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

Optical Magic Lenses and Polarization-Based Interaction Techniques

Manuela Waldner, Michael Kalkusch, Dieter Schmalstieg

Institut für Maschinelles Sehen & Darstellen, TU Graz, Austria

Abstract

In this paper we present a novel approach of the magic lens user interface metaphor for large-scale projector-based displays. We altered a standard polarization-based passive stereo projection setup and employed a standard LCD panel as a purely optical, tangible magic lens device. Due to the properties of polarized light, the modified passive stereo setup can be used to separate two views – a primary and a secondary layer – of the projected data. A non-powered LCD panel serves as magic lens filter, as it rotates the direction of polarized light 90 degrees, providing the user a different view on the projected data. The system is arbitrarily scalable for multiple users and can be applied to numerous applications. Based on the two projection layers resulting from the proposed setup we explored interaction techniques and present some examples of the system.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Methodology and Techniques]: Interaction techniques I.3.7 [Three-Dimensional Graphics and Realism]: Virtual reality

1. Introduction

Large-scale display walls built with commodity projectors are becoming increasingly popular for applications involving the visual analysis of complex data sets. The large display surface allows showing a large amount of data, and also provides sufficient space to arrange data items according to the user's preferences. Collaboration is naturally supported by allowing enough physical space in front of the display to accommodate multiple users.

Large screen displays afforded by projector based systems allow for efficiently applying a large variety of visualization and interaction techniques. However, a number of problems remain which can only be partially addressed with established techniques:

Direct spatial interaction with the display requires that the user employs a spatial tracking device, typically an expensive 3DOF/6DOF tracking device if interaction should not be strictly limited to 2D. Simultaneous interaction of multiple users requires multiple tracking devices as well as a software design that can handle multiple input devices. Such a solution increases system complexity and does not scale well with a larger group of users, such as encountered in a public space setup. Moreover, simultaneous interaction of multiple users in the same display space can lead to unwanted situations where users compete for a particular region of the display or get into each other's way.

Visualization of complex datasets is often addressed with multiple view systems and focus+context techniques combined with linking and brushing inside datasets. While a wall-sized display (or multiple side-by-side displays) may offer ample space for arranging such views, the multiple view as well as the focus+context approach leave it to the user to mentally merge corresponding elements from multiple spatially separated views. Moreover, navigating in multiple separate views operating at multiple scales requires additional mental effort because correspondences may not always be obvious.

Magic lenses [BSP*93] can be described as an in-place display for information filtering, similar in its effect to using multiple views. A magic lens is a user controlled widget which provides a filtered view of the content underneath. Magic lenses make direct use of spatial context and consequently are a very powerful visualization technique. However, all existing implementations have some of the following disadvantages for collaborative applications:

- They do not allow multiple users to apply per-user filtering or subjective views [Smi96]. Users can still get into each others way when manipulating the data.

- They require a per-user electronic input device, adversely affecting scalability and system complexity.
- The physical nature of the display, in particular the display resolution, cannot be changed by the magic lens. Physical focus+context displays try to overcome this problem by superimposing multiple displays inside the lens. Magnification, rendering style or even the sematic of the underlying data is changed inside the lens. However, these approaches either use an inflexible static display arrangement [AR05], [BGS01] or require a complex mechanic setup [YGH*01] [SAKH06].

Our approach describes a set of interaction techniques relying on visual channel separation using polarization effects. Polarization filters are well-known for their use in passive stereoscopic viewing systems, where they are used to separate left and right stereo channels. We retarget the channel separation effect achieved in a two-projector setup to separate *content* channels – namely, the standard or primary view and the secondary view – through the lens. A simple 90 degree polarization retarder (an optical device altering the polarization direction of light, such as a non-powered liquid crystal panel) allows to select among the content channels by looking at a point on the screen either through the retarder or not. The retarder can be cast in several forms – for example, a handheld *optical magic lens* prop, a view-finder mounted to a console or a camera-mounted filter for computer vision-based interface techniques. Figure 1 shows a system sketch.

