

C1x6: A Stereoscopic Six-User Display for Co-located Collaboration in Shared Virtual Environments

Alexander Kulik^{*,1} André Kunert^{*,1} Stephan Beck^{*,1} Roman Reichel^{*,1} Roland Blach^{†,2}
Armin Zink^{‡,3} Bernd Froehlich^{*,1}
¹ Bauhaus-Universität Weimar ² Fraunhofer IAO ³ digital IMAGE

Abstract

Stereoscopic multi-user systems provide multiple users with individual views of a virtual environment. We developed a new projection-based stereoscopic display for six users, which employs six customized DLP projectors for fast time-sequential image display in combination with polarization. Our intelligent high-speed shutter glasses can be programmed from the application to adapt to the situation. For instance, it does this by staying open if users do not look at the projection screen or switch to a VIP high brightness mode if less than six users use the system. Each user is tracked and can move freely in front of the display while perceiving perspective correct views of the virtual environment.

Navigating a group of six users through a virtual world leads to situations in which the group will not fit through spatial constrictions. Our augmented group navigation techniques ameliorate this situation by fading out obstacles or by slightly redirecting individual users along a collision-free path. While redirection goes mostly unnoticed, both techniques temporarily give up the notion of a consistent shared space. Our user study confirms that users generally prefer this trade-off over naïve approaches.

CR Categories: B4.2 [INPUT/OUTPUT AND DATA COMMUNICATIONS]: Input/Output Devices—Image Display; H5.1 [INFORMATION INTERFACES AND PRESENTATION]: Multimedia Information Systems—Artificial, augmented, and virtual realities

Keywords: Virtual Reality, Display Technology, Collaboration

Links:  DL  PDF  WEB

1 Introduction

3D television sets and 3D cinemas display only a single stereoscopic image stream, and thus there is only a single location from which a person observes a perspective correct view of the displayed scenes. All of the other spectators perceive the 3D scene more or less distorted, which inhibits precise spatial perception of the displayed geometry. While this may not matter much in movie theaters, these distortions significantly hamper the acceptance of

3D technology in many other application areas. In order to compensate for this, each user must be provided with individual stereoscopic image pairs, which are rendered for the exact position of the user in front of a display. While the computing power to generate multiple views of interactive content is available, the display technology for presenting large individual stereoscopic images for multiple users is still lacking

We developed the C1x6, a new projection-based stereoscopic display for six users (Figure 1). Our system consists of six customized DLP projectors, each of which projects six fast time-sequential images in one of the primary colors. By differently polarizing the light output of the first set of three single color projectors (red, green, blue) than those of the second set, we are able to project twelve separable full-color images onto a projection screen. Our intelligent high-speed shutter glasses can be fully controlled from the application level. This feature is used to keep the glasses open if users look away from the screen or for supporting a VIP high brightness mode if less than six users are involved. We developed the software and hardware infrastructure to generate, warp and feed the stereoscopic images for the six tracked users into the projectors.

Multi-user displays enable co-located collaborative work in shared virtual environments. For collaborative design reviews we developed the Spheron, an input device which makes interactions transparent to other co-located users. However, when navigating a group of users through a virtual building, many situations arise in which there is not enough space to place the users in the virtual world in the same way as they are positioned relative to each other in the real world. This problem did not exist in common projection-based virtual reality systems, where all the observers share the same perspective as the head-tracked navigator. Therefore, we present several approaches to facilitate group navigation in such situations by avoiding collisions of group members with surrounding objects such as walls and other obstacles.

The main contributions of our work fall into three areas:

- A synchronized DLP projector array that is capable of displaying twelve high resolution (1920x1200) full color image streams at 360Hz—60Hz per user. Left and right eye images are simultaneously projected and separated by polarization.
- Application-level programmable shutter glasses consisting of double-cell liquid crystals that enable intelligent shutter control and provide fast switching speeds as well as high-contrast.
- Augmented group navigation techniques that avoid collisions when traveling through narrow spaces. Our user study reveals that these techniques are preferred over naïve approaches.

Besides these central contributions, significant amounts of engineering are necessary to build and run such a complex system, including a custom digital video-multiplexing hardware, synchronization of all the components and real-time color convergence through image warping. Measurements show that our six-user system achieves almost the same brightness per user as a stereoscopic single-user display based on the same type of projector would.

*e-mail:<first name>.<last name>@uni-weimar.de

†email:roland.blach@iao.fraunhofer.de

‡email:zink@digital-image.de

ACM Reference Format

Kulik, A., Kunert, A., Beck, S., Reichel, R., Blach, R., Zink, A., Froehlich, B. 2011. C1x6: A Stereoscopic Six-User Display for Co-located Collaboration in Shared Virtual Environments. *ACM Trans. Graph.* 30, 6, Article 188 (December 2011), 11 pages. DOI = 10.1145/2024156.2024222 <http://doi.acm.org/10.1145/2024156.2024222>

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2011 ACM 0730-0301/2011/12-ART188 \$10.00 DOI 10.1145/2024156.2024222 <http://doi.acm.org/10.1145/2024156.2024222>

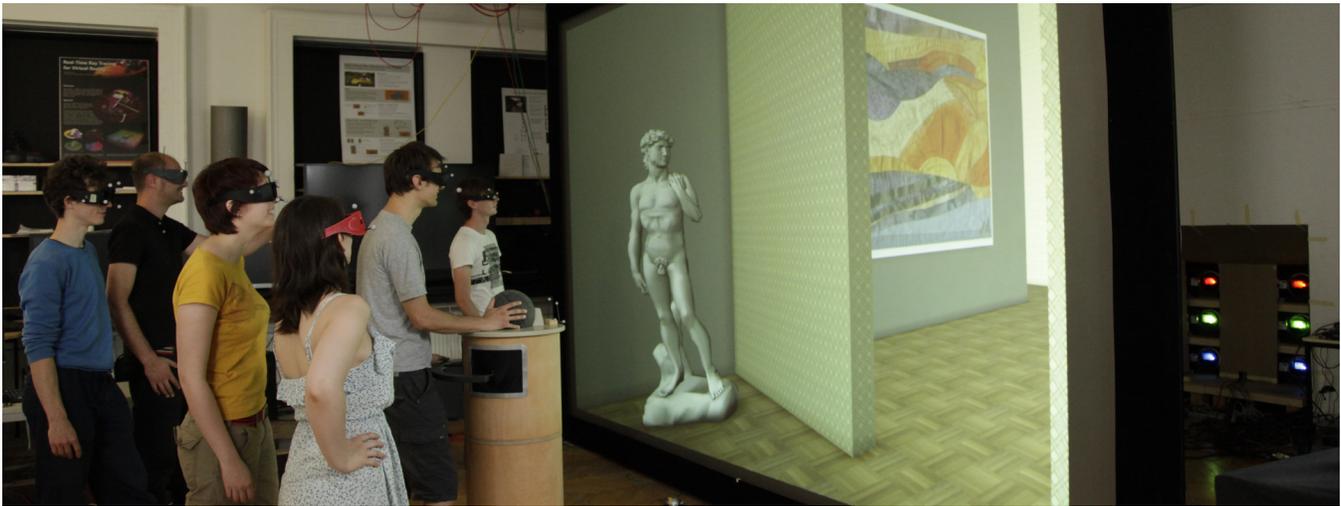


Figure 1: *The six-user projection system. The lenses of the six projectors appear as bright color spots on the right side. Each of the six users is tracked and provided with a perspectively correct image. The Spheron, our group navigation device, is centrally placed in front of the display. 12 different images are projected onto the screen, only one image is shown here.*

2 Related Work

The most straightforward way to provide multiple users with individual views of a shared virtual world is the use of personal displays such as head-mounted displays (HMDs) or handheld displays. The Studierstube system by [Schmalstieg et al. 2002] supported co-located collaborative augmented reality using see-through HMDs. [Aspin and Roberts 2005] performed a co-located architectural design review by providing multiple users with HMDs. [Hua et al. 2004] equipped multiple users with head-mounted projectors, which projected onto retro-reflective walls.

