user are compensated for the shear-

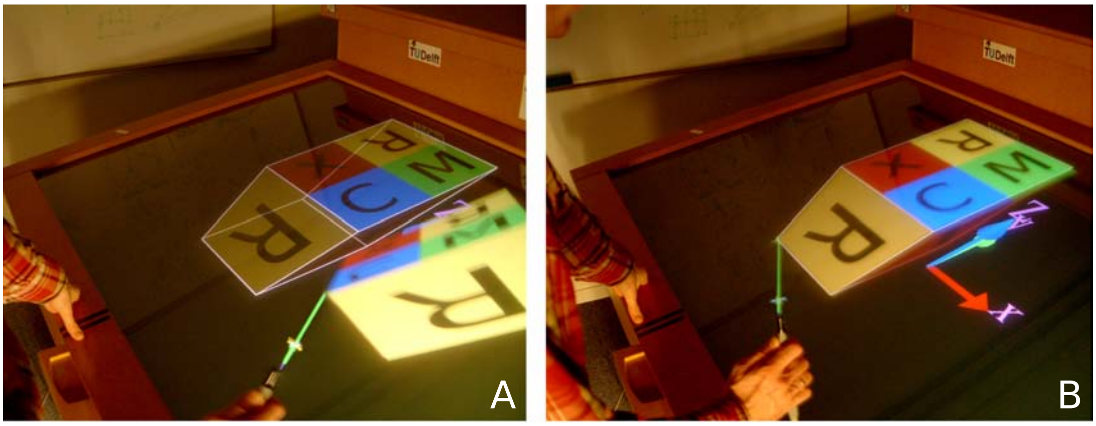

Figure 10: *Sheared view: Rotation in the docking task without shearing correction for rotation. (a) Intermediate state, (b) final manipulation state; Note the* ***rotation alignment error*** *between the stylus ray and the edge of the manipulated cube.*

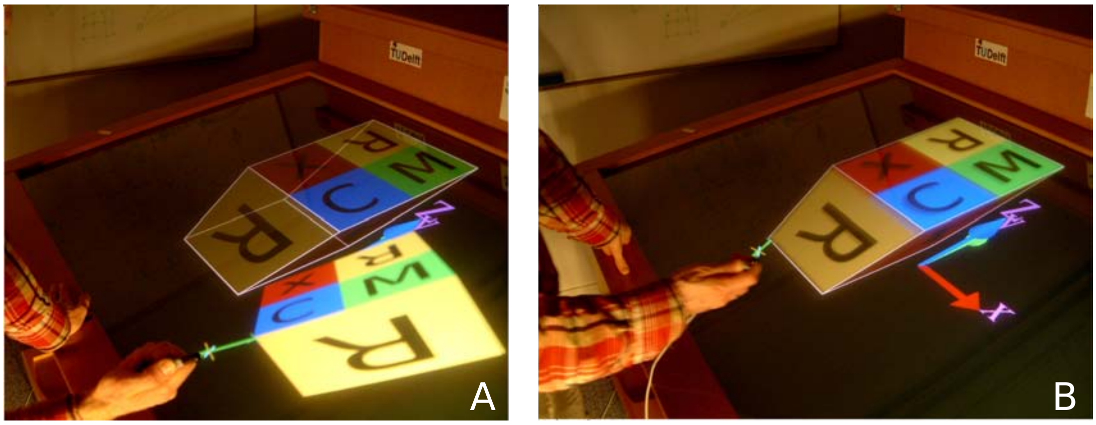

Figure 11: *Sheared view: Rotation in the docking task with shearing compensation for rotation. (a) Intermediate state, (b) final manipulation state; Note the* ***correct rotation alignment*** *between the stylus ray and the edge of the manipulated cube.*

ing distortion and are consistent with his view perception. Both users can dynamically change their viewpoints. Therefore, the amount of actual shearing distortion is dynamic as well. The viewpoint compensation ensures a constant alignment of the interaction ray and device with the manipulated object. The rotation angles work according to the perceived space.