The unique advantage of our approach is that no electronic input device or tracking system is required to perform basic operations – physical placement of the retarder prop by the human user is simple, error free and scalable. To illustrate our ideas, we report on several experiments that have been conducted with the polarization-based interaction techniques, including a visual data mining application with two-user optical magic lenses and an augmented-reality style annotation application.

2. Related Work

2.1. Magic Lenses and focus+context techniques

Magic lenses were first introduced by Bier et al. [BSP*93]. They represent movable visual filters that change the representation of the underlying data. Magic lenses can be used as virtual magnifying glasses, to reveal hidden information, to filter the data representation, or to modify the visualization mode on the given position. This concept was later extended to the third dimension by Viega et al. [VCWP96]. The idea of a magic lens remains a popular choice for organising complex information displays, including collaborative applications [FMM*05] and very large displays [BB05].

In information visualization, focus+context is an established technique to allow a large overview integrated with details. A moveable focus area is surrounded by a context area, which is often compressed or distorted [RM93] [Fur86]. Focus+context and magic lens techniques are related, but focus+context seems to have a more general meaning in terms of spatial arrangement. While focus+context approaches offer continous transitions between regions, magic lenses have a defined boundary.

2.2. Physical magic lenses and focus+context displays

A significant amount of investigation is concerned with the embodiment of magic lens techniques through physical objects and tangible interfaces. A first category of visualization systems concern the provision of high-resolution focus insets in projector-based displays, which are not directly manipulable by the user. The focus+context screen [BGS01] uses a carefully aligned flat screen for this purpose. The Escritoire [AR05] is a table-top display composed of two superimposed projections generating a high resolution projected image within a lower resolution one. Unlike these completely static insets, PixelFlex [YGH*01] and later [SAKH06] use a pan-tilt projector to achieve a moveable high resolution area within a lower-resolution projected display.

A second category aims at providing handheld magic lens props, literally interpreting the “lens” as a tangible device. One approach is to use a passive, tracked prop together with a single projected display. The display depicts the regular view with a magic lens rendered on top of it. Representatives of this category are props made from transparent material used with a stereo-backprojection [SES99] and similar retroreflective props used with a head mounted projector display [BHG03].

A third category of magic lens displays are active magic lenses based on handheld displays. The pioneering work in this category is the Chameleon [FZC93] which relied on a simple tethered, tracked LCD screen. Recent work has used self-contained graphics and fiducial tracking using camera-equipped handheld computers [WS03] to achieve augmented reality effects. This technology was used later for maps [SKM06]. Similarly, the Peephole [Yee03] uses an externally tracked handheld computer for two handed interaction with documents anchored in 2D space. The Interaction Lens described in [MPL*02] combines a digitizer tablet with a handheld computer for integrating the work on paper-based and virtual documents. A final variety uses handheld tracked displays together with larger projected displays – the metaDESK [UI97] has a mechanically suspended screen positioned above a tabletop display, while the Ubiquitous Graphics system [SH06] combines a tracked Tablet PC with a projected display.