Projection-based stereoscopic displays such as the CAVE [Cruz-Neira et al. 1992] have a long tradition in the virtual reality domain, but there have only been a few approaches providing multiple tracked users with individual stereoscopic images. The two-user Responsive Workbench [Agrawala et al. 1997] displays four different images in sequence on a CRT projector at 144Hz, which results in 36Hz per eye per user. They also developed custom shutter glasses for cycling between four eyes. This system was the first demonstration of a two-user system, but suffered from flicker, low brightness and crosstalk. [Barco 1999] developed the "Virtual Surgery Table", which provides two users with stereoscopic images by differently polarizing the light output of two active stereo projectors. This approach was also used for a large projection wall in [Riege et al. 2006]. All these systems are limited to two users and cannot be easily extended to support more users. [Bolas et al. 2004] presented a modified DLP projector, which is capable of running at 120Hz. They also briefly mention the integration of such single-chip displays with the optics of a 3-chip DLP system to achieve a three-user system running at 120 Hz per eye in a single projector. Unfortunately, no technical details were provided. Our approach is similar in that it also uses customized 120Hz projectors, but we realized a complete fully synchronized six-projector system and demonstrate its use for six users.

[Kunz and Spagno 2002] employed a pair of shuttered LCD projectors to generate an active stereo display for their blue-c system [Gross et al. 2003]. [Fröhlich et al. 2005] extended this approach to support four users by using eight shuttered LCD projectors. To limit the shutter frequency, they used shuttering to cycle among the users and polarization for separating the left and right eye images.

However, shuttering projectors is not a very light-efficient approach considering that each projector is blocked for most of the time. For example, in the four user setup, each projector is blocked for three-quarters of the time and thus 75% of the light output is lost. This is also the reason why it does not scale well to more users.

There are also a number of special purpose multi-viewer displays. The PIT [Arthur et al. 1998], the Illusionhole [Kitamura et al. 2001] and the Virtual Showcase [Bimber et al. 2001] use different approaches to assign a separate partition of a projection screen to each user, where the stereoscopic images for each person are displayed. The PIT uses two orthogonal screens, in which each user looks at only one of the screens. The Illusionhole uses a circular mask on top of a tabletop projection. By looking through the mask the users positioned around the table see different areas of the screen, where their individual images are presented. The Virtual Showcase consists of a tabletop projection with a truncated half-silver mirror pyramid (or cone) placed in the middle of the table. By looking into the mirror, users positioned around the table see a reflected image off the tabletop. The stereoscopic images are rendered such that the virtual objects appear inside this Virtual Showcase. The Joint Space Station [Mulder and Boschker 2004] uses a similar approach based on separate Virtual Workbench displays [Poston and Serra 1994] facing each other. These displays are limited to two to four users with a small overlap of the users' viewing frustra, which considerably limits the size of the objects that can be displayed in the shared virtual space. [Maksakov et al. 2010] simply used separate viewports on a larger screen for each user to provide individual head-tracked monoscopic views of a 3D scene. Such approaches effectively discard the notion of a locally shared space and require similar interaction techniques and affordances as do distributed multi-viewer systems.

[Dodgson 2005] and [Favalora 2005] provided an introduction and overview of the many other types of multi-view displays, in particular autostereoscopic and holographic systems. While the use of such technology for displaying large, interactive and full color 3D images for multiple non-stationary users is the ultimate goal, all of these systems pose different limitations. However, various recent developments are convincing solutions for particular application domains. [Cossairt et al. 2007] and [Jones et al. 2007] developed similar approaches for occlusion-capable parallax multi-view

3D displays. Their systems use modified DLP projectors to project fast time-sequential images onto a rotating anisotropic projection surface. Both systems achieve about one degree of angular resolution and support a 180 and 360 degree field of view, respectively. [Jones et al. 2009a] reported on a further refined prototype of such a system and showed its use in a very convincing real-time one-to-many teleconferencing application [Jones et al. 2009b]. Due to the use of a single projector, the bandwidth of these systems is limited, which results in a small color depth of one bit color or even only black and white depending on the used DLP projector type. In addition, such a system is difficult to scale to a larger size due to the rotating display surface.

The research surrounding collaborative virtual environments (CVEs) has mostly focused on distributed collaboration (e.g. [Benford et al. 2001] reviews the history of CVEs). [Otto et al. 2006] and [Wolff et al. 2007] provided a solid analysis of the requirements for supporting closely coupled collaborative tasks in a shared virtual workspace for non-co-located users, which also apply to a certain extent to co-located collaboration. However, there is limited work on co-located collaboration in projection-based multi-user virtual reality. The original two-user Responsive Workbench work [Agrawala et al. 1997] suggested the use of specialized views, which were used to provide different information to each user, as in a teacher-student scenario. [Riege et al. 2006] suggested the use of a bent pickray to visualize the constraints that are involved when two users are jointly manipulating an object with six degrees of freedom. [d’Angelo et al. 2008] showed that stereoscopic display in combination with collaborative manipulation improve task performance and are clearly preferred in a complex assembly task involving two users. [Argelaguet et al. 2010] demonstrated the use of specialized views to reduce the problem of interpersonal occlusion. All these approaches consider only two collaborating users and focus on joint manipulation. It is not clear how these approaches scale to more users.

[Bowman et al. 2005] provide an overview and introduce a taxonomy for the large variety of navigation techniques for virtual environments. However, the problem of navigating multiple co-located users with individual views through a shared virtual world has not yet been addressed. Group navigation as it is defined here – moving multiple people simultaneously through a virtual environment – is a new problem that is closely linked to the introduction of stereoscopic multi-viewer systems. Augmented group navigation techniques to mitigate associated issues are orthogonal to general single-user navigation techniques. In our setup each head-tracked person can independently walk in front of the display, but apart from that, does not independently travel within the environment since otherwise the group would no longer share a consistent virtual space.

3 Synchronized 12-View Projector Array

Our goal was to build a fast time-sequential full color DLP-based system which also exploits polarization. Our approach is based on the following ideas:

- Color wheel DLP projectors project the different primary colors as fast time-sequential images. There are various color wheel versions; we assume a basic three-segment color wheel consisting of three color filters, one in each primary color: red, green and blue. If the color wheel is removed, we can project three monochrome time-sequential views (Figure 2) instead of the different primary colors of a single view. By using three projectors and equipping each projector with a primary color filter, we regain full color images for three views.

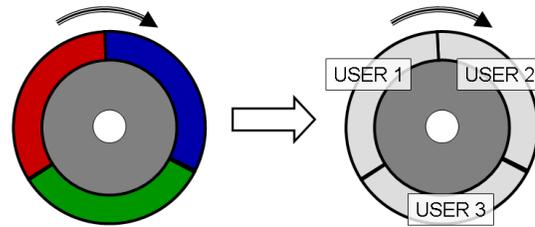


Figure 2: A three-segment color wheel. We display individual images for three eyes instead of time-sequential colors.

- Most DLP projectors rotate the color wheel at least twice per video frame and are thus effectively running at 120Hz while 60Hz input is provided. However, at the time of our development, a 1920x1200 pixels resolution projector was not yet available, which would accept a 120Hz stereo signal. Thus we had to extend an existing projector to process a 120Hz image stream or to interleave two 60Hz streams. This way we could project six different views at 360Hz (three views times two rotations times 60Hz).
- Polarization can be effectively used in combination with shuttering to double the number of views, thus allowing 12 views to be achieved using two times three projectors.