6. Conclusions and Future Work

We developed multi-user stereo display techniques, where the camera is dynamically placed in between users and interaction is based on their current viewpoint. Any user-private-view objects: interaction cursors, ray, icons, etc. can be rendered without distortions for a specific user, while the rest of the scene is rendered according to the viewpoint of a shared camera. The interaction is compensated for in such a way that the user, despite seeing shearing distortions in scene elements, can still manipulate the scene as if the viewing were correct. Without these compensations, the user would have difficulty interacting with the VR, see our user study in [dHMKP07].

Working with correctly perceived tools in distorted space still has some side effects. The user has to be aware that up and down directions in the sheared-perceived world are actually slanted. It is therefore necessary to provide additional visual and interaction cues, providing appropriate feedback.

In this paper we extensively described additional compensations required in sheared space to make rotation predictable and visually consistent. We explained the details of the problems with selection, translation and rotation in this situation.

Without viewpoint compensation during rotational tasks, visual inconsistencies occur, see Figure 12 b,c). Our methods serve to mitigate these effects, object deformations still occur, see Figure 12 e,f).

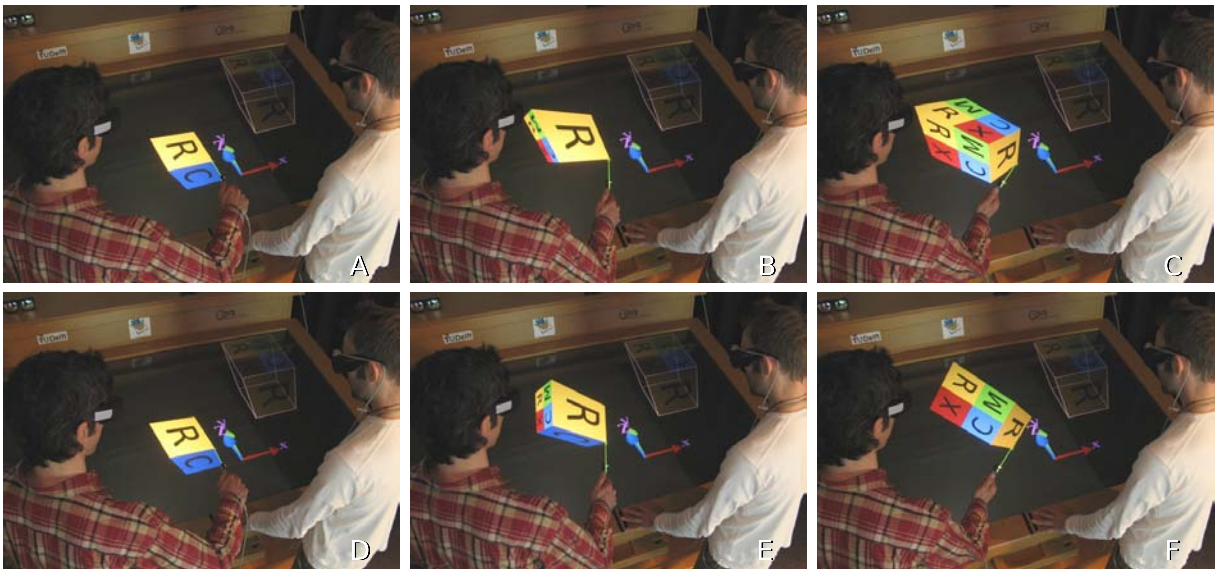

Figure 12: *Viewpoint compensation method for rotation in a two-user scenario at the Workbench. The scene is rendered from the right user's viewpoint and the viewpoint compensation is done for the photo camera position for the purpose of this presentation. The left user is translating and rotating an object in a visually sheared VE. Sub-figures a,d) show the situation at the moment of object selection. Note, that the interaction ray is aligned with one of the edges. Sub-figures a,b,c) show the case where no viewpoint compensation for stylus rotation is used. See the "flip-flop" effect. Sub-figures d,e,f) show the compensated case. Both translations and rotations behave consistently with the sheared perception of the scene.*

With our viewing methods, multiple users can successfully interact with one virtual display system. We are currently planning qualitative and quantitative user evaluations for rotational docking tasks to see if the additional compensation techniques improve the user's experience and performance.

