JanusLights

A Camera-Projection System for Telematic Omni-Presence with Correct Eye Gaze

Chris Hermans[1], Maarten Dumont[1,2], Philippe Bekaert[1]
{chris.hermans, maarten.dumont, philippe.bekaert}@uhasselt.be

[1]Hasselt University - Expertise Centre for Digital Media
transnationale Universiteit Limburg
Wetenschapspark 2, 3590 Diepenbeek, Belgium

[2]Interdisciplinary institute for BroadBand Technology
IBBT - Expertise Centre for Digital Media
Diepenbeek, Belgium

ABSTRACT

We present a camera-projection system for single- or multi-party telepresence which allows for correct eye gaze, and unlike standard videophony, provides a great deal of spatial context.
The unique feature of our system is a combination of omnidirectional video capture and display from corresponding projection centers. In essence, this creates a virtual overlap between the screen and the camera, which results in the participant looking directly into the camera whilst looking at the display. The novelty in this approach is that we correct eye gaze *without* the need of interpolating multiple views.
With a given external calibration of the camera-projector setup, the mapping used for this system has to be calculated only once. This makes it possible for the algorithm to be implemented on commodity hardware and the GPU.

Categories and Subject Descriptors

H.4.3 [**Information Technology and Systems Applications**]: Communications Applications—*Computer conferencing, teleconferencing, and videoconferencing*
; I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Color,Texture*; I.4.8 [**Image Processing and Computer Vision**]: Scene Analysis—*Color,Sensor fusion,Stereo*

Keywords

JanusLights, teleconferencing, videoconferencing, telepresence, eye gaze, omnidirectional video

1. INTRODUCTION

In most standard videoconferencing systems it is impossible for participants to make eye contact, or to infer at whom or what the other participants are gazing. This loss of gaze awareness has a profound impact on the communications that take place [8]. Also, because of the limited spatial nature of a planar viewing window, only a very small amount of context information is available (fig.1, left).

On the other hand, when looking head-on into a spherical mirror, a person looks into her own eyes. This implies correct eye gaze, as her gaze immediately locks onto her reflections gaze. This happens without losing a clear view of the environment (fig.1, right).
Exchanging such views between two (groups of) people creates a natural way of communication. This is the core idea behind our system (fig.2).
In order to achieve this goal, we combine the capture of omnidirectional video by filming a spherical mirror with the corresponding projection onto a diffuse white spherical screen. Both capture and display are performed on a single sphere, forming a single communication device. During the rest of the paper, we will refer to this device as a JanusLight.

The rest of this paper is organized as follows. In section 2 we will discuss the related work. Next, in sections 3 and 4, we will explain the conceptual setup and a more detailed mathematical background for the different mappings used between the input and the output of our system. The prototypes we created for testing purposes are then described in section 5. This is followed by their different application scenarios (section 6), after which we will conclude with results (section 7) and future work (section 8).

2. RELATED WORK

2.1 Image-Based Illumination

Our work shows similarities with work done in the area of image-based illumination. Debevec [7] presented a method that uses measured scene radiance and global illumination in order to add new objects to light-based models with correct lighting. They compute the scene radiance by filming a reflective sphere in the center of the distant scene. Raskar *et al.* [9] animated neutral objects - diffuse white objects with a defined geometry, but without any texture information - with image-based illumination. This is similar to our method on the receiver side, where the neutral spherical screen is illuminated as if it were a reflective mirror.