Such a system maintains the brightness of a single user active stereo system since we are using six projectors for six users. In addition, we retain full color depth, full resolution (1920x1200) and a 60Hz refresh rate. Figure 3 shows an overview of our setup.

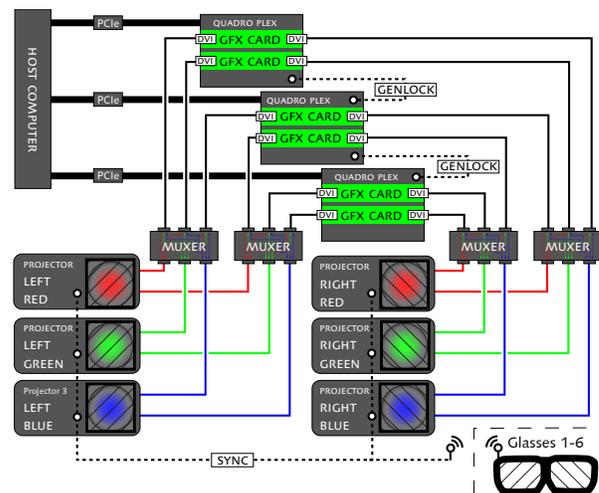


Figure 3: The projector array is driven by a single computer. Three synchronized NVIDIA Quadro Plex 7000 graphics systems are connected to the host computer via separate PCIe interfaces. Each Quadro Plex consists of two graphics cards with two DVI outputs each. It produces the left and right eye images for two users. Sets of three DVI outputs carrying the images for three eyes are connected to the video multiplexers (muxer), which rebin the image streams by color and send them to the respective projectors. The left three projectors display the left eye images for the six users, while the right three projectors display the right eye images. The two sets of projectors emit differently polarized light which matches the polarization of the users’ left and right eye shutters. External synchronization is provided to the projectors and to the radio-controlled shutter glasses.

In building our 12 view projection system, we had to develop the following main components:

- A 360Hz projector with external synchronization capable of projecting six different views.
- Video-multiplexing hardware for feeding each projector with six different views, which are generated on different graphics cards.
- Real-time image warping for geometric alignment of images, since the color components of the images are projected from different projectors.

3.1 The 360Hz Projector

We modified six projectiondesign F32 DLP projectors [Projectiondesign 2011] to accept two 60Hz input streams and alternate the display of these image streams at 120Hz using an approach similar to the one presented in [Hopp 2005]. The F32 projector contains a separate input and scaler board, which accepts the DVI image stream and scales it to the resolution of the Digital Micromirror Device (DMD). The resulting 1920x1200 images are sent on to the formatter board, which reformats the image stream and sends it to the DMD. The main modification was the addition of a second input and scaler board, as well as a formatter board. The second input pipeline accepts a second image stream. From the formatter boards, the image stream is sent to the DMD via an LVDS interface. We added an LVDS switch in front of the DMD, which is connected to the two formatter boards for switching between the two image streams as shown in Figure 4. In our experiments we found that the F32 projector uses exactly the same DMD patterns for the first and second rotations of the color wheel. Thus there is no reduction in color depth involved by using two different input streams. Our projectors have the basic three-segment color-wheel firmware installed.

Doubling the input and scaler boards and switching between these two input streams is necessary if the projector accepts only a monoscopic image stream at e.g. 60Hz. This major hardware modification is not needed if a stereoscopic projector is available, which uses alternating rotations of the color wheel for projecting the left and right eye images. However, it also needs to have a three-segment RGB color wheel firmware installed for mapping each of the primary colors to a different user.

A dedicated IO board is used in the F32 projector for taking care of the fans, lamp power control and temperature sensors. IO status is routed from the IO board through the formatter board to the input and scaler board. We developed a microcontroller board for interfacing with the IO board, and doubling its inputs and outputs to keep the second input pipeline alive. The microcontroller operates the crossbar and its power supply.

The last major modification of the projector involves the color wheel. Since the F32 projector uses two lamps, there were also two color wheels. Since we only project a single color per projector, we removed the color wheels and use a fixed color filter in one of the primary colors. Cooling and infrared filtering had to be added to avoid heat problems. Color wheel projectors receive their internal sync from the rotating color wheel. We provide an external synchronization signal to all the projectors to keep them in perfect sync. The same sync is also provided to our custom shutter glasses.

3.2 Video-Multiplexing and Geometric Alignment

We built a digital video-multiplexer hardware, which takes three full color image streams from the graphics cards, rebins them by

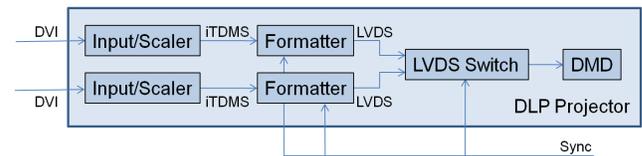


Figure 4: Main electronics components and signal flow inside our modified DLP projector. Two synchronized 60Hz DVI signals are fed into the projector. The input and scaler board adjusts the size and refresh rate of an input stream to the requirements of the formatter board and sends it via an iTDMS connection to the formatter board. The formatter board formats the bitstream for the DMD chip and sends it to the DMD via an LVDS link. The LVDS switch alternates between the two LVDS streams from the formatter boards. The sync signal is externally generated and keeps all the components in sync.

color and sends them to the respective projectors. We use four of our video-multiplexer units to route the color components of the 12 DVI streams to the respective projectors (Figure 3).

Since the primary colors are displayed by different projectors, the colors of the projected images will be misaligned. Although shift lenses and software keystone correction can help in some configurations, the overall image quality is not acceptable for our needs. Therefore we use a real-time image warping technique that warps each pixel of a rendered image to the correct location on the screen. The multiplexers route each image to three different projectors and thus each color component has to be warped to a different location to achieve color convergence. This warping step is based on individually precomputed look-up tables, one for each projector, and is performed in real-time as a post processing step in a fragment shader program.

Our algorithm automatically calibrates multiple projectors onto a single target projector. This involves several computation steps. First, we capture Gray code patterns with a monochromatic 5 megapixel camera to derive camera-to-projector maps for each projector, similar to [Damera-Venkata and Chang 2007]. Based on these maps, our algorithm computes the largest projection area that can be covered by each projector. Finally, look-up tables are computed for each projector so that every location within the projection area maps to a location of a rendered image. In this calibration process, the acquisition of the Gray code patterns is the most crucial step. For achieving a good signal-to-noise ratio, our calibration tool simultaneously renders the Gray code patterns on all six graphics cards contributing to the projection of a single projector. We achieve a precise per-pixel alignment of the six projector images throughout the projection area.

4 Intelligent Shutter Glasses

Shutter glasses are unavoidable in our approach. They need to work at 360Hz and the left and right eyes need to be differently polarized. Since shutter glasses consist of two crossed polarizers with liquid crystal material in between, the left and right eye shutters just need to be rotated against each other by 90 degrees to achieve orthogonal filtering capabilities. Regular liquid crystal shutters are not suitable for our system since they have asymmetric opening and closing properties. They close quickly (e.g. less than 0.2ms depending on the operating voltage) and open slowly (e.g. longer than 2ms), too slow for 360Hz cycles. The standard solution to circumvent this problem is the use of ferro-electric (FLC) shutters with symmetric opening and closing times of less than 0.1ms. While such a solution is easily capable of running at 360Hz, FLCs are much more expensive.

sive than standard liquid crystal (LC) shutters and they are very fragile. FLCs are also designed to work with symmetric open/close timings, which is not the case in our setup.

