Eurographics Italian Chapter Conference (2011)
A. F. Abate, M. Nappi, G. Tortora (Editors)

Visual enhancements for improved interactive rendering on light field displays

Marco Agus[1], Giovanni Pintore[1], Fabio Marton[1], Enrico Gobbetti[1], and Antonio Zorcolo[1]

[1] Visual Computing Group, CRS4, Pula, Italy

Abstract

Rendering of complex scenes on a projector-based light field display requires 3D content adaptation in order to provide comfortable viewing experiences in all conditions. In this paper we report about our approach to improve visual experiences while coping with the limitations in the effective field of depth and the angular field of view of the light field display. We present adaptation methods employing non-linear depth mapping and depth of field simulation which leave large parts of the scene unmodified, while modifying the other parts in a non-intrusive way. The methods are integrated in an interactive visualization system for the inspection of massive models on a large scale 35MPixel light field display. Preliminary results of subjective evaluation demonstrate that our rendering adaptation techniques improve visual comfort without affecting the overall depth perception.

1. Introduction

Recent advances in 3D display design have demonstrated that high resolution 3D display technology able to reproduce natural light fields are now able to closely render the perceptual quality and the unique aura of real 3D objects. These automultiscopic displays offer viewing of high-resolution stereoscopic images from multiple positions without glasses, by generating view-dependent pixels that reveal a different color to the observer based on the viewing angle [ZMDP06]. The most evolute version of this kind of displays are continuous multiview (light field) displays, which provide extremely compelling 3D images [AGG*08], continuous horizontal parallax, and they have been demonstrated to provide better depth discrimination capabilities with respect to standard discrete stereoscopic systems [AGIM10]. However, given the physical limitations of projector-based light field displays and the aliasing due to the discrete nature of the system, visual discomfort is likely to occur, generated by crosstalks, when images are projected too far with respect to the display screen. In this paper, we aim to investigate if rendering adaptation can improve the image quality and reduce discomfort for scenes with depth range greater than the display comfort area. Specifically, we introduce techniques for adapting to light field display characteristics the geometry and appearance of rendered 3D objects in order to deliver comfortable and compelling visualizations. Our main contributions are:

- a non-linear depth remapping scheme to map the scene depth to a predefined range on the display to avoid excessive perceived depth;
- a depth of field simulation method, for blurring the part of the scenes in background and far from the comfort viewing range;
- integration of the display adaptation schemes in an interactive massive model rendering system capable to drive the light field display in order to provide users compelling rendering of objects floating in the space;
- preliminary results of subjective evaluation of the proposed display adaptation method.

Results clearly show that visual comfort is sensibly improved by our adaptation methods, without affecting the overall depth perception. It is, to our acknowledge, the first attempt to study and improve the visual comfort in light field displays by modifying the scene depth content, and for this reason, it represents a substantial enhancement to the state-of-the-art. The rest of the paper is organized as follows: section 2 describes the state of the art in display retargeting and light field geometry generation methods, while section 3 provides an overview on lightfield technology, the sources of visual discomfort, as well as the requirements for correctly projecting geometry. Section 4 describes our display adaptation methods specifically targeted for the light field display, and section 5 provides a description and preliminary results of subjective testing of the proposed methods.

2. Related work

Our work is strictly related to 3D content retargeting for stereo systems and to the image generation methods for projector-based light field display. Here we provide an overview of the state of the art in these subjects.

3D content retargeting for stereo systems Visual discomfort is the main issue arising in disparity-based stereo viewing, mostly caused by crosstalk. In order to achieve comfortable display conditions, various depth adjustments ranging from simple parallax shifting to more advanced methods of new view synthesis are often used [DHH11]. In general, most of the methods presented compare the Geometry of Perceived Depth (GPD) to the Comfortable Viewing Range (CVR). Specifically, these methods employ fixed or dynamic mapping from scene depth to GPD [SH09,JLHE01], or fixed and dynamic Depth of Field simulation [BFSC04, LKC08, LKC09, Kra11] to blur the parts of the scenes falling out of the CVR [SH09]. Other solutions are instead targeted for high-quality off-line stereo post-production, and they are based on classical view interpolation schemes [BGRR09, CBR*07], view morphing methods [MHM*09], and warping techniques originally created for optimizing image or video content [CAA09, KLHG09, WFS*09, LHW*10]. In this paper, we propose methods combining adaptive nonlinear depth modification, and visual methods for providing adequate cues of field of view limitations, such as fading, and depth-of-field based blurring. Our methods are specifically targeted to be employed with light field projector-based display systems.