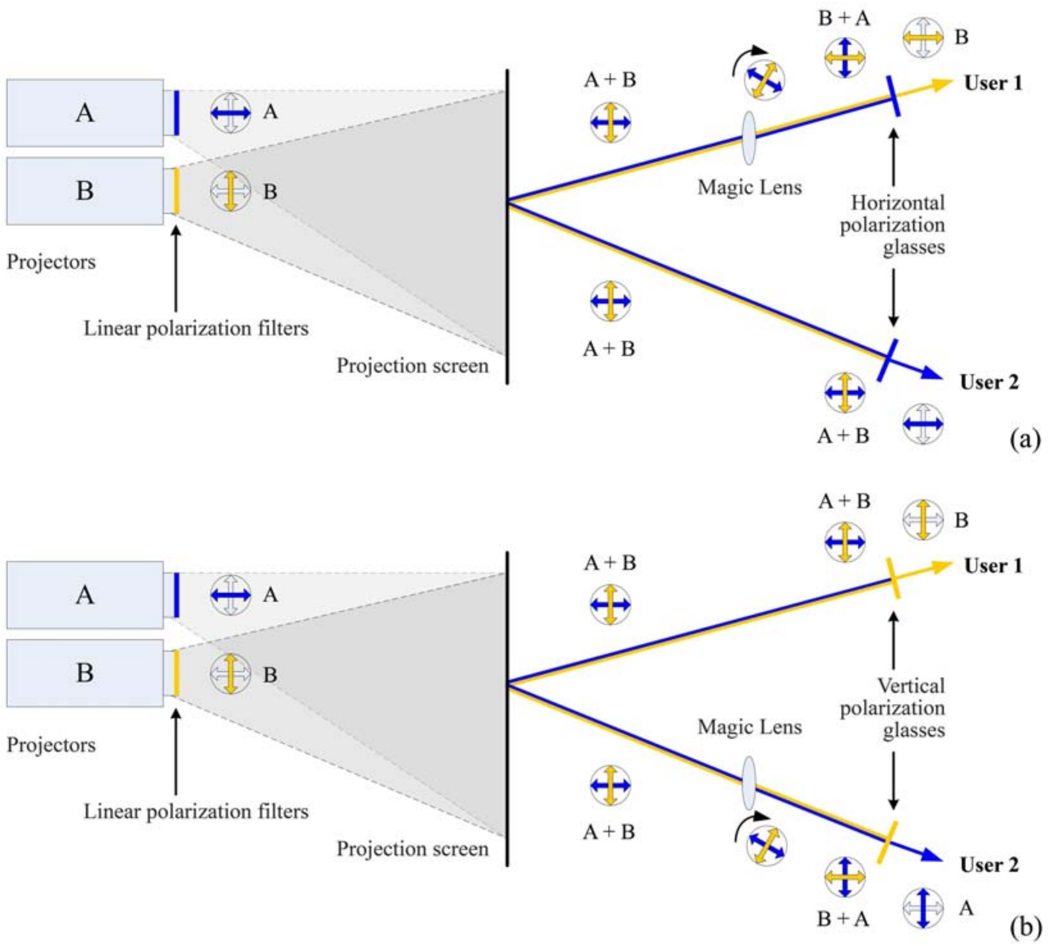

Figure 1: *The optical magic lens in a rear projection setup: The two projectors are equipped with linear polarization filters with 90 degrees offset (projector A horizontal, projector B vertical). Users looking at the screen using horizontal polarization glasses only see the image of projector A (a), whereas using vertical polarization glasses, they only see the image of projector B (b). By combining the polarization filter with the magic lens, which turns the polarized light, the image of the occluded projector becomes visible and the image of the other projector disappears.*

2.3. Subjective views and multi-user displays

Subjective views were introduced by Smith for collaborative virtual environments [Smi96]. The basic idea is to relax consistency in a shared dataset to allow per-user customized views. Snowdon and Jää-Aro [SJA97] also employ subjective views, in this case to show different representations of datasets. Butz et al. [BBF98] describe user-interface widgets to make users more aware of the state of their subjective and shared views.

For projection-based collaborative virtual environments such as workbench or CAVE displays, a problem conceptually similar to subjective views arises: How to present individual perspectively correct stereo views for multiple users. Agrawala et al. [ABM*97] approaches this problem with the two-user responsive workbench, which is based on four-way time interleaved display shuttering. Later Fröhlich et al. [FBS*05] described a multi-viewer stereo display for up to four users based on a combination of active shutter based and passive polarization based stereo.

3. System description

The objective of this work was to develop a simple and inexpensive system for using passive tangible props as magic lenses. The basic interaction with the data, such as panning and zooming the filtered lens area, should not require tracking devices.

3.1. Channel separation

The setup uses two commodity projectors as a conventional passive stereo projection system: In such a system, two projectors create superimposed images on a polarization-preserving screen. The projectors are for example connected to a computer with dual-headed graphics output and fitted with one horizontal polarization (HP) and one vertical polarization (VP) filter. The viewer wears inexpensive glasses fitted with an HP filter for the left eye and a VP filter for the right eye, to obtain proper channel separation for the two eyes.