As an alternative to FLCs, we built our shutter glasses based on a novel double-cell shutter design, which consists of two layers of differently configured regular LC shutters. The first layer is a regularly cross-polarized LC shutter (normally white (NW)), which is transparent if no voltage is applied. The second layer has equally oriented polarization filters on both sides and thus it is opaque (normally black (NB)) if no voltage is applied. This combination of shutters functions so that the NB shutter opens quickly while the NW shutter closes quickly. These shutters are ideally suited for an asymmetric use case: our shutters need to be open for only $1/360th$ of a second and closed for $5/360th$ of a second. During the longer closing time, both shutters relax one after the other: first, the NB shutter closes fully and then the NW shutter opens completely (Figure 5). In addition, using a stack of two shutters improves the contrast ratio, an important property in the context of our system.

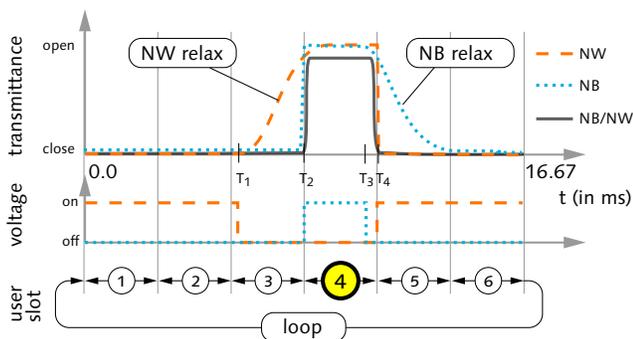


Figure 5: Illustration of the double shutter functionality (top), the electrical shutter driving pattern (middle) and the time slots for each user (bottom). At 60Hz a 16.67ms time frame is divided into six adjacent user time slots of equal length. The diagram shows the timings of the double cell shutter of user 4, who receives an image during the fourth slot lasting from $T_2=8.34ms$ to $T_4=11.12ms$. The NW shutter is switched off at T_1 about 2ms before T_2 to ensure its relaxation and thus maximum light transmission at the beginning of the following opening period. The NB shutter is still blocking light during this NW relaxation phase and immediately opens when the voltage is applied at the beginning of the 4th time slot (T_2). At the end of the opening period of 2.78ms (T_4) the NW shutter is immediately blocking the light transmission as the voltage is applied. The NB shutter is switched off for relaxation slightly before ($T_3=11ms$).

In a six-user stereo-projection system each individual shutter must blank 11 of 12 displayed images. For a left eye shutter of a particular user three distinct cases can be considered (similarly for a right eye shutter):

1. The user's right eye image is separated by polarization.
2. The left eye images of the other 5 users are blocked by the shutter operation.
3. The right eye images of the other 5 users are blocked by polarization and shuttering.

The first case contributes only the relatively low crosstalk of standard polarization-based systems. The second case is addressed by our new double-cell shutter design, which provides fast switching times and high contrast to avoid crosstalk. In our setup the shutters in closed state must block five times more light as compared to the case of active stereo displays. Double-cell shutters help with



Figure 6: Our custom shutter glasses consist of two double cell shutters, a Zigbee radio module, a rechargeable lithium-polymer battery and the shutter driving circuit. The housing also contains multiple threaded holes for assembling different IR-reflective marker configurations.

this requirement since the total contrast ratio is the product of the contrast ratios of the NW cell and the NB cell. The third case contributes at least one order of magnitude less crosstalk than the other two cases since the light is blocked by shuttering and polarization.

We designed our wireless shutter glasses (Figure 6) such that their principal state can be controlled from the application, independent of the basic clocking. The communication to the shutter glass controller is realized by using the μ racoli implementation [URACOLI 2011] of the two lower levels of the IEEE 802.15.4 protocol stack for wireless personal area networks (which form the basis of the Zigbee protocol). There are two different aspects that can be programmed:

- The general assignment to one or more of the six-user time slots. This control can be used to implement a VIP (Very Important Person) mode by assigning two or more time slots to a single person. We often have the case that the system is used by less than six individuals and thus we use this control to increase the brightness by assigning more than one time slot to one or more users.
- A transition from shutter mode to full-open mode and vice versa. In regular operation the shutters are open for only $1/6th$ of the time and thus everything but the display is perceived as quite dark. However, if six people are in front of the display discussing various aspects of their application, it quite often happens that they look at each other or do not look at the display at all. They may even move to a whiteboard to continue discussion. In these cases we open the glasses and turn off the shutter mode using simple heuristics based on the head tracking information.

There are many other uses for application-controlled shutter glasses. Particularly in multi-display environments (e.g. [Pirchheim et al. 2009], [Kitamura et al. 2009]), where users interact with a variety of displays, shutters need to sync to the currently faced display and should be turned off if it is a 2D display or only 2D content is presented.

5 Augmented Group Navigation

Projection-based multi-viewer systems expose each user to an individual pair of images, which can be computed in a way that the following two properties are ensured:

- *Perspectively-correct perception of the virtual world by each user.* The perspective projection is defined by the relationship between the physical position of a user's eye and the physical position and size of the projection screen. Since the physical position is different for each user, the perspective projections and the computed images are also different. The correct perspective enables correct size, shape and distance perception in the virtual world. Regular stereoscopic displays present the same stereoscopic image pair to all users. Since the perspective can only be correct for at most one user, all the other users perceive distorted versions of the virtual objects [Agrawala et al. 1997].
- *Perception of a consistent virtual world among all users.* The users and the projection screen are placed in the virtual world in exactly the same spatial configuration as in the real world, apart from a global scaling factor. Only this configuration ensures that the virtual world is consistently perceived by all users (e.g. they are seeing a virtual model as if they were standing around a real version of the same model). Bare-handed pointing becomes possible and each user interprets the pointing gesture as in the real world, accompanied by some limitations with respect to accuracy [Salzmann et al. 2009].

These two properties enable a group of co-located users to perceive a virtual model as if it were real.

Navigation is a basic interaction capability of almost all virtual reality applications. However, the availability of projection-based multi-viewer systems introduces a new problem for the navigation through a virtual world, which did not exist in classical VR systems with a single head-tracked user. In such single-user systems, all the users share the same perspective and if the navigator is moving along a collision-free path, all the group members also perceive the path as being collision-free. In multi-viewer systems all of the users are simultaneously moved around along with a virtual representation of the projection screen. Since projection displays can be quite large (e.g. our display is 4.3m wide), the users are typically distributed around the space in front of the display. In this configuration, they may not fit through constrictions such as doors or aisles even if the navigating person chooses a collision-free path. The co-travelers might be passing through walls, which can be annoying and irritating and the navigator may not even be aware of these problems.

Unfortunately, there is no general solution to this problem, which maintains a correct perspective *and* the consistency of the shared virtual world for all users without requiring that all users are looking from the same position as the navigator, which is physically impossible. However, we suggest various approaches to mitigate the problem:

- *Stop and crowd.* As a simple solution we perform collision detection for each user and stop the navigation if at least one user collides. To continue, users need to resolve the collision (e.g. by moving closer to the navigating user).
- *Distort.* On a path towards an obstacle, the head position of the user is moved towards the head position of the navigator until the collision is avoided (Figure 7). Changing the head position incurs a distortion of the perspective and thus the surrounding space is no longer correctly perceived. In an extreme case, the system collapses to a two-view system and all

the users see the same images, which still results in a different perception of the size, shape and position of virtual objects.