Image generation for light field displays Automultiscopic displays require generating multiple views for delivering perspective correct images. State-of-the art rendering methods for such displays exploit perspective correction, multiple center of projection (MCOP) geometries [JMY*07] or adaptive sampling [AGG*08] to fit with the display geometry and the finite angular resolution of light beams.However, none of these techniques take into account the problem of fitting scenes with a large depth within the limited depth-of-field of such displays, as required for the visualization of massive and complex scenes. Since now, to our knowledge, no solutions specialized for these kind of displays have been presented. We propose here a solution specialized for real-time presentation of images on projector-based light-field displays which combine adaptive nonlinear depth modification, and visual methods for providing adequate cues of field of view limitations, such as fading, and depth-of-field based blurring. Furthermore, a number of user studies have been carried out to evaluate 3D-TV systems [KKL*08], as well as nonlinear disparity mapping methods for post-processing [LHW*10], or the effects of blur in perceiving depth in binocular stereoscopic images [WBRC11]. We considered similar subjective evaluation methods aimed at demonstrating the comfort improvements of our display adaptation method on light field systems.

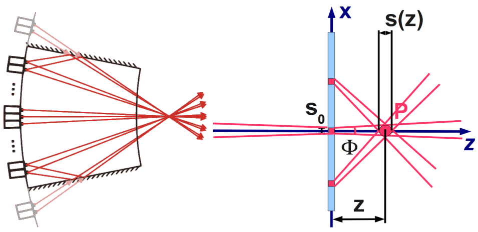

Figure 1: ***Light field display concept.*** *Left: the display is based on projection technology and uses a specially arranged projector array and a holographic screen. Right: the finite angular size of the light beams determines the voxel dimension as a function of distance from the screen.*

3. Light field display overview

The light-field display hardware employed in this work is based on commercially available technology developed by Holografika (see `www.holografika.com`). It is based on projection technology and uses a specially arranged projector array controlled by a PC cluster and a holographic screen (see Fig. 1 left). The projectors are densely arranged at a fixed constant distance from a curved (cylindrical section) screen. All of them project their specific image onto the holographic screen to build up a light field. Mirrors positioned at the side of the display reflect back onto the screen the light beams that would otherwise be lost, thus creating virtual projectors that increase the display field of view. The holographic screen has a holographically recorded, randomized surface relief structure able to provide controlled angular light divergence: horizontally, the surface is sharply transmissive, to maintain a sub-degree separation between views determined by the beam angular size Φ. Vertically, the screen scatters widely, hence the projected image can be viewed from essentially any height. With this approach, a horizontal-parallax-only display is obtained. By appropriately modeling the display geometry, the light beams leaving the various pixels can be made to propagate in specific directions, as if they were emitted from physical objects at fixed spatial locations. Following [JMY*07, AGG*08], we employ a multiple-center-of-projection (MCOP) technique for generating images with good stereo and motion parallax cues. Our technique is based on the approach of fixing the viewer's height and distance from the screen to those of a virtual observer in order to cope with the horizontal parallax only design (see Fig. 2). We assume that the screen is centered at the origin with the y axis in the vertical direction, the x axis pointing to the right, and the z axis pointing out of the screen. Given a virtual observer at **V**, the ray origin passing through a point **P** is then determined by $\mathbf{O} = (E_x + \frac{P_x - E_x}{P_z - E_z}(V_z - E_z), V_y, V_z)$, where **E** is the position of the currently considered projector. The ray connecting **O** to **P** is then used as projection direction to transform the model in normalized projected coordinates. The parameters used for mapping screen pixels to screen 3D points can be determined by automated multi-projector calibration techniques [AGG*08].