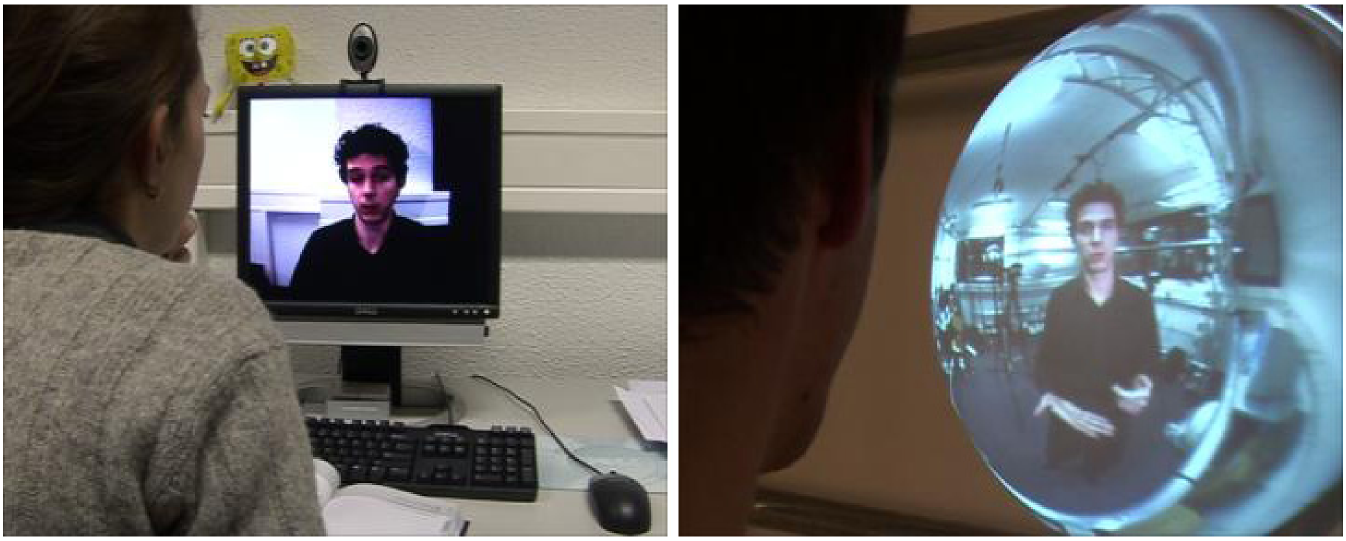

Figure 1: (left) Standard videophony fundamentally suffers from incorrect eye gaze. Because a person has to choose between looking at her computer screen -or- at the camera filming her, direct eye contact is impossible. (right) Our goal: the viewer can watch the spherical screen, which corresponds to looking into the omnidirectional camera.

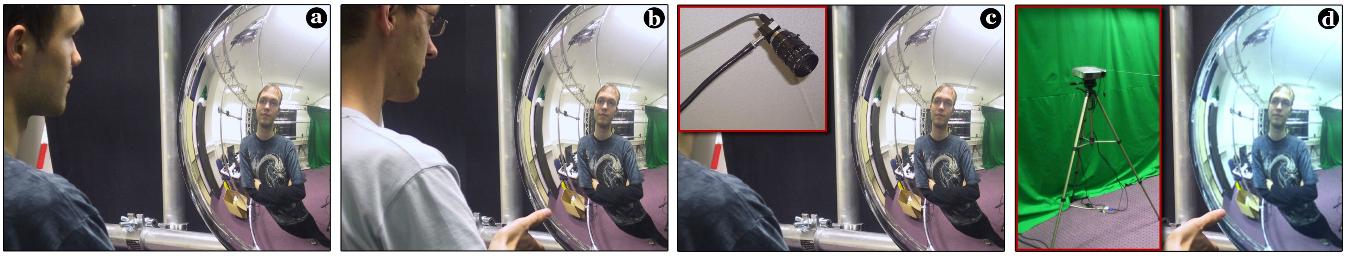

Figure 2: When looking head-on into a spherical mirror, a person looks into her own eyes without losing a clear view of her environment (a). Exchanging such views between two people creates a natural way of communication (b). We try to achieve this core idea by combining the capture of omnidirectional video from a spherical mirror (c), with the corresponding display by projection onto a diffuse white spherical screen (d).

2.2 Display Systems

Spherical display systems with the projectors stored on the inside, such as the Omniglobe [3] or the Magic Planet [2], provide excellent alternatives for our current display system, even though they are more difficult to construct. Using internal projections would obviously alter our mapping method, as the output mapping will become that of a convex mirror, but the main principles would remain the same.

2.3 Gaze Correction in Videoconferencing

A lot of work has already been done on gaze correction in peer-to-peer teleconferencing. In many cases these algorithms boil down to novel-view synthesis from a pair of images, acquired from a stereo camera setup. Previously proposed solutions to handle this problem can be broadly categorized as *model-based* or *image-based*.

For example, Vetter [12] and Yang *et al.* [13] apply a model-based approach. Their methods both use a detailed head model, which they reproject into the cyclopean view.

A more general approach is to use some form of low level stereo matching, followed by an image-based rendering approach (IBR) [5]. These techniques are more widely applicable [4]. In the context of gaze correction, the aim is to synthesize a novel view from a virtual camera that is located approximately where the image of the head will be displayed on the screen for each participant, thus achieving eye contact [6].

This paper demonstrates that the same effect can be achieved without view interpolation, by capturing and projecting video on appropriately curved rather then planar surfaces.