- *Detour.* On a path towards an obstacle the system interferes and moves the user along a collision-free path while maintaining a perspectively correct rendering (Figure 8). This mode has three degrees of freedom if a planar movement is assumed. The user along with the screen representation can be translated in two dimensions as well as rotated around the up-axis. We use a heuristic to determine the translation, which moves the user towards the navigator as much as necessary to avoid a collision. The rotation is ignored up to this point. This approach temporarily gives up the consistency of the shared virtual world, while the perspectively correct perception is maintained.
- *Fade.* If a user is on a path towards an obstacle, this object is slowly faded out during approach. The opacity of the obstacle is defined as a function of the distance to the respective viewpoint and the current velocity of the user. In our current implementation object geometry less than four seconds away starts fading out such that it is invisible one second before the linearly predicted collision. Thus the irritating collision with the obstacle is avoided while users may still examine details of an object when standing in close proximity. However, during travel the consistency of the shared model is affected since other users may see an unaltered object.

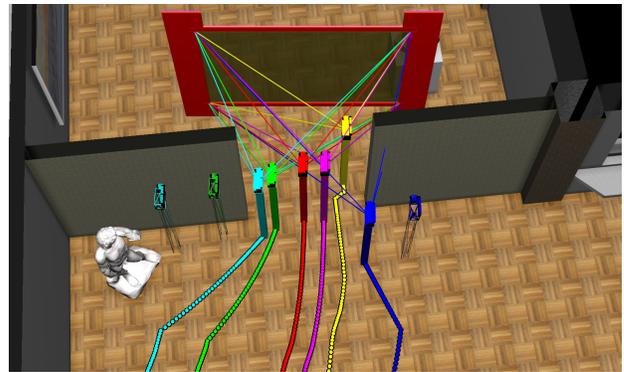


Figure 7: Distort: The viewpoint position (opaque representations) is shifted towards the open passage if the physical head position of a user (drawn as wireframe) would collide on its original course. Users are represented as colored icons. Their viewing frustum is defined by their position and the virtual screen representation. Shifting the head position away from the correct position leads to a distorted perception of the virtual world.

The last three approaches return to an artifact-free rendering if there are no more constrictions. They use smooth transitions in and out of these modes to avoid disorientation. *Stop and crowd* and *Fade* depend only on each user's individual viewpoint position and the motion velocity. *Detour* and the *Distort* redirect a user from a given navigation path. In these cases we assume that the navigating user chooses or is forced along a collision-free path. For all other users we use ray-object intersection tests to find obstacles in the steering direction. In case that intersections closer than a velocity-dependent distance threshold are found an obstacle-avoiding movement toward the path of the navigating user is added to the prevailing viewpoint motion. The distance threshold is a function of the steering velocity. The velocity of the obstacle-avoiding movement is defined as a function of the distance to the main navigation path and the current steering velocity ensuring that it is zero once the predicted collision would occur. In our current implementation we

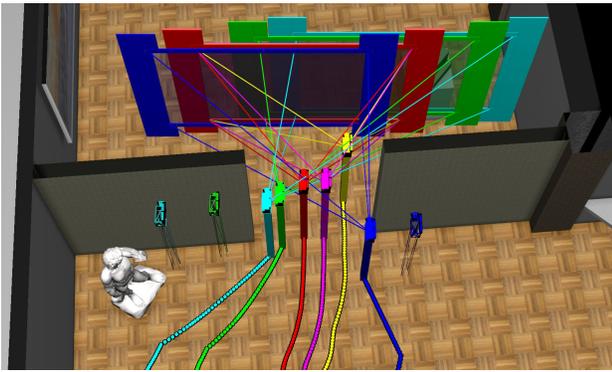


Figure 8: *Detour:* The viewing frustum associated with a particular user is shifted towards the open passage if the physical head position of the user would lead to a collision. Users no longer observe the virtual world through the same window.

compute the distance threshold such it considers objects up to five seconds ahead (e.g. 3.5m at a velocity of 0.7m/s).

As a consequence of these viewpoint adaptations we must deal with two representations of each user's position: one is corresponding to the original head position defined by the tracking system and the other corresponds to the adjusted head point position for collision avoidance (Figures 7 and 8). Once looking ahead from the adjusted head position does not predict any more collisions, we do not move the user further toward the path of the navigator. As a result the collision-avoiding motion tries to minimize the deviation from the original path as can be seen in Figure 8. The adjusted head position can be moved back to the physical head position if the ray-object intersection tests for the original head position do not predict any further collisions and if there is no obstacle between the original head position and the adjusted head position.

This heuristic approach for collision avoidance has obviously many limitations including the shape and orientation of obstacles that can be robustly detected to avoid collisions. However, in practice it works well in typical use cases such as an architectural walk through or a tour through a virtual museum. In particular, it was sufficient for our purpose of evaluating the general applicability and user acceptance of these different collision handling techniques for group navigation.

The *Detour* and *Distort* techniques can be compared to the concept of redirected walking [Razzaque et al. 2002; Nitzsche et al. 2004; Steinicke et al. 2008], since they redirect a user from a given navigation path. Redirected walking aims to extend the range of physical walking motions in virtual environments by redirecting the user from an actively controlled walking path. In our system a group of users is being navigated by an operator as if being passengers in a vehicle that is driven by somebody else. The path of the driver is unaffected by the redirection techniques. Thus the passive passengers do not have a particular expectation of a movement through the virtual environment or a proprioceptive reference that would allow them to judge the extent of redirection. The only reference are the positions of the other users in front of the display and the size of the constriction in the virtual world. This makes it very difficult to even detect the redirection.

The described augmented group navigation techniques are all orthogonal to the actual navigation technique being used by the navigating user. In our demonstrations and user studies we use predefined navigation paths or direct steering techniques. As a result of unsatisfactory experiments with handheld steering devices we

developed a stationary group navigation device, the Spheron (Figure 9). The most prominent feature of the Spheron is a large 3D trackball for rotating the view. An elastic handle is mounted at the base of the device for controlling the movement through the scene. The Spheron is centrally placed within the shared action space and is easily accessible by all users so that taking turns while navigating is easily achieved. The physical presence of the device itself and the required manual operation foster mutual awareness among the group of the interactions performed by the navigator and allow others to quickly take over control and interfere if desired.

The Spheron is inspired by the CAT [Hachet et al. 2003], which was also developed for group interaction in large-screen environments. Compared to the CAT the Spheron has a stronger affordance for rotation through the 3D trackball. Both, examining objects from different sides and looking around in the environment are particularly necessary in the context of group navigation using a single wall-sized display. As another important difference to the CAT the trackball can be equally well accessed from all directions whereas the steering wheel of the CAT can be awkward to grasp if it is tilted. The 3D trackball of the Spheron always remains at its position. It can even be accessed blindly and thus facilitates frequent changes of the operator among a group of users.



Figure 9: *The Spheron.* Moving the elastically suspended handle induces a corresponding rate-controlled motion through the scene. Rotation of the 3D trackball results in changing the view direction. Alternatively, for viewing a single model, the trackball rotations can be mapped to rotations of the model.

6 Results and Evaluation

We set up the six-user system in our lab with a rear-projection onto a 4.3m by 2.7m screen (Figure 1). The system is driven by two computers in a master-client architecture. All application logic is computed on the master node and the state updates are distributed over the network to the clients. The client applications render the scene for the individual users based on head tracking information directly fed to the clients. Using this setup we typically achieve application frame rates above 30Hz for the scenarios in our user study. The client applications run on a single HP Z800 computer

equipped with two Intel Xeon X5680 six-core processors running at 3.33GHz, 96GiB of main memory and three NVIDIA Quadro Plex 7000 graphics subsystems. The system configuration of the rendering computer can be seen in Figure 3. Our demonstrators are based on the free software AVANGO^{NG} [AvangoNG 2011; Kuck et al. 2008] under 64 bit Ubuntu 9.04. The end-to-end latency from tracking the user's motion input to the display update is about 80-120ms depending on the actual rendering frame rate. This is very similar to a single-user system, since the actual rendering threads run in parallel on separate cores and GPUs.