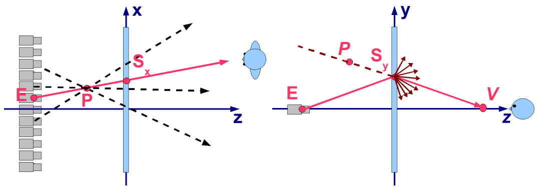

Figure 2: ***Light field display geometry.*** *Left: horizontally, the screen is sharply transmissive and maintains separation between views. Center: vertically, the screen scatters widely so the projected image can be viewed from essentially any height.*

Sources of visual discomfort The main source of geometry error in projector-based light field displays is due to the discrete number of projectors and to the discrete characteristics of the screen, leading to aliasing artifacts [ZMDP06]. Furthermore, in order to drive a light field display, it is necessary to exactly know its geometry. In our case, geometric calibration data is derived by emplying an automated multi-projector calibration technique. However, this procedure is able to provide subpixel errors on the display screen, but this error is degraded as the distance from screen increases, and leads to visible crosstalks and eye discomfort. For both these reasons, aliasing and geometry calibration, we can assume for light field displays limitations similar to those experienced in standard disparity-based stereo systems. Hence, the viewer Geometry of Perceived Depth (GPD) should be limited into a defined volume, the so-called Comfortable Viewing Range (CVR). If the GPD on the display exceeds the limits of CVR, viewers are likely to experience visual discomfort in form of eye strain or headache, and in general their experience will be drastically degraded.

Rendering requirements The 3D display characteristics and limitations impose constraints to the interaction and rendering system in order to have a compelling visualization and reduce rendering artifacts. Specifically, the following requirements have to be taken into account for the implementation of a specialized rendering system:

- the spatial resolution of the display is variable with respect to depth, approximately according to the equation $s(z) = s_0 + 2\|z\|\tan(\frac{\Phi}{2})$, where z is the distance to the holographic screen, and s_0 is the pixel size on the screen surface [AGG*08] (see Fig. 1 right);
- the calibration technique minimizes errors on the screen surface; thus, the effective field of depth of the display is reduced not only because of the diminishing spatial resolution but also of the calibration accuracy;
- because of the display geometry, the angular field of view is limited, and allows presentation of objects only within well defined angular bounds.

Thus, the best viewing performances are obtained when (a) the scene is kept centered with respect to the screen; (b) the scene remains inside a limited depth range; and (c) the frequency details of the objects are adapted to the display spatial accuracy. In the next section, we'll see how to best respond to these constraints by adapting the rendered 3D content (Sec. 4).

4. Rendering adaptation for light-field displays

Given the geometric characteristics of a projector-based light field display, rendering of complex scenes requires 3D content adaptation in order to provide comfortable viewing experiences in all conditions. In our approach, we modify both the visual appearance and the geometry of rendered models to improve visual experiences while coping with the limitations in the effective field of depth and the angular field of view of the light field display.

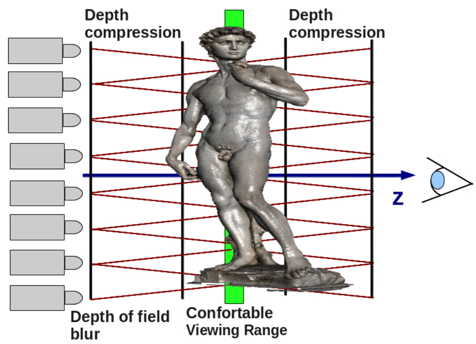

Figure 3: *A schematic view of the display adaptation method: non-linear depth remapping is applied for parts of the scene out of comfort range, while a depth of field blur is applied for the parts of the scene behind the comfort range.*

Overview Since the limited comfortable viewing range of the light field display, rendering adaptation schemes are needed to reduce the discomfort due to the parts of the scene which overcome that area. Furthermore, another target of the method should be to drive the attention of observer only to parts of the scene inside the comfort area. To reach these goals, we apply the following adaptation schemes:

- a depth remapping method, which modifies the rendered geometry in order to reduce the scene depth range and to constrain most of the scene to stay inside the comfortable viewing range. Non-linear depth compression is employed for objects out of the comfort area.

- a depth of field simulation scheme, which blurs the parts of the scene in the background, as to drive users to focus on foreground area.
- a frame fade-out method, in order to reduce clipping artifacts due to the borders of the light-field display.