3. CONCEPTUAL SETUP

Our JanusLights - partly diffuse, partly specular spheres - are used in one of two configurations:

- The *many-to-many configuration* has one camera at a distance on top, and a projector from below, while the viewers are located in the lower horizontal space around the globe. A typical setup is depicted in the

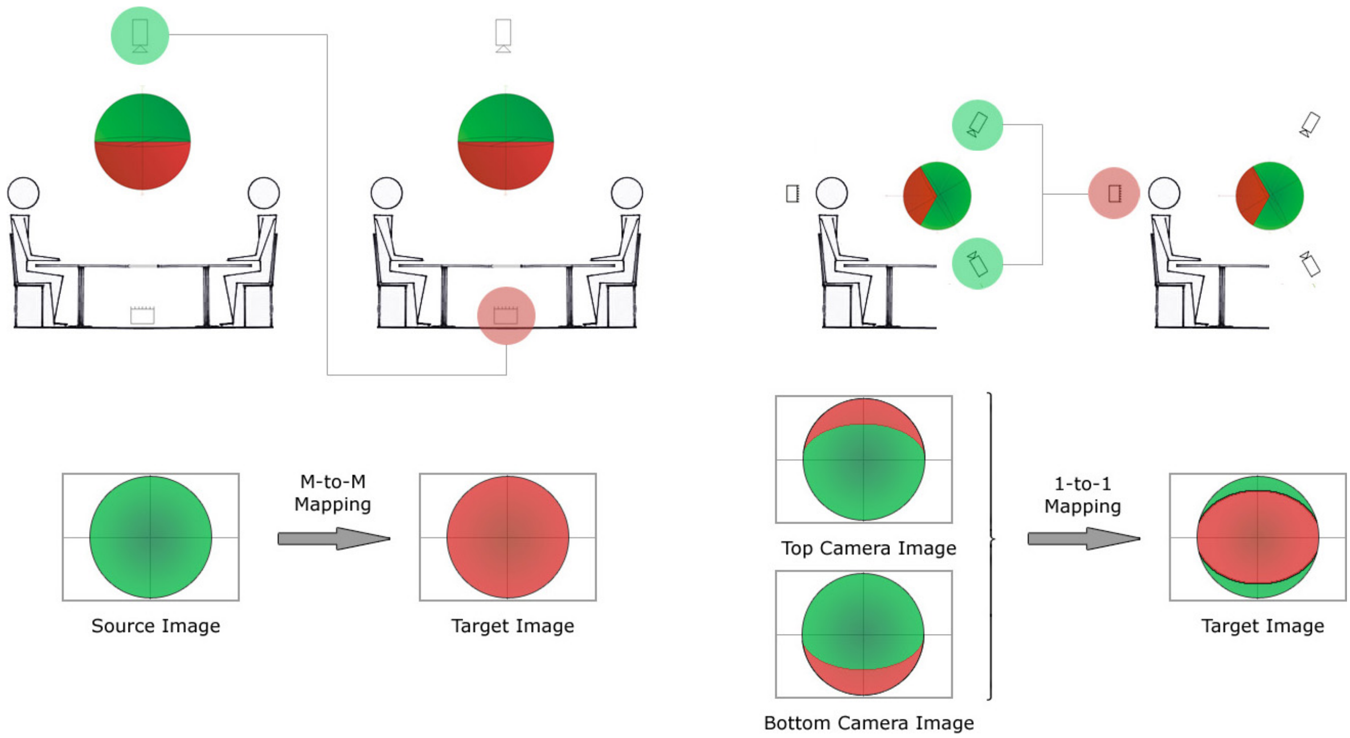

Figure 3: JanusLights are used in one of two configurations. (left) The many-to-many configuration has one camera at a distance on top, and a projector from below, while the viewers are located in the 360 degree space around the globe. (right) The one-to-one configuration has a triangular camera-projector setup with two cameras from behind the globe, and a projector in front of the display screen. The viewers direction is the same as that of the projector, so we put the projector above the shoulder of the user. Green areas indicate specular surfaces viewed by the camera(s), while red stands for the diffuse white surfaces used for projector display.

left part of fig.3, where the users sit around a meeting table with a hole in the center. The image recorded by the camera is transformed by the many-to-many image mapping, and the resulting image is projected on the diffuse lower hemisphere of the JanusLight.

- The *one-to-one configuration* has a triangular camera-projector setup with two cameras behind the globe, and a projector in front of the display screen. The viewers direction is the same as that of the projector, so we put the projector above the shoulder of the user. The images recorded by the cameras are transformed according to the one-to-one image mapping, and the resulting image is projected on the display surface of the JanusLight. In order to acquire the needed information from both cameras, we need to slightly reduce the angle of the vertical field of view for this setup. This results in an eye-like display screen, as can be seen in the right part of fig.3.

4. IMAGE MAPPING

In this section we will first describe a general camera-projector mapping, of which we later specify the variables for both the many-to-many and the one-to-one configuration. For the rest of this section, the used notation refers to vectors shown in fig.4.