We use a large-field tracking system from Personal Space Technologies [Pst 2011] consisting of six units for tracking the head positions of the users and further input devices. Each unit is an independent 6DOF tracking system that is synchronized to the other tracking devices in the compound. The resulting tracking range is defined by the union of the viewing frusta of the six tracking units, which require only small overlaps for the handover between adjacent frusta. For the image separation we currently use different prototypes of shutter glasses. Three of them consist of FLC shutters, the other three are based on the novel LC double-cell design.

6.1 Technical Evaluation

We used a Spectroradiometer (Konica-Minolta CS-1000A) to measure the relative impact of the display components on the overall optical quality. We were particularly interested in the perceived brightness at the user's eye behind the shutter glasses and the potential crosstalk from images of the other users. We also wanted to verify that adding polarization filters in front of the DLP projectors only slightly decreases the brightness level per user. Our measurements show that it is in fact only decreased by 12% since the shutter glasses are polarizing the light anyway. We found that the brightness linearly increases with the length of the opening periods of the shutters. This behavior verifies that each user receives 1/6th of the total brightness of the six projectors, which is equivalent to the brightness of a single projector. These tests were done with FLC shutters and double-cell shutters with precisely adjusted timings, which resulted in similar behavior.

As explained earlier, ghosting images are much more an issue in multi-user projection systems than in common single-user stereo displays. In particular, ghosting of images from other users is more noticeable than left-right eye ghosting since the other users move independently of oneself. Thus the ghosts move within ones images even while standing still. To avoid ghosting we developed the double-cell shutter glasses, which provide fast switching times and a high contrast ratio. An informal comparison of the novel double-cell LC-shutter glasses to the FLC-shutter glasses indicates that double-cell shutters eliminate the ghosting to a non-perceptible level in most scenarios while FLC-shutters show slight but perceptible leakage in dark image areas, which is not visible with the double-cell shutters.

6.2 User Feedback

In general we observe that people are more enthusiastic about exploring a virtual environment as part of a group with perspective correct views for each user than in a regular stereoscopic environment. Even if the displayed content does not directly correspond to their interest, it becomes a relevant part of their shared reality through the immediate exchange with others. The Spheron turned out to be an easy to use navigation device, which does not require any explanations in most cases. For evaluating our augmented group navigation techniques we performed a pilot study followed by a formal users study.

6.2.1 Pilot Study

For the first study we invited two groups of five users. One group consisted of students of industrial design, the other group of students and alumni from the department of architecture. We introduced the five participants of each group to the technology and advised them to notify us of any problems such as excessive crosstalk, loss of orientation or nausea. An instructor also wearing shutter glasses steered the group through the model of a museum with several rooms and exhibits using the Spheron device. The different group navigation techniques were presented in a predefined sequence starting with the *Stop and crowd* mode.

For the evaluation we chose a semi-structured group interview, allowing us to introduce new questions during the discussion. We were interested in two main topics: the overall usability of the system and, more specifically, the users' experience with the different augmented group navigation techniques. The different techniques were rated on five-point Likert scale. Besides an overall very positive assessment of the system the main observations of the pilot study are:

- Ghosting was not explicitly reported as a problem. Nevertheless, the participants remarked that they generally preferred the double-cell shutter glasses over the FLC-glasses for reasons of image quality.
- Everyone agreed that bumping through opaque walls is not acceptable and augmented group navigation techniques are needed.
- The *Distort* technique was consistently rated low for two major reasons: the distorted geometry breaks the realistic impression of space, and the animated transitions in and out of the *Distort* mode can contribute to a feeling of dizziness.
- The *Stop and crowd* solution provides the most realistic experience. However, it quickly becomes annoying if the group needs to pass through a number of doors and aisles.
- The opinions on the *Fade* mode varied. Some liked it for being a convenient option to maintain one's relative position in the group, while others found that it would deteriorate the appearance of the model.
- The architects preferred *Detour* as it provided the most realistic experience of the virtual building without the necessity of crowding.

Statistical analysis of the user ratings for the different techniques using the Kruskal-Wallis test for ordinal data and the Mann-Whitney U test for post-hoc comparisons revealed a significant preference for the *Detour* technique over *Distort* ($p < 0.05$). The other reported preferences did not become significant, but there are clear trends visible in the data. The strong preference for *Detour* as compared to *Distort* demonstrates the users' aversion to perspective distortions, which was the main reason to develop the multi-user projection system.

6.3 User Study

Overall the pilot study revealed that *Detour* and *Fade* were considered the most useful augmented group navigation techniques. We compared these two techniques in a controlled user study to evaluate their usability in relation to the baseline of untreated collisions during group navigation.

6.3.1 Experimental Setup

We modeled a simple virtual architecture containing full-size 3D models of cars and photos of their real counterparts (Figure 10). Navigation was automated and took the participants with a fixed velocity of 0.7m/s on a predefined path through the environment. The automated tour eliminated any bias of an operator steering in a slightly different way in each test condition. It also allowed testing with groups of six users. During the tests the users could move freely in front of the screen to assume desired viewpoints.

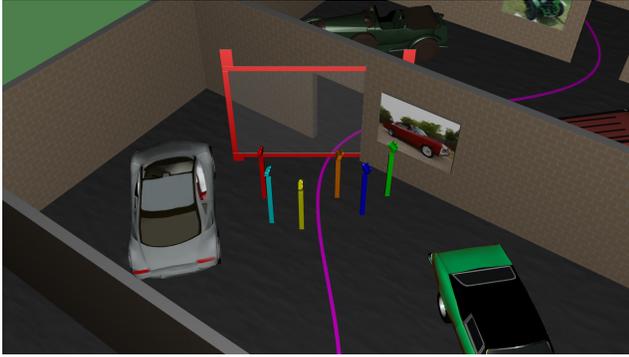


Figure 10: Study task: The group is automatically moved along a given path through a virtual car exhibition. Participants were asked to detect a set of features on the 3D car models or in the exhibited pictures.

The participants were asked to perform a search task in the virtual environment. Each user received a sheet of paper showing 15 pictures with details from virtual car models and their real counterparts. We informed them that exactly seven of these details could be found in the presented scene. The users' task was to memorize the pictures and search for the respective details during the automated tour with one of both group navigation techniques applied. Thereafter, recognized details should be marked on the paper. During a repetition of the same tour everybody got the chance to verify the findings, while another group navigation technique was applied. This primary task ensured that they would concentrate on the application content rather than focusing for the perceptual artifacts of the tested group navigation technique.

Besides the factor *navigation technique* we included *information exchange* as another independent variable. We were interested to see if solving the memory task individually in the *single* condition or exchanging information within the *group* would affect the discovery. While we did not permit gesturing or talking in the *single* condition, mutual information exchange during the tour was encouraged in the *group* condition. In both conditions everybody noted the recognized detail individually again. Our hypothesis was that individual users would perform better if they are able to exchange information about recognized details in the virtual environment. We used two different sets of models and pictures with comparable difficulty in both conditions.

6.3.2 Participants

Seven female and 17 male users aged between 20 and 30 years participated in this study. None of them reported problems with stereoscopic vision. Six of them claimed to have extensive experience with interactive 3D computer graphics while three reported to have no prior experience with the technology. All 24 participants were university students.

6.3.3 Design and Procedure

First, we introduced the participants of each group to the general characteristics of our multi-user VR system. The interior of a car was shown and the participants were asked to touch items like the rear view mirror or the steering wheel. As everybody could see what the others were doing, it became clear that the displayed virtual environment was consistent for all involved users. We also introduced the mentioned collision problem for multi-user navigation. The automated navigation was started without any collision handling. We made sure that the users were changing their position in front of the screen such that everybody experienced the effect of bumping through virtual walls. Directly after this experience we asked the participants to rate the severity of the collision issue on a five-point Likert scale (ranging from -2 (unacceptable) to $+2$ (very good)). We also asked whether anybody felt symptoms of cybersickness during this training tour, offering to abort the study if necessary.