Unlike in a passive stereo setup, in our system the glasses are fitted with either two HP or two VP filters, so under normal viewing conditions, only one of the two projector images is visible. This way, we separate two content channels which we denote as primary and secondary view.

By the use of linear polarization the system suffers from a restriction known from normal passive stereo display systems: users need to look relatively straight to the display surface and may not tilt their head much in order to preserve a proper polarization. In practice this means that users cannot move fully to either side of a wide display area.

3.2. Magic lens

We rely on the unique property of liquid crystal display (LCD) panels: Without power applied, they are designed to retard linearly polarized light by 90 degrees. In conventional LCD screens this property is used together with a polarization filter to control light flow through the panel by applying power selectively to portions of the LCD panel to make it transparent or opaque.

By using passive unpowered LCD panels as 90 degree retarders, the channel separation is effectively inverted through the panel in a two-projector display system using polarization filters. Polarized light from a projection screen that passes through the retarder, which can, for example, have the shape of a handheld magic lens prop, changes from VP to HP or vice versa. See figure 1 for a system sketch.

A similar effect can be obtained with cellophane [Iiz03], but the retardation is typically less than 90 degrees, leading to ghosting effects. The degree of rotation of polarized light depends on the wavelength and thus the polarization rotation varies for different colors and results in a color shift, depending on the used retarder material. When using the LCD lens, this effect is hardly perceiveable.

The separation is surprisingly effective and does not require any electronics held or worn by the user. Since it is purely based on optical properties, typical problems of virtual reality systems such as lag or tracking errors are eliminated.

3.3. Calibration

In order to align the two overlapping display areas of the two projectors we use a simple homography-based calibration procedure. In our case, the calibration is conducted manually. The user has to click four points for each projector image representing the largest area illuminated by both projectors or the region of interest. Using these four point correspondences for each projector, the homography matrix for the geometric warping of the projector images can be calculated. The projected images for the application are then rendered with a two pass rendering technique. Alternatively, this procedure can be automated using a structured light approach as described, for example, in [RvBC02].

4. Interaction Techniques and Applications

4.1. Optical magic lens interaction

The optical magic lens is made from an LCD retarder attached to a handle, in the shape of a conventional reading lens. Panning the filtered view is managed by just moving the lens in front of the display wall. Accordingly, zooming works implicitly as the user moves the lens towards his face or further away. Thus, no tracking data evaluation and computation for basic visual interaction is needed. A limitation of the purely passive prop is that the application is not aware of the user's activities, and feedback to the application needs to be provided through other input devices, such as, for example, a wireless trackball.

The optical magic lens can be employed as a visual filter for a given data representation. Examples are given in section 4.2. Filtered regions by the magic lens are invisible to other users and do not distract them from the given overall view. Every user can have an independent view on the screen and employ a personal optical magic lens, so the information display can be examined by multiple users without disturbing or influencing each other. The number of active users is not restricted by technological constraints such as the tracking system, which is one of the major benefits of our system in contrast to existing approaches. Our setup is only limited to a primary and secondary view that can be perceived by an arbitrary number of users.

For collaborative operations, two users can be equipped with separate subjective views. Rather than letting both users wear the same type of glasses, the first user can choose HP filters, while the glasses of the second user are equipped with VP filters. This allows the users to see either the primary or the secondary view depending on the type of glasses. For example, a user interface can be presented in the native language of each user, while they are simultaneously looking on the same display. This principle also works with a larger number of users divided into two groups. If the application accepts simultaneous input from two users, any disturbances from temporary changes in the display caused by one user such as pop-up menus can be suppressed in the view of the second user. By looking through the optical magic lens, a user may choose to see what the other user is currently doing. See figure 2 for an example.