4.1 General

In our computations, we assume the (x, y) image coordinates of the input cameras to be centered around the image of the captured sphere. On the projector side we make the similar assumption that the pixel at $(0, 0)$ is projected onto the center of the diffuse surface of the sphere. Both these translations can be achieved easily in a preprocessing step that takes care of perspective distortions. Therefore we make these assumptions without loss of generality.

4.1.1 Output

In both configurations we will do a backward mapping, starting in the coordinate system of the receiver with the origin corresponding to the center of the receiving JanusLight, ending in the coordinate space of the sending sphere. To accommodate for possible imperfections in the constructed hardware, we model the JanusLight as a spheroid Q_{JL}, rather then a sphere. This provides us with a first equation for points $[X, Y, Z, 1]^T$ on the surface.

$$\begin{bmatrix} X & Y & Z & 1 \end{bmatrix} \begin{bmatrix} \frac{1}{a^2} & 0 & 0 & 0 \\ 0 & \frac{1}{b^2} & 0 & 0 \\ 0 & 0 & \frac{1}{a^2} & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = 0 \quad (1)$$

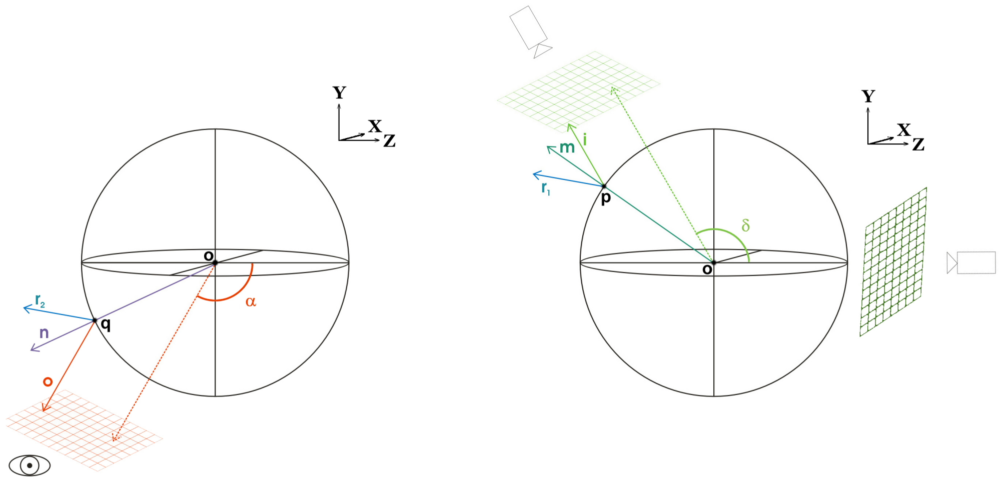

Figure 4: In both configurations, we use a backward mapping from projected pixels to the camera pixels. Rays coming from the projecting device (direction $\vec{o}$) are mapped to their normal $\vec{n}$ or more naturally, their reflection direction $\vec{r_2}$. Using the corresponding reflection $\vec{r_1}$ on the input side and the direction of the camera $\vec{i}$, we can find the intersection of the normal $\vec{m}$ and the coordinate of the wanted pixel in the camera image.

The input of a JanusLight is omnidirectional video, with a projection center in the middle of the sphere. For our calculations, this implicates the use of an affine camera model. This assumption is valid in case of a substantially large camera-sphere and projector-sphere distance, compared to the radius of the sphere. The larger the JanusLight, the larger the required distances. To retain uniformity in our computations, we make a similar assumption on the output end. We define the affine camera matrix as follows:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2}$$

We also know the positions of the cameras, which can be reached with a rotation around the X-axis with angle α. Combining this knowledge with the affine model gives us:

$$\begin{bmatrix} x_{out} \\ y_{out} \\ 1 \end{bmatrix} = A\ R_X(\alpha) \begin{bmatrix} X_q \\ Y_q \\ Z_q \\ 1 \end{bmatrix} \tag{3}$$

Putting these first three equations together, we can see that equating the light intensities arriving at point $q(X_q, Y_q, Z_q, 1)$ on the JanusLight ($q \in Q_{JL}$) to the image coordinates (x_{out}, y_{out}) of the projector device, results in the following group of constraints:

$$\begin{cases} X_q = x_{out} \\ Y_q = \frac{y_{out} - sin(\alpha) Z_q}{cos(\alpha)} \\ Z_q = \frac{y_{out} - cos(\alpha) Y_q}{sin(\alpha)} \geq 0 \\ \frac{X_q^2 + Z_q^2}{a^2} + \frac{Y_q^2}{b^2} = 1 \end{cases} \tag{4}$$

Our goal on the receiving side is to mimic the illumination of a mirror on the sending side. In order to do this, we need to calculate the reflection direction for each point $q \in Q_{JL}$.