Figure 3 shows a schematic view of the adaptation methods proposed here.

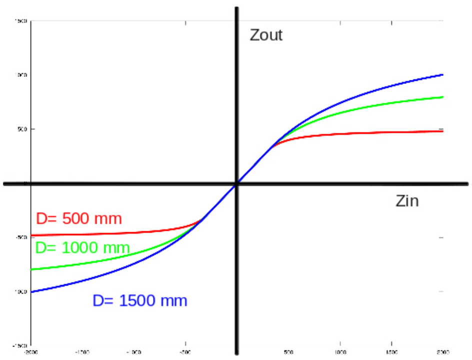

Figure 4: ***Depth mapping function.*** *In order to reduce artifacts when objects fall out of the display comfortable depth range, a non-linear depth-mapping function is applied (see eq. (1)).*

Depth remapping The field of depth of the display determines the largest depth a display can visualize with a defined minimum resolution. As we have seen, the smallest feature (voxel) the display can reconstruct is a function of the distance from the screen, the angular resolution, and the calibration accuracy. A trivial solution for coping with field of depth limitations, would be to clip all the geometry primitives which fall outside a predefined depth range. This approach preserves the 3D shape of rendered objects, but clipping planes introduce annoying artifacts that reduce the quality of experience. On the other hand of the spectrum is the approach of dynamically remapping the depth range of the scene to the comfortable viewing range. Using the common approach of linearly remapping the geometry to the relatively narrow depth range of the display works better, but causes an important "flattening" of the objects, reducing their 3D quality. In our approach, we propose a non-linear depth correction, which leaves unmodified the depth near the screen plane, while drastically reducing its range as the geometry is moved further with respect to the screen. Specifically, the remapped depth is computed according to the equation

$$z_{out} = \begin{cases} B + \frac{z_{in}-B}{1+\frac{z_{in}-B}{D_B-B}} & \text{if } z_{in} < B \\ z_{in} & \text{if } B \le z_{in} \le F \\ F + \frac{z_{in}-F}{1+\frac{z_{in}-F}{D_F-F}} & \text{if } z_{in} > F \end{cases} \quad (1)$$

where F and B are the front and back distance until which the objects are assumed to have comfortable appearance without any remapping, while and D_F and D_B are the front and back largest comfortably renderable depths. The non-linear depth mapping function is represented in Fig. 4: the maximum output depth asymptotically tends to a constant value D (in figure 4 we considered $F = B = 300mm$ and $D_F = D_B = 500, 1000, 1500mm$). Fig. 5 shows an example of the effects of the application of this depth mapping function ($F = 300mm$ and $D_F = 500mm$).

The parts of the object close the the display preserve their geometry, and are thus perceived as "fully 3D", while the parts far from it are flattened.

Depth of field Quality of experience can be improved by also modifying visual appearance, rather than only geometry, for coping with the limitations in the effective field of depth and the angular field of view of the light field display. As indicated in figure 3, we apply a depth-dependent blur to adapt the frequency content of the scene to the display spatial resolution. In this way, aliasing artifacts due to the degradation of spatial accuracy can be drastically reduced. To this end, we consider a typical depth of field effect, commonly employed to simulate the effect of a real camera. According to this effect, focused objects appear to be sharp within the distance range, called the depth-of-field (DOF), around the focal plane, whereas the rest looks blurred with respect to distances to the focal plane. This effect can be usefully employed in light field displays, where the depth-dependent angular resolution can be related to the Circle of Confusion of real lens systems [LKC09]. Nevertheless, the light field display discomfort area affects also parts of the scene very close to the viewer, and in that portion a depth of field blur would result in unrealistic artifacts. For this reason, the effect can be applied only to the background area, as shown in figure 3. In our system, in order to exploit hardware capabilities of modern graphics card, we considered an image-based two-passes depth-of-field simulation method similar to [Ngu07]. First of all, we computed the Circle of Confusion with respect to the comfort viewing range of the light field display and linearly dependent to the depth-dependent spatial resolution of the light field display. Given the circle of confusion, we next employ a post-processing pixel shader which takes as input the original image and a downsampled and pre-blurred version of the same, and uses a variable size kernel approximating the circle of confusion to blend between the original and the pre-blurred image. Stochastic Poisson sampling is considered in order to reduce the number of samples, and the blend between values is weighted by a blurriness value, which is linearly dependent with respect to the depth of the samples. Figure 6 shows an example of the effect of depth-of-field blur to the background parts of the scene.