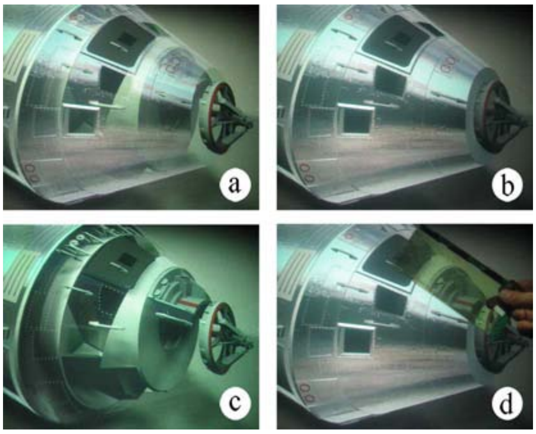

Figure 2: *The projected 3D model seen without polarization filter (a), the first person's view (b), the second person's view (c), and the first person's view with magic lens (d) (3D model courtesy by discreet 3D Studio Max).*

4.2. Application examples

The proposed magic lens can be applied for any 2D or monoscopic 3D visualization application that requires switching between two different views. Examples include:

- arial photographs overlaid by maps (cf. figure 3),
- pairs of ancient and up to date images of cities or buildings,
- Xray vision to see inside an object (cf. figure 2),
- sketches or photos augmented by additional text information,
- different views on complex datasets or
- showing a movie with subtitles in two different languages.

Figure 4 shows GeneView, a visual data mining tool gene data for microarray measurements. Measured protein activity is represented using several visualization techniques such as 2D scatterplots, 3D scatterplots, heatmaps, graphs and cluster based hierarchy visualization. Since the exploration of the data is focused on collaborative analysis comparing selections of genes is very important. Figure 4 shows a heatmap with selected genes marked with blue areas. For collaborative analysis two user select a subset of the gene data and analyse different parts of the gene datasets in their primary view while the marked regions of their collaborator is visible through their magic lens. By switching between the views the collaborators can identify patterns, comparing the two selections. Again, our proposed setup allows users to select their preferred view and collaborate via the magic lens tool.

Figure 3: *A satellite image augmented with a map view using the magic lens (Image courtesy by Google Maps).*

Similarly, individual user interface settings can be used in collaborative work. Figure 5 shows this concept with a mock-up painting application.

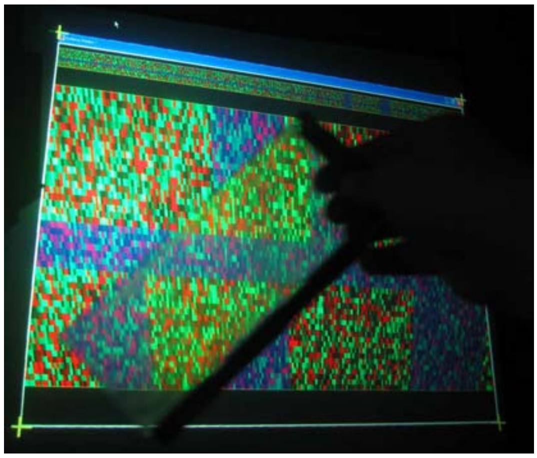

Figure 4: *Gene data visualization via heatmaps.*

All experiments shown were conducted using a front projection setup.

4.3. Augmented Reality applications

Augmented Reality displays can be used to annotate the real world with computer-generated information [RAH98]. Polarization-based interaction allows two novel variations of bringing together virtual and real content in an Augmented Reality display:

- An image printed in black on a transparent foil can be placed over the silver screen, for example showing a physical map. By using only one projector, a digital image can be projected to the physical image rather than a second