We know that for each point X on a quadric Q, the tangent plane π_X is found:

$$\pi_X = QX \tag{5}$$

The tangent plane π_q in point q defines the normal $\vec{n}$:

$$\vec{n} = \left[\frac{X_q}{a^2}, \frac{Y_q}{b^2}, \frac{Z_q}{a^2}, 0\right]^T \tag{6}$$

Given this normal and the projector (output) direction $\vec{o}$, the corresponding reflection direction $\vec{r}$ is calculated as:

$$\vec{r} = 2\vec{n}(\vec{n} \cdot \vec{o}) - \vec{o} \tag{7}$$

4.1.2 *Input*

On the sending side, we need a mapping from the reflection direction to the corresponding image pixel in the camera. First, we calculate the normal from the reflection direction and the camera direction:

$$\vec{m} = \frac{\vec{r} + \vec{i}}{\|\vec{r} + \vec{i}\|} = [X_m, Y_m, Z_m, 0]^T \tag{8}$$

The light ray reflecting on the JanusLight according to the given reflection direction, touches the globe at a point p. The corresponding tangent plane in p is defined as

$$\pi_p = [a^2 X_m, b^2 Y_m, a^2 Z_m, -1]^T \tag{9}$$

From this equation, we can derive p itself

$$\begin{gathered} p(X_p, Y_p, Z_p, 1) \in \pi_p \cap Q \\ \Downarrow \\ p = \frac{[a^2 X_m, b^2 Y_m, a^2 Z_m, \sqrt{a^2 X_m + b^2 Y_m + a^2 Z_m}]^T}{\sqrt{a^2 X_m + b^2 Y_m + a^2 Z_m}} \end{gathered} \tag{10}$$

Once we know the location of the surface point p, we can seek its projection in the corresponding camera - with angle

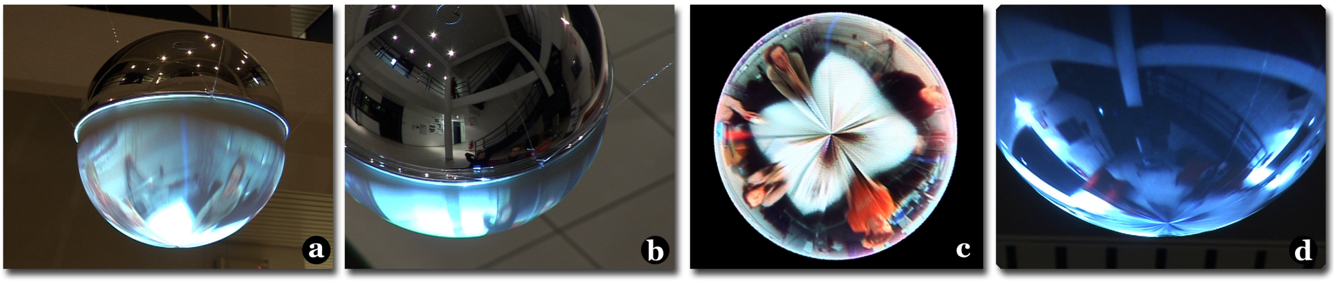

Figure 5: Our many-to-many prototype, as seen from multiple viewpoints: (a) a bottom-view, as seen by the typical user; (b) a top-view; (c) the actual image projected onto the spherical scene; (d) a close-up from another view, originating from a hallway.

δ around the X-axis - thus completing the transformation.

$$\begin{bmatrix} x_{in} \\ y_{in} \\ 1 \end{bmatrix} = A\ R_X(\delta) \begin{bmatrix} X_p \\ Y_p \\ Z_p \\ 1 \end{bmatrix} \tag{11}$$

4.2 Many-to-Many

In this configuration a single camera is positioned at angle $\delta = \frac{\pi}{2}$ in the YZ-plane, while the projector is located at an angle $\alpha = -\frac{\pi}{2}$. The many-to-many mapping uses a small variation to the one mentioned earlier. The reason for this is twofold: (a) viewers will look at the JanusLight from the sides, instead of standing right 'in front' of it ($\alpha = -\frac{\pi}{2}$); (b) multiple viewers will use it at the same time, so the projected image has to be view independent.