Both augmented navigation techniques were presented two times to each participant, once for each condition of the factor *information exchange*. The order of all conditions was equally distributed among the four test groups using a reduced Latin square design.

6.3.4 Results

Users strongly complained about bumping through walls and consistently rated this baseline condition negatively with a mean of -0.5 ($sd = 0.71$) while both augmented group navigation techniques were rated positively with a mean of $+0.5$ for *Detour* ($sd = 1.1$) and $+0.7$ for *Fade* ($sd = 0.71$) (Figure 11). Statistical analysis using the Kruskal Wallis test and the Mann-Whitney U test for post-hoc comparisons revealed this difference to be highly significant ($p < .01$). There was no statistically significant difference between the ratings for *Detour* and *Fade*. We observe, however, that 22 of 24 users have a clear preference for one of both techniques. The 10 participants voting for *Detour* explain their choice with a more realistic appearance of the architecture and many among them claim that they did not even realize the redirection. *Fade* instead is preferred by 12 participants who identified the consistency of the virtual view to their physical position in the group as the most relevant feature of our multi-user VR system.

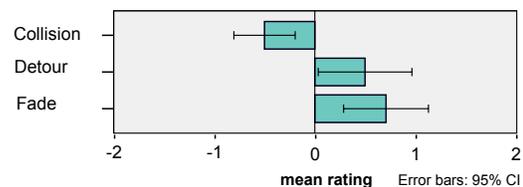


Figure 11: Mean ratings and confidence intervals for the three tested group navigation techniques.

We also find evidence that users can indeed benefit from immediate communication within our virtual environment. The success ratio for recognizing details in the virtual environment was about nearly twice as high in the *Group* condition ($mean = 0.43$, $sd = 0.16$) as compared to the *Single* case ($mean = 0.25$, $sd = 0.22$). This statistically significant result ($T_{(46)} = 3.437$, $p < 0.05$) was expected. It indicates that our multi-user VR system indeed enables users to exploit common advantages of collaboration.

7 Conclusions and Future Work

We designed and implemented the first large screen stereoscopic multi-viewer display for six tracked users, which provides precise horizontal and vertical parallax. The system runs at 360Hz, which results in 60Hz per user. Left and right eye separation is provided by polarization. This efficient combination of active and passive stereo achieves almost the same brightness per user as a single-user active stereo system based on the same projectors would. Our intelligent shutter glasses enable application-level control of the shutter timing for better usability and higher brightness if less than six users use the system. Novel augmented group navigation techniques facilitate the exploration of environments with constricted navigation paths by avoiding collisions of fellow travelers.

Our augmented group navigation can be further refined in many directions. In particular, formulating the combination of different collision-avoiding strategies as an optimization problem could result in a good heuristic to select the best strategy for each situation and provide smooth transitions in between. So far we have focused on keeping the shared space intact as much as possible while navigating. While this seems highly desirable in most cases, there might also be applications where it is beneficial to temporarily relax this requirement. Each user might stroll through a museum on an individual path, but has the possibility to rejoin the group if desired. Tele-collaboration systems such as the blue-c [Gross et al. 2003] have only been developed for a single tracked user. Our approach would enable immersive group-to-group collaboration. However, navigating a group consisting of local and remote users in a consistent and transparent way remains a considerable challenge.

Our six projector array is currently a unique prototype system. The integration of three separate projectors into a three-chip DLP projector is the next step forward. This would reduce the complexity of setting up and running such a system. In addition, it is much more energy efficient and produces much less heat considering a three-chip DLP projector generates the different primary colors by splitting up the white light beam from the projector lamp instead of using color filters. Our approach of turning a regular monoscopic DLP projector into a stereoscopic projector by switching between two input and formatter boards could be applied twice, and thus we could switch between four input streams. While this may come with the disadvantage of losing color depth and reduced brightness, it would allow us to build an active stereo two-user single-chip color wheel projector, an active stereo six-user three-chip DLP projector or a twelve-user active-passive system consisting of two such three-chip projectors. With only minor modifications, all of these systems can be built from readily available components.

Multi-viewer projectors would significantly improve the usability of various immersive system designs, which typically involve groups of three to six users. Multiple users could gather around a table-top display like the Responsive Workbench [Krüger et al. 1995]. Since the invention of the CAVE [Cruz-Neira et al. 1992], the original design has been considerably improved by using more walls and higher resolution (e.g. the C6 [VRAC 2007]). We are looking forward to seeing our technology serve as an instrumental part in the evolution towards the first C6x6.

Acknowledgments

This work was supported in part by the European Union under grant 217140 (project IMVIS), the German Federal Ministry of Education and Research (BMBF) under grant 03IP704 (project Intelligentes Lernen) and the Thuringian Ministry of Education and Cultural Affairs (TKM) under grant B514-08028 (project Visual Analytics in Engineering). We thank our IMVIS partners LC-TEC

DISPLAYS AB, Personal Space Technologies and Centro Ricerche FIAT for the excellent collaboration throughout the project, Frank Pudlowsky and Jörg Krall from NVIDIA's Professional Solutions Group as well as Joachim Girke from Hewlett-Packard GmbH for providing extraordinary support with our complex hardware setup, and the members and students of the Virtual Reality Systems group at Bauhaus-Universität Weimar for their help with getting everything to work. The model of the David statue in Figure 1 is courtesy of Stanford's Digital Michelangelo Project.

References

- AGRAWALA, M., BEERS, A. C., MCDOWALL, I., FRÖHLICH, B., BOLAS, M., AND HANRAHAN, P. 1997. The two-user responsive workbench: support for collaboration through individual views of a shared space. In *Proceedings of ACM SIGGRAPH '97*, 327–332.
- ARGELAGUET, F., KUNERT, A., KULIK, A., AND FROEHLICH, B. 2010. Improving co-located collaboration with show-through techniques. In *3D User Interfaces (3DUI), 2010 IEEE Symposium on*, 55–62.
- ARTHUR, K., PRESTON, T., TAYLOR, R., BROOKS, F., WHITTON, M., AND WRIGHT, W. 1998. Designing and building the pit: a head-tracked stereo workspace for two users. In *Proceedings of IPT '98, Workshop CD-ROM*.
- ASPIN, R., AND ROBERTS, D. 2005. A model for distributed, co-located interaction in urban design/review visualisation. In *Proceedings of IEEE DS-RT 2005*, 70–77.
- AVANGONG, 2011. Cross-platform software framework for virtual reality, <http://www.avango.org/>.
- BARCO, 1999. Virtual surgery table, <http://www.barco.com>.
- BENFORD, S., GREENHALGH, C., RODDEN, T., AND PYCOCK, J. 2001. Collaborative virtual environments. *Commun. ACM* 44 (July), 79–85.
- BIMBER, O., FROEHLICH, B., SCHMALSTIEG, D., AND ENCARNAÇÃO, L. M. 2001. The virtual showcase. *IEEE Computer Graphics and Applications* 21, 48–55.
- BOLAS, M., MCDOWALL, I., AND CORR, D. 2004. New research and explorations into multiuser immersive display systems. *IEEE Computer Graphics and Applications* 24, 18–21.
- BOWMAN, D., KRUIFF, E., VIOLA, J. L., AND POUPLYREV, I. 2005. *3D User Interfaces: Theory and Practice*. Addison-Wesley.
- COSSAIRT, O. S., NAPOLI, J., HILL, S. L., DORVAL, R. K., AND FAVALORA, G. E. 2007. Occlusion-capable multiview volumetric three-dimensional display. *Appl. Opt.* 46, 8 (Mar), 1244–1250.
- CRUZ-NEIRA, C., SANDIN, D. J., DEFANTI, T. A., KENYON, R. V., AND HART, J. C. 1992. The cave: audio visual experience automatic virtual environment. *Commun. ACM* 35 (June), 64–72.
- DAMERA-VENKATA, N., AND CHANG, N. L. 2007. Realizing super-resolution with superimposed projection. In *CVPR: Computer Vision and Pattern Recognition*.
- D'ANGELO, D., WESCHE, G., FOURSAs, M., AND BOGEN, M. 2008. The benefits of co-located collaboration and immersion on assembly modeling in virtual environments. In *Advances in*