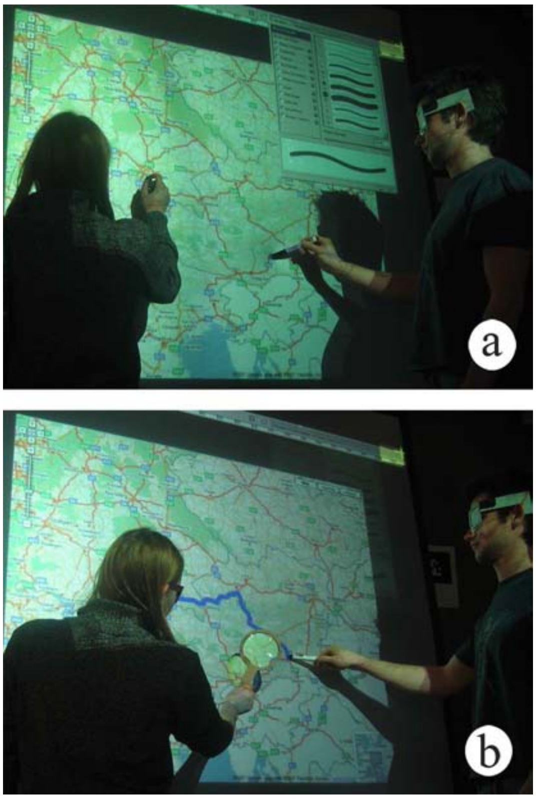

Figure 5: *The right person's view (a) and the left person's view (b) of a collaborative painting application mock-up.*

projected image in a registered way. The physical image is always visible even if the projection system is turned off, and can therefore serve an everyday function for example as an overview plan in a building lobby. The projector is equipped with an HP filter, while the user wears glasses with VP filters. Therefore the user does not perceive the projection unless viewed through an optical magic lens, thereby selectively providing detail or additional information for the physical image.

- An empty transparent foil without any printing is placed over the silver screen. The user can annotate the image projected with HP by writing on the foil with a black felt-tip pen. In order to digitize the physical ink from the felt-tip pen, a video camera with a VP filter takes a snapshot of the screen, thereby eliminating the polarized image. The resulting-black and white image is processed to be used for digital annotations of the projected content. The camera can either be registered offline or on the fly by detecting black fiducials placed on the foil to mark the corners of the projected image. The remaining cross-talk from the projector after capturing the image with imperfect retardation can be filtered using a histogram-based threshold operator. In the result image 6 (b) thresholding was done offline, but can be easily implemented in real-time software. The digitized binary image can either be attached to the digital content directly, or potentially be forwarded to a handwriting recognition software.

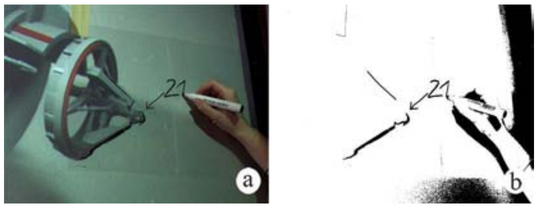

Figure 6: *A person annotating a projected polarized image without polarization glasses (a) and the camera's view with polarization filter and adjustment for post-processing (b).*

5. Discussion

We built a prototype front projection setup to test our ideas and invited people to try out our examples (cf. figure 7). We observed that by having the required projection setup and a sufficent number of polarization glasses and magic lenses, a lot of people could smoothly interact with our system without any initial explanations. We therefore see a great potential of our system to be employed at public spaces, such as museums or schools, where applications as described above could easily be shown as part of an installation in combination with a large-scale display wall.

Figure 7: *Two users exploring an augmented satellite image (Image courtesy by Google Maps).*

The interaction techniques emerging from the proposed tangible magic lens are more intuitive than comparable approaches, as our system relies purely on natural optical properties. This emphasises the ease of use for a big group of potentially inexperienced users. However, the interaction techniques are limited to sole visual interaction such as changing the visible section of the magic lens view. Further interaction with the application data itself requires additional input devices. Another drawback is the fact that all users have to wear polarization glasses in order to use the magic lens as intented.