Our proposed approach is to map the reflection $\vec{r}$ in the receiving JanusLights to the normal $\vec{n}$ in the sending one.

$$\begin{aligned} &(X_q = x_{out}) \wedge \left(Y_q = \sqrt{b^2 - \tfrac{b^2(x_{out}^2+y_{out}^2)}{a^2}}\right) \wedge (Z_q = y_{out}) \\ &\Downarrow \\ \vec{n} &= [X_n, Y_n, Z_n, 0]^T \\ &= \left[\frac{x_{out}}{a^2}, \frac{\sqrt{b^2 - \frac{b^2(x_{out}^2+y_{out}^2)}{a^2}}}{b^2}, \frac{y_{out}}{a^2}, 0\right]^T \end{aligned} \tag{12}$$

According to the Law of Reflection ($\vec{r} = 2\vec{m}(\vec{m}\cdot\vec{i}) - \vec{i}$), this normal gets mapped to reflection $\vec{r}$ in the sending JanusLight:

$$\begin{aligned} (\vec{m}\cdot\vec{i})\vec{m} &= \tfrac{\vec{r}+\vec{i}}{2} \\ &\Downarrow \\ Z_m[X_m, Y_m, Z_m, 0] &= -\tfrac{[X_n, Y_n, Z_n - 1, 0]}{2} \\ &\Downarrow \\ Z_m \leftarrow \sqrt{\tfrac{1-Z_n}{2}} \ \wedge\ X_m &\leftarrow -\tfrac{X_n}{2Z_m} \ \wedge\ Y_m \leftarrow -\tfrac{Y_n}{2Z_m} \end{aligned} \tag{13}$$

Derivation of the corresponding image point from the normal on the senders side can be done analogue to the derivation in the general mapping.

4.3 One-to-One

In this configuration, we can directly apply the general mapping described earlier. More specifically, the cameras are positioned at angles $\alpha = 0$ and at angle $\alpha = \frac{2\pi}{3}$ in the YZ-plane, while the viewer looks at the scene from $\delta = -\frac{2\pi}{3}$ degrees. Depending on the pixel position in the projected image ($y_{out} \leq 0$ or $y_{out} > 0$), the system chooses the correct camera to perform the mapping.

5. APPLICATIONS

Depending on the configuration used, we can have different application scenarios. Examples of these include:

- The combination of a diffuse and a specular half-sphere, mapping reflections on the camera side to normals on the projector side, generates sensible images from any point of view around the device. This allows for *multi-party many-to-many* telepresence. A meeting room with a central table would an excellent candidate for the use of our JanusLights as videoconferencing devices. Fig.5a shows an example of the above application scenario.

- A triangular camera-projector setup is suited for *one-to-one communication*. Typical applications would be standard videophony, be it a stand-alone system like a telephone, or a webcam-like extension for a personal computer. An example of this application can be seen in fig.6

- JanusLights can also be used as *decorative lighting devices*, offering a point of view into (network-) linked spaces (e.g. pubs, halls, theaters, events,...). For example, fig.7 shows the many-to-many setup at work as part of an artistic exposition at a local theater.

6. IMPLEMENTATION

For our prototypes, we used off-the-shelf half-spherical mirrors, the kind typically used for shop surveillance. Because these mirrors are neither perfectly spherical nor specular, image quality suffered because of the resulting distortions. As mentioned earlier, we tried to compensate for these spatial deformations by modeling our sphere as a spheroid, with