Acknowledgements

Part of research has been funded by the Dutch BSIK/BRICKS project. This work was also carried out in the context of the Virtual Laboratory for e-Science project (www.vl-e.nl). This project is supported by a BSIK grant from the Dutch Ministry of Education, Culture and Science (OCW) and is part of the ICT innovation program of the Ministry of Economic Affairs (EZ).

References

[ABM*97] AGRAWALA M., BEERS A. C., MCDOWALL I., FROEHLICH B., BOLAS M., HANRAHAN P.: The two-user responsive workbench: support for collaboration through individual views of a shared space. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1997), ACM Press/Addison-Wesley Publishing Co., pp. 327–332.

[APT*98] ARTHUR K., PRESTON T., TAYLOR R., BROOKS F., WHITTON M., WRIGHT W.: *Designing and Building the PIT: a Head-Tracked Stereo Workspace for Two Users*. Tech. rep., University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, 1998.

[BFSE01] BIMBER O., FRÖHLICH B., SCHMALSTIEG D., ENCARNACAO L. M.: The virtual showcase. *IEEE Computer Graphics and Applications 21*, 6 (2001), 48–55.

[BLCN02] BLOM K., LINDAHL G., CRUZ-NEIRA C.: Multiple active viewers in projection-based immersive environments. In *Proceedings, Seventh Annual Symposium on Immersive Projection Technology (IPT 2002)* (March 24 - 25, 2002).

[BMC04] BOLAS M., MCDOWALL I., CORR D.: New research and explorations into multiuser immersive display systems. *IEEE Comput. Graph. Appl. 24*, 1 (2004), 18–21.

[dHMKP07] DE HAAN G., MOLENAAR R., KOUTEK M., POST F. H.: Consistent Viewing and Interaction for Multiple Users in Projection-Based VR Systems. *Computer Graphics Forum, to appear 26*, 3 (2007).

[FHH*05] FROEHLICH B., HOCHSTRATE J., HOFFMANN J., KLUEGER K., BLACH R., BUES M., STEFANI O.: Implementing multi-viewer stereo displays. In *WSCG*

'05, International Conference in Central Europe on Computer Graphics and Visualization (2005), pp. 139–146.

[KKYK01] KITAMURA Y., KONISHI T., YAMAMOTO S., KISHINO F.: Interactive stereoscopic display for three or more users. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 231–240.

[MB97] MCDOWALL I. E., BOLAS M. T.: New developments for virtual model displays. *SIGGRAPH Comput. Graph. 31*, 2 (1997), 50–52.

[MB04] MULDER J. D., BOSCHKER B. R.: A modular system for collaborative desktop vr/ar with a shared workspace. In *VR '04: Proceedings of the IEEE Virtual Reality 2004 (VR'04)* (Washington, DC, USA, 2004), IEEE Computer Society, p. 75.

[Mol07] MOLENAAR R.: Viewing and Interaction for Multiple Users in Projection-Based VR Systems. In *MSC. thesis, Delft University of Technology, Data Visualization Group* (april 2007).

[Sim05] SIMON A.: First-person experience and usability of co-located interaction in a projection-based virtual environment. In *VRST '05: Proceedings of the ACM symposium on Virtual reality software and technology* (New York, NY, USA, 2005), ACM Press, pp. 23–30.

[Sim07] SIMON A.: Usability of multiviewpoint images for spatial interaction in projection-based display systems. *IEEE Trans. Vis. Comput. Graph. 13*, 1 (2007), 26–33.

[SS05] SIMON A., SCHOLZ S.: Multi-viewpoint images for multi-user interaction. In *VR '05: Proceedings of the 2005 IEEE Conference 2005 on Virtual Reality* (Washington, DC, USA, 2005), IEEE Computer Society, pp. 107–113.

[WHR99] WARTELL Z. J., HODGES L. F., RIBARSKY W.: The analytic distortion induced by false-eye separation in head-tracked stereoscopic displays. In *GVU Lab Technical Report, no. 99-01, 1999.* (1999).