Frame fade-out In addition to depth range limitations, another source of discomfort is the limited angular workspace of the projector-based display, which causes the abrupt disappearance of the object when the viewer is looking at positions not sampled in the display ray space. This happens when the ray connecting the eye position to a object point

Figure 5: ***Depth compression.*** *Left: rendering quality for object portions too far from the screen surface is insufficient because of limited spatial accuracy and precision in calibration. Right: non-linear depth compression reduces artifacts while minimizing modification of shape in the neighborhood of the screen surface.*

Figure 6: ***Depth-of-field blur.*** *Left: rendering quality for object portions too far from the screen surface is insufficient because of limited spatial accuracy and precision in calibration. Right: An image-based depth-of-field blur is applied to reduce the effect of depth compression and to adapt scene frequency content to the light field display spatial resolution.*

does not reach the projectors, and causes discomfort if the virtual object is in front of the display (see Fig. 7 left).

We have found that a simple but effective method to reduce discomfort consists of applying a color blending which fades the color object to the background in areas nearby unsampled positions (see Fig. 7 right). This makes the object smoothly fade-out, having a ghost-like appearance, rather than abruptly clipping them. As for depth-dependent, this effect is realized in a fragment shader operating as a post-pass, and taking as input the shaded color and the depth buffer.

5. Results and evaluation

Implementation The display adaptation schemes presented in this paper have been tested in a interactive system capable of rendering massive surface models, implemented on Linux using C++, OpenGL and GLSL. The integrated system enables multiple naked-eye users to perceive detailed giga-triangles models as floating in space, responsive to their actions. Given the size of the model, adaptive out-of-core structures are used both for rendering and for the geometric queries required by the interaction. We use a distributed image generation system implemented on a cluster, with a front-end PC coordinating many rendering back-end PCs. The light field 3D display is capable of visualizing $35M\,pixels$ by composing images generated by 72 SVGA LED commodity projectors illuminating a $160 \times 90cm$ holographic screen. The display provides continuous horizontal parallax within a 50° horizontal field-of-view with 0.8° angular accuracy. The pixel size on the screen surface is $1.5mm$. The rendering back-end runs on an array of 18 Athlon64 3300+ Linux PCs equipped with two NVIDIA 8800GTS 640MB (G80 GPU) graphics boards running in twin-view mode. Each back-end PC has thus to generate $4 \times 800 \times 600$ pixels using two OpenGL graphics boards based on an old G80 chip. Front-end and back-end nodes are connected through Gigabit Ethernet and communicate through OpenMPI 1.2.6. In this paper, we show the results obtained with the inspection of the *David*$0.25mm$ model, composed of $970M$ triangles.

Qualitative results We tested our schemes under various conditions and with different parameters. Here we provide pictures taken by hand-held photo camera. It is really difficult to convey the impression provided by the light field

Figure 8: ***Live capture.*** *Top row: no depth remapping (left), depth remapping with* $D_F = 1500mm$ *(center), depth remapping with* $D_F = 500mm$ *(right). Bottom row: no adaptation (left), depth remapping (center), depth remapping with depth of field blur (right)*

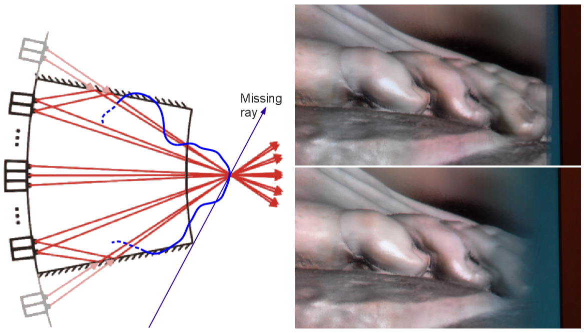

Figure 7: ***Workspace limitation.*** *Left: Objects in front of the display suddenly disappear when looking at them from positions not sampled by projector rays. Right: color blending fades color to background when objects are close to the display vertical and horizontal borders. The effect is depth-dependent, and takes into account the position of the virtual object with respect to the screen.*