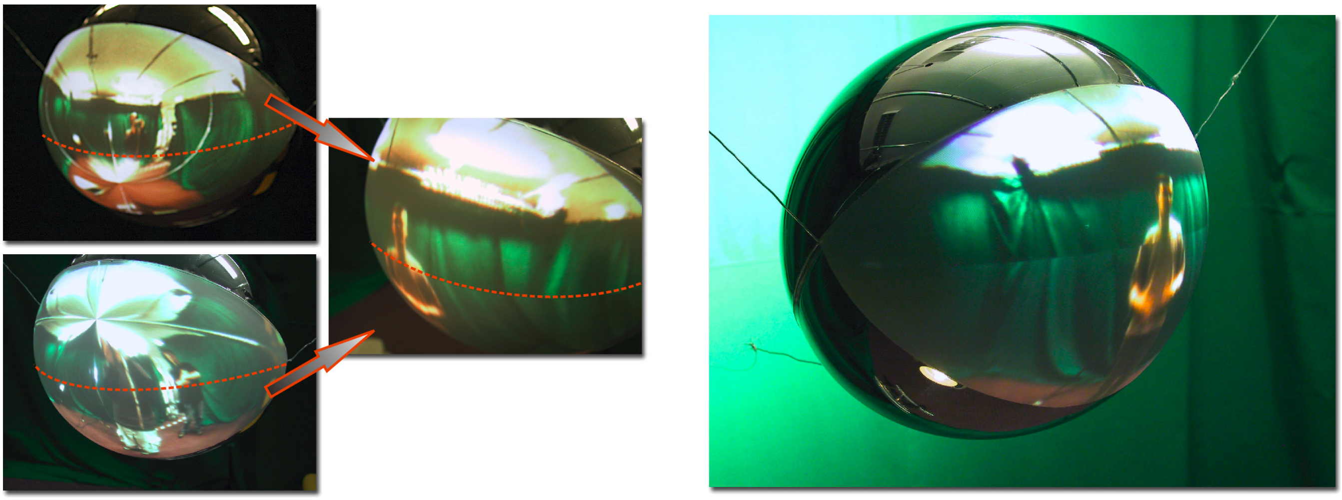

Figure 6: (left) Two camera views are blended into a single image. We merge the halves of the transformed camera images that do not contain their respective vanishing point, and use soft blending to mask artefacts. (right) An example of our one-to-one prototype in action.

a vertical axis of a different length then the horizontal ones. Another hardware limitation on image quality consists of the inherent resolution restrictions of the projector and cameras. It is important to note that both issues are related to the hardware of the prototype, rather than the system concept.

In order to balance the perceived intensities of the pixels when projected onto the spherical display, each pixel in the projected image needs to be multiplied by an attenuation factor that depends on the distance from the projection center. This factor results in an increased luminance value for the outer pixels when compared to the center pixels.

All employed image manipulations are straightforward 2D or 3D operations. This means that they need to be precalculated only once ahead of the rendering step (e.g. most of the mapping code can be calculated as a preprocessing step, which produces a reference table that can be uploaded to the GPU as a texture), or they are well suited for implementation and execution on the GPU.

Several extensions were made to the basic system in order to provide the user with more options for personal use. These additions have no significant impact on system performance.

- During many-to-many communication, users can easily rotate the view by performing a simple 2D rotation on the captured scene. When used as decorative lighting device, applying a small constant rotation results in an increased feeling of affiliation with the linked space, as a non-moving person is able to gradually get a full mental image of the other side.

- The many-to-many output can be divided in a number of equally sized slots - much like a pie-chart - equal to the number of participants on the viewing side. Together with the rotations mentioned in the previous extension, this gives each participant an individual viewing window which she can control.

Our implementation, using the GPU to optimize speed, operates at a real-time rendering speed of 25 frames per second, with no mentionable delays other than network and frame grabber delays. This also accounts for the extensions mentioned above.

7. RESULTS AND DISCUSSION

Overall, the results from our prototype systems (fig.5, fig.6) look promising. We shall now explain and address any remaining artefacts of our prototypes.

7.1 Many-to-Many

As the light rays emitted from the projection device are bundled like a cone, projection on the edges of the display side is suboptimal (fig.5a). There are two possible ways to resolve this issue.

The first method consists of employing carefully placed concave mirrors to bend the projected rays to their corresponding 3D coordinate on the sphere. This method still uses an external light source, but it will require additional careful calibration.

The second method consists of placing the projection device *inside* the sphere, as has been successfully done by commercial system like Omniglobe [3] and Magic Planet [2]. These systems project their imagery onto a convex mirror inside the globe, dispersing the rays along the curved display surface.

Another noticeable issue is the vanishing point at the center of the projected image. This occurs because of the poor quality of the information gathered from the edge of the captured sphere. This is understandable, as there is no way to view the content exactly under the sphere. Fortunately for our applications, discomfort is minimized due to two facts: (a) in an office scenario a JanusLight is usually placed above a uniformly colored meeting table, which means the pixels associated with the artefact are roughly the same color as

Figure 7: JanusLights can be used as decorative lighting devices, offering a point of view into (network-) linked spaces (e.g. pubs, halls, theaters, events,...). (left) A JanusLight located at the top floor of a local theater. (center) A linked JanusLight located at the bottom floor of the same theater. (right) A close-up of a transformed image.

their neighbors; (b) no participant is positioned at the viewing angle where the artefact occurs in the first place. This can be seen in fig.5c, where the vanishing point is located at the center of the white conferencing table.

7.2 One-to-One

Our one-to-one setup cannot afford to suffer from the vanishing point artefact, as it would be located in the center of the display screen. Hence we opted for a two-camera setup, where we merge the halves of the transformed camera images that do not contain their respective vanishing point (fig.6, left). The use of two cameras at the same time results in an improved image quality, as our view of the needed reflection surfaces and the associated light intensities is now available at a better angle.