Summarising, our approach provides the following advantages over other standard magic lens implementations:

- No tracking system is required for basic visual interaction with the projected imagery.
- Instead of using an expensive self-contained handheld device as magic lens, our system employs an inexpensive passive prop to change the visual representation of projected data. Thus, the system scales very well with a high number of users.
- The magic lens itself is light-weight and does not require cable connections which makes the system comfortable to use, even over a longer period of time.
- No tracking data evaluation and re-calculation of visual data for each user is needed which results in a smooth data representation without any delay.
- Only the view "inside" the lens is modified, which is not visible to other users not looking through the lens. The primary view is entirely unmodified leading to undisturbed interaction for other users.
- The basic visual interaction such as panning and zooming the filtered view works implicit and does not require understanding of computer systems and experiences with 3D tracking devices.
- Due to the polarization-based channel separation, the system provides the possibility of employing two subjective views by simply choosing the appropriate polarization glasses.

However, the proposed system also results in a number of restrictions, namely:

- The system requires a two projector setup, as used in passive stereo systems.
- Linear polarisation is depending on the view direction. Tilting the head or looking at the screen from a narrow angle may result in an unclean channel separation.
- Users have to wear polarization glasses in order to achieve channel separation.
- Interaction is restricted to mere visual changes of the data representation using the magic lens. More complex interaction requires additional interaction devices.
- Only two separate views are possible – for the magic lens filtering as well as for the subjective views.

These restrictions imply that the system is rather inappropriate for applications requiring highly complex interaction with the data itself. However, we presented a number of examples where our system seems to be preferrable to existing technologies. In particular, the setup can encourage untrained users to participate in collaborative applications. We see the unique advantages listed above to be a great new opportunity for public installations.

6. Conclusion and Future Work

We introduced a novel interaction tool, the optical magic lens, for large projector-based displays. By using polarization to separate two views and an LCD panel to rotate the direction of polarized light, the lens reveales a view of the hidden secondary layer when using uniform polarization glasses. The optical magic lens is a purely passive prop and does not require expensive 6DOF tracking devices. The arising interaction with the magic lens is very natural as it is handled like a real handheld lens. The two separated views emerging from this approach permit new collaboration techniques relying on subjective views. Since our system does not require a tracking system, interaction with the application software has to be done through other means, such as conventional 2D interaction devices.

In our examples we only showed monoscopic data. However, the setup can be extended to a stereoscopic 3D display by combining the polarization approach with active, time-multiplexed shuttering by using a variation of the technique from [FBS*05]. While resulting in a slightly more complex setup, polarizing shutter glasses can perform channel separation both in the time domain (for stereo vision) as well as in the polarization domain (for separating normal and magic lens image). This allows to build a four channel image (two stereoscopic views and two content channels).

In the future, we see a great potential in using the optical magic lens in massive multi-user presentations and evaluation of complex datasets. Also, in museums or other public institutions this user interface metaphor seems to be promising, as it allows ad-hoc interaction of a great amount of people who do not have any experience with 3D input devices.

7. Acknowledgements

This work is supported by the Austrian Science Fund FWF under contracts Y193, DK W1209-N15 and the Austrian Ministry of Traffic, Innovation and Technology through the FIT-IT project Genopticum. Thanks to Martin Braun for his contributions to the magic lens idea.

References

[ABM*97] AGRAWALA M., BEERS A. C., MCDOWALL I., FRÖHLICH B., BOLAS M., HANRAHAN P.: The two-user responsive workbench: support for collaboration through individual views of a shared space. In *SIGGRAPH* (1997), ACM, pp. 327–332.

[AR05] ASHDOWN M., ROBINSON P.: Escritoire: a personal projected display. *IEEE Multimedia 12*, 1 (2005), 34–42.

[BB05] BEZERIANOS A., BALAKRISHNAN R.: View and space management on large displays. *Computer Graphics and Applications 25*, 4 (2005), 34–43.

[BBF98] BUTZ A., BESHERS C., FEINER S.: Of vampire mirrors and privacy lamps: Privacy management in multiuser augmented environments. In *UIST* (1998), ACM, pp. 171–172.

[BGS01] BAUDISCH P., GOOD N., STEWART P.: Focus plus context screens: combining display technology with visualization techniques. In *UIST* (2001), ACM, pp. 31–40.