display, but here we want to highlight how our display retargeting methods can successfully reduce artifacts and provide better comfort to users when interacting with the system. The top row of figure 8 shows the effects of using our depth remapping scheme when part of the scene is out of the comfort range in front of the screen: at left, the image rendered without adaptation, at center the image rendered with $D_F = 1500mm$, and at right with $D_F = 500mm$. It appears evident that lower values for D_F improve visual quality, but with the cost of reducing depth perception. The bottom row of figure 8 shows instead the difference in applying depth remapping and depth of field blurring for parts of the scene which are behind of the screen and outside with respect to the screen. In this case, left image is obtained without applying any adaptation, center image is obtained with depth remapping ($D_B = 1000mm$) , and finally right image is obtained with depth remapping together with depth of field blurring. It appears evident that depth of field blurring reduces the effect of depth compression and drives users to focus on the part of the scene inside the comfortable viewing range.

Evaluation Furthermore we also carried out a subjective evaluation in order to find whether proposed adaptation schemes can attenuate discomfort due to the limitations of light field displays. To this end, we considered perceptual evaluation methods similar to those employed in [WBRC11, Kra11]. We asked 22 persons to provide their opinion about static images rendered on our interactive visualization system. Subjects mostly with no experience with light field displays were asked to compare static rendering of David surface model under different conditions. A 2-Interval Forced Choice design was considered, and users were faced against 6 different static stimuli, presenting the model in positions out of the comfort range of the display (3 times in front of the screen surface, and 3 times behind the screen surface). In each trial, users were randomly shown two different renderings of the model (with adaptation, and with no adaptation), and were asked to indicate which images provided better viewing comfort and better depth perception. For all tests, after trial and error tuning, we decided to employ the following parameters: $F = B = 300mm$ and $D_F = D_B = 1000mm$. We obtained the following results: with respect to visual comfort, for stimuli in front of the screen, in 100% of cases subjects revealed to feel more comfortable when non-linear depth compression is active, while for stimuli behind the screen, in 75% of cases they felt more comfortable when depth compression and depth of field blur is employed. Instead, with respect of depth perception, surprisingly, even if the scene depth range is drastically reduced, in 57% of cases users revealed to have better depth and 3D perception when the visual adaptation is employed for stimuli in front of the screen, and in 60% of cases for stimuli behind the screen. These preliminary results can lead us to the following conclusions: our display adaptation methods clearly improve visual comfort in light field displays, and they do not heavily affect depth perception, even when the depth range is drastically reduced. Discussions with subjects also revealed another interesting effect related to the perception of depth of field in background: some subjects revealed to feel uncomfortable with it, while others revealed to have better depth perception, since a blurred background drives users attention to the foreground scene. In any case, these effects should be more deeply studied, and further evaluation is needed.

6. Conclusions and Future Work

We have presented methods for adapting to light field display characteristics the geometry and appearance of rendered 3D objects in order to deliver comfortable and compelling visualizations. The proposed display adaptation method consists of a non-linear depth remapping scheme and depth of field simulation technique, and it is integrated in an interactive massive model rendering system capable to drive the light field display in order to provide users compelling rendering of objects floating in the space. We also performed subjective testing, and our preliminary results show that visual comfort is sensibly improved by our adaptation methods, without affecting the overall depth perception. As future work, we aim to carry out a parameters analysis of our display retargeting method, and a more extensive perceptual evaluation in order to find the effects of the non-linear remapping and of the depth-of-field blur.

Acknowledgments This work is partially supported by the EU FP7 Program under the DIVA project (290277). We also acknowledge the contribution of Sardinian Regional Authorities. The David dataset is courtesy of the Digital Michelangelo Project.