Merging the two different views however comes at a cost. In order to create a uniform image, we need to carefully align the transformed images. Inaccuracies in the camera positions can result in image distortions. Therefore external camera-calibration is more of a recommendation, rather than a luxury. For our prototypes, we settled for a soft blending approach to mask such artefacts.

7.3 Evaluation and Future Work

Overall, our system seems to give promising results. Our experiments have shown that the concept works, even though there are still technical difficulties. When pixel resolution was high enough to see the eyes of the correspondent, we have witnessed eye gaze to be accurate. In addition, the spherical nature of the display screen made a large amount of context information available to the users. Nevertheless, we are currently looking into a number of ways to improve our results.

First of all, we have employed an affine camera model during our entire approach. We are now looking into the possibility of using perspective cameras in our calculations. While this might complicate calculations to some extent, it should be able to solve the vanishing point artefact in the many-to-many setup. The missing image information at the vanishing point coordinates can then be interpolated.

Secondly, during our calculations we have assumed a known external calibration of our camera-projector configuration. For our prototype systems, achieving this alignment has proved to be an elaborate manual process. An automatic calibration is needed in order to decrease the currently required setup time, at the same time removing image distortions originating from misalignment. We are looking into the work of Svoboda [10, 11] to achieve this calibration.

8. CONCLUSION

We have presented a new camera-projection system for single and multi-party telepresence, based on the combination of omnidirectional video capture with display on a spherical surface.

Unfolding the environmental reflection filmed on a specular sphere yields omnidirectional video with a projection center approximately located at the center of the sphere. During video capture, this provides us with the light intensities arriving at the center of the sphere, coming from all directions. As a result, the full spatial context of the correspondents is unveiled. Compared to the limited information provided by a 2D viewing window, we believe this to be a significant enhancement.
Similar to the calculations used to compute environmental reflection, we can map points on the diffuse spherical display to directions of reflected light (as if the display was't diffuse at all, but a perfect mirror) or their normal direction on the display surface.

If we combine these mappings, a person automatically looks into the camera when watching the display. This accounts for the superior eye-gaze quality of our technique. As long as the direction of the viewer meets the center of projection (which it should, as the viewer looks at the sphere head-on), eye gaze should be correct.

Our implementation, using the GPU to optimize speed, operates at a rendering speed of 25 frames per second, with no mentionable delay other than network and frame grabbing delays.
Whilst using only off-the-shelf materials for our prototypes, we were still able to conceive promising results, showing it to be a working concept. Our work has thus introduced the

possibilities of achieving eye gaze correction and - unlike previous approaches - accomplished this without the need to perform view interpolation.

9. ACKNOWLEDGMENTS

The authors acknowledge financial support on a structural basis from the European Regional Development Fund (ERDF) and the Flemish Government. Part of this research was funded by the BOF-projectfund of Hasselt University.

Furthermore we would like to thank Eric Joris and the people of CREW [1], as well as our colleagues for their help and inspiration.

10. REFERENCES

[1] Crew - multimedia as a protheses. http://www.crewonline.org.

[2] Global imagination. magic planet. http://www.globalimagination.com/products.html.

[3] ARC. Omniglobe: A self-contained spherical display system. ACM Siggraph 2003 Emerging Technologies, 2003.

[4] S. Avidan and A. Shashua. Novel view synthesis in tensor space, 1997.

[5] S. E. Chen and L. Williams. View interpolation for image synthesis. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 279–288, New York, NY, USA, 1993. ACM Press.

[6] A. Criminisi, J. Shotton, A. Blake, and P. H. S. Torr. Gaze manipulation for one-to-one teleconferencing. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 191, Washington, DC, USA, 2003. IEEE Computer Society.

[7] P. E. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of SIGGRAPH98*. ACM, 1998.

[8] J. Gemmell, K. Toyama, C. L. Zitnick, T. Kang, and S. Seitz. Gaze awareness for video-conferencing: A software approach. *IEEE MultiMedia*, 7(4):26–35, 2000.

[9] R. Raskar, G. Welch, K. Low, and D. Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. 1999.

[10] T. Svoboda. *Central Panoramic Cameras, Design, Geometry, Egomotion*. PhD dissertation, Czech Technical University, Sept. 1999.

[11] T. Svoboda, T. Pajdla, and V. Hlav. *Epipolar geometry for central panoramic catadioptric cameras.* In Panoramic Vision: Sensors, Theory and Applications. Springer Verlag, New York, 2001.

[12] T. Vetter. Synthesis of novel views from a single face image. *Int. J. Comput. Vision*, 28(2):103–116, 1998.

[13] R. Yang and Z. Zhang. Eye gaze correction with stereovision for video-teleconferencing. In *ECCV (2)*, pages 479–494, 2002.