[BHG03] BROWN L. D., HUA H., GAO C.: A widget framework for augmented interaction in scape. In *UIST* (2003), ACM, pp. 1–10.

[BSP*93] BIER E. A., STONE M. C., PIER K., BUXTON W., DEROSE T. D.: Toolglass and magic lenses: the see-through interface. In *SIGGRAPH* (1993), ACM, pp. 73–80.

[FBS*05] FRÖHLICH B., BLACH R., STEFANI O., HOCHSTRATE J., HOFFMANN J., KLUEGER K., BUES M.: Implementing multi-viewer stereo displays. In *WSCG* (2005).

[FMM*05] FURUICHI M., MIHORI Y., MURAOKA F., ESTENTHER A., RYALL K.: Dtmap demo: Interactive tabletop maps for ubiquitous computing. In *UbiComp* (2005).

[Fur86] FURNAS G.: Generalized fisheye views. In *CHI* (1986), ACM, pp. 16–23.

[FZC93] FITZMAURICE G. W., ZHAI S., CHIGNELL M. H.: Virtual reality for palmtop computers. *TIS 11*, 3 (1993), 197–218.

[Iiz03] IIZUKA K.: Cellophane as a half-wave plate and its use for converting a laptop computer screen into a three-dimensional display. *Review of Scientific Instruments 74*, 8 (2003), 3636–3639.

[MPL*02] MACKAY W. E., POTHIER G., LETONDAL C., BØEGH K., SØRENSEN H. E.: The missing link: augmenting biology laboratory notebooks. In *UIST* (2002), ACM, pp. 41–50.

[RAH98] REKIMOTO J., AYATSUKA Y., HAYASHI K.: Augment-able reality: Situated communication through physical and digital spaces. In *ISWC* (1998), pp. 68–75.

[RM93] ROBERTSON G. G., MACKINLEY J. D.: The document lens. In *UIST* (1993), ACM, pp. 101–108.

[RvBC02] RASKAR R., VAN BAAR J., CHAI J.: A low-cost projector mosaic with fast registration. In *ACCV* (2002).

[SAKH06] STAADT O. G., AHLBORN B. A., KREYLOS O., HAMANN B.: A foveal inset for large display environments. In *VRCIA* (2006), ACM, pp. 281–288.

[SES99] SCHMALSTIEG D., ENCARNAÇÃO L. M., SZALAVÁRI Z.: Using transparent props for interaction with the virtual table. In *SI3D* (1999), ACM, pp. 147–153.

[SH06] SANNEBLAD J., HOLMQUIST L. E.: Ubiquitous graphics: combining hand-held and wall-size displays to interact with large images. In *AVI* (2006), ACM, pp. 373–377.

[SJA97] SNOWDON D., JÄÄ-ARO K.-M.: A subjective virtual environment for collaborative information visualization. In *Virtual Reality Universe* (1997).

[SKM06] SCHÖNING J., KRÜGER A., MÜLLER H. J.: Interaction of mobile camera devices with physical maps. In *Pervasive Computing* (2006).

[Smi96] SMITH G.: Cooperative virtual environments: lessons from 2d multi user interfaces. In *CSCW* (1996), ACM, pp. 390–398.

[UI97] ULLMER B., ISHII H.: The metadesk: models and prototypes for tangible user interfaces. In *UIST* (1997), ACM, pp. 223–232.

[VCWP96] VIEGA J., CONWAY M. J., WILLIAMS G., PAUSCH R.: 3d magic lenses. In *UIST* (1996), ACM, pp. 51–58.

[WS03] WAGNER D., SCHMALSTIEG D.: First steps towards handheld augmented reality. In *ISWC* (2003), IEEE, pp. 127–135.

[Yee03] YEE K.-P.: Peephole displays: Pen interaction on spatially aware handheld computers. In *CHI* (2003), pp. 1–8.

[YGH*01] YANG R., GOTZ D., HENSLEY J., TOWLES H., BROWN M.: Pixelflex: a reconfigurable multi-projector display system. In *VIS* (2001), IEEE, pp. 167–554.