References

[AGG*08] Agus M., Gobbetti E., Guitián J. A. I., Marton F., Pintore G.: GPU accelerated direct volume rendering on an interactive light field display. *Computer Graphics Forum 27*, 2 (2008), 231–240. 1, 2, 3

[AGIM10] Agus M., Gobbetti E., Iglesias Guitián J. A., Marton F.: Evaluating layout discrimination capabilities of continuous and discrete automultiscopic displays. In *Proc. Fourth International Symposium on 3DPVT* (2010). 1

[BFSC04] Bertalmio M., Fort P., Sanchez-Crespo D.: Real-time accurate depth of field using anisotropic diffusion and programmable graphics cards. In *Proc. 3DPVT* (2004), pp. 767–773. 2

[BGRR09] Bleyer M., Gelautz M., Rother C., Rhemann C.: A stereo approach that handles the matting problem via image warping. In *Proc. IEEE CVPR* (2009), pp. 501–508. 2

[CAA09] Carroll R., Agrawal M., Agarwala A.: Optimizing content-preserving projections for wide-angle images. *ACM Trans. Graph. 28* (2009), 43:1–43:9. 2

[CBR*07] Criminisi A., Blake A., Rother C., Shotton J., Torr P. H.: Efficient dense stereo with occlusions for new view-synthesis by four-state dynamic programming. *Int. J. Comput. Vision 71* (2007), 89–110. 2

[DHH11] Daly S., Held R., Hoffman D.: Perceptual issues in stereoscopic signal processing. *Broadcasting, IEEE Transactions on 57*, 2 (june 2011), 347 –361. 2

[JLHE01] Jones G. R., Lee D., Holliman N. S., Ezra D.: Controlling perceived depth in stereoscopic images. In *Proc. SPIE Stereoscopic Displays and Virtual Reality Systems VIII* (2001), pp. 42–53. 2

[JMY*07] Jones A., McDowall I., Yamada H., Bolas M. T., Debevec P. E.: Rendering for an interactive 360 degree light field display. *ACM Trans. Graph 26*, 3 (2007), 40. 2

[KKL*08] Kapteina R. G., Kuijsters A., Lambooij M. T., IJsselsteijn W. A., Heynderickxa I.: Performance evaluation of 3d-tv systems. In *Proc. of SPIE Electronic Imaging* (2008). 2

[KLHG09] Krähenbühl P., Lang M., Hornung A., Gross M.: A system for retargeting of streaming video. *ACM Trans. Graph. 28* (2009), 126:1–126:10. 2

[Kra11] Kraus M.: *Using Opaque Image Blur for Real-Time Depth-of-Field Rendering*. Institute for Systems and Technologies of Information, Control and Communication, 2011, pp. 153–159. 2, 6

[LHW*10] Lang M., Hornung A., Wang O., Poulakos S., Smolic A., Gross M.: Nonlinear disparity mapping for stereoscopic 3d. *ACM Trans. Graph. 29* (2010). 2

[LKC08] Lee S., Kim G. J., Choi S.: Real-time depth-of-field rendering using point splatting on per-pixel layers. *Comput. Graph. Forum 27*, 7 (2008), 1955–1962. 2

[LKC09] Lee S., Kim G. J., Choi S.: Real-time depth-of-field rendering using anisotropically filtered mipmap interpolation. *IEEE Trans. Vis. Comp. Grap. 15* (2009), 453–464. 2, 4

[MHM*09] Mahajan D., Huang F.-C., Matusik W., Ramamoorthi R., Belhumeur P.: Moving gradients: a path-based method for plausible image interpolation. *ACM Trans. Graph. 28* (July 2009), 42:1–42:11. 2

[Ngu07] Nguyen H.: *Gpu gems 3*. Addison-Wesley Professional, 2007. 4

[SH09] Sun G., Holliman N.: Evaluating methods for controlling depth perception in stereoscopic cinematography. In *Proc. SPIE Stereoscopic Displays and Virtual Reality Systems XX* (2009). 2

[WBRC11] Wang J., Barkowsky M., Ricordel V., Callet P. L.: Quantifying how the combination of blur and disparity affects the perceived depth. In *Proc. SPIE* (2011). 2, 6

[WFS*09] Wang Y.-S., Fu H., Sorkine O., Lee T.-Y., Seidel H.-P.: Motion-aware temporal coherence for video resizing. *ACM Trans. Graph. 28* (2009), 127:1–127:10. 2

[ZMDP06] Zwicker M., Matusik W., Durand F., Pfister H.: Antialiasing for automultiscopic 3d displays. In *in Eurographics Symposium on Rendering* (2006). 1, 3