- Visual Computing*, vol. 5358 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 478–487.
- DODGSON, N. A. 2005. Autostereoscopic 3d displays. *Computer* 38, 8, 31–36.
- FAVALORA, G. E. 2005. Volumetric 3d displays and application infrastructure. *Computer* 38 (August), 37–44.
- FRÖHLICH, B., BLACH, R., STEFANI, O., HOCHSTRATE, J., HOFFMANN, J., KLÜGER, K., AND BUES, M. 2005. Implementing multi-viewer stereo displays. In *WSCG (Full Papers)*, 139–146.
- GROSS, M., WÜRMLIN, S., NAEF, M., LAMBORAY, E., SPAGNO, C., KUNZ, A., KOLLER-MEIER, E., SVOBODA, T., VAN GOOL, L., LANG, S., STREHLKE, K., MOERE, A. V., AND STAADT, O. 2003. blue-c: a spatially immersive display and 3d video portal for telepresence. In *ACM SIGGRAPH 2003 Papers*, ACM, 819–827.
- HACHET, M., GUITTON, P., AND REUTER, P. 2003. The cat for efficient 2d and 3d interaction as an alternative to mouse adaptations. In *Proceedings of the ACM VRST'03*, ACM, 205–212.
- HOPP, A. 2005. Using a single spatial light modulator for stereoscopic images of high color quality and resolution. In *IPT/EGVE*, Eurographics Association, E. Kjems and R. Blach, Eds., 15–22.
- HUA, H., BROWN, L., AND GAO, C. 2004. Scape: supporting stereoscopic collaboration in augmented and projective environments. *Computer Graphics and Applications, IEEE* 24, 1 (Jan-Feb), 66–75.
- JONES, A., MCDOWALL, I., YAMADA, H., BOLAS, M., AND DEBEVEC, P. 2007. Rendering for an interactive 360 degree light field display. *ACM Trans. Graph.* 26 (July).
- JONES, A., LANG, M., FYFFE, G., YU, X., BUSCH, J., MCDOWALL, I., BOLAS, M., AND DEBEVEC, P. 2009. Achieving eye contact in a one-to-many 3d video teleconferencing system. *ACM Trans. Graph.* 28 (July), 64:1–64:8.
- JONES, A., LANG, M., FYFFE, G., YU, X., BUSCH, J., MCDOWALL, I., BOLAS, M., AND DEBEVEC, P. 2009. Headspin: a one-to-many 3d video teleconferencing system. In *ACM SIGGRAPH 2009 Emerging Technologies*, ACM, New York, NY, USA, SIGGRAPH '09, 13:1–13:1.
- KITAMURA, Y., KONISHI, T., YAMAMOTO, S., AND KISHINO, F. 2001. Interactive stereoscopic display for three or more users. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 231–240.
- KITAMURA, Y., SAKURAI, S., YAMAGUCHI, T., FUKAZAWA, R., ITOH, Y., AND KISHINO, F. 2009. Multi-modal interface in multi-display environment for multi-users. In *Proceedings of the 13th International Conference on Human-Computer Interaction. Part II: Novel Interaction Methods and Techniques*, Springer-Verlag, Berlin, Heidelberg, 66–74.
- KRÜGER, W., BOHN, C.-A., FRÖHLICH, B., SCHÜTH, H., STRAUSS, W., AND WESCHE, G. 1995. The responsive workbench: a virtual work environment. *Computer* 28, 7, 42–48.
- KUCK, R., WIND, J., RIEGE, K., AND BOGEN, M. 2008. Improving the avango vr/ar framework: Lessons learned. In *5th Workshop of the GI-VR/AR Group*, Shaker, 209–220.
- KUNZ, A. M., AND SPAGNO, C. P. 2002. Technical system for collaborative work. In *Proceedings of the Eurographics Workshop on Virtual Environments 2002*, EGVE '02, 73–80.
- MAKSAKOV, E., BOOTH, K. S., AND HAWKEY, K. 2010. Whale tank virtual reality. In *Proceedings of Graphics Interface 2010*, Canadian Information Processing Society, GI '10, 185–192.
- MULDER, J. D., AND BOSCHKER, B. R. 2004. A modular system for collaborative desktop vr/ar with a shared workspace. *Virtual Reality Conference, IEEE*, 75–83.
- NITZSCHE, N., HANEBECK, U. D., AND SCHMIDT, G. 2004. Motion compression for telepresence walking in large target environments. *Presence: Teleoper. Virtual Environ.* 13 (February), 44–60.
- OTTO, O., ROBERTS, D., AND WOLFF, R. 2006. A review on effective closely-coupled collaboration using immersive cve's. In *Proceedings of the ACM VRCIA '06*, ACM, 145–154.
- PIRCHHEIM, C., WALDNER, M., AND SCHMALSTIEG, D. 2009. Desktotheque: Improved spatial awareness in multi-display environments. In *Virtual Reality Conference, 2009. VR 2009. IEEE*, 123–126.
- POSTON, T., AND SERRA, L. 1994. The virtual workbench: dextrous vr. In *Proceedings of the conference on Virtual reality software and technology*, World Scientific Publishing Co., Inc., River Edge, NJ, USA, 111–121.
- PROJECTIONDESIGN, 2011. <http://www.projectiondesign.com>.
- PST, 2011. Personal space tracker, <http://www.ps-tech.com>.
- RAZZAQUE, S., SWAPP, D., SLATER, M., WHITTON, M. C., AND STEED, A. 2002. Redirected walking in place. In *EGVE '02: Proceedings of the Eurographics Workshop on Virtual Environments 2002*, Eurographics Association, 123–130.
- RIEGE, K., HOLTAMPER, T., WESCHE, G., AND FROEHLICH, B. 2006. The bent pick ray: An extended pointing technique for multi-user interaction. In *IEEE Symposium on 3D User Interfaces.*, 62–65.
- SALZMANN, H., MOEHRING, M., AND FROEHLICH, B. 2009. Virtual vs. real-world pointing in two-user scenarios. In *Proceedings of the 2009 IEEE Virtual Reality Conference*, IEEE Computer Society, Washington, DC, USA, 127–130.
- SCHMALSTIEG, D., FUHRMANN, A., HESINA, G., SZALAVÁRI, Z., ENCARNACÃO, L. M., GERVAUTZ, M., AND PURGATHOFER, W. 2002. The studierstube augmented reality project. *Presence: Teleoper. Virtual Environ.* 11 (February), 33–54.
- STEINICKE, F., BRUDER, G., JERALD, J., FRENZ, H., AND LAPPE, M. 2008. Analyses of human sensitivity to redirected walking. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*, VRST '08, 149–156.
- URACOLI, 2011. *uracoli - The μ Controller Radio Communication Library*, <http://www.nongnu.org/uracoli/>.
- VRAC, 2007. Virtual reality application center, Iowa State University, <http://www.vrac.iastate.edu/c6.php>.
- WOLFF, R., ROBERTS, D. J., STEED, A., AND OTTO, O. 2007. A review of telecollaboration technologies with respect to closely coupled collaboration. *Int. J. Comput. Appl. Technol.* 29 (June), 11–26.

