A PROTOTYPE FOR PRACTICAL EYE-GAZE CORRECTED VIDEO CHAT ON GRAPHICS HARDWARE

Maarten Dumont[1], Steven Maesen[1], Sammy Rogmans[1,2] and Philippe Bekaert[1]

[1]*Hasselt University – tUL – IBBT, Expertise centre for Digital Media, Wetenschapspark 2, 3590 Diepenbeek, Belgium*
[2]*Multimedia Group, IMEC, Kapeldreef 75, 3001 Leuven, Belgium*

Keywords: Prototype, practical, video chat, eye-gaze correction, GPU, real-time.

Abstract: We present a fully functional prototype to convincingly restore eye contact between two video chat participants, with a minimal amount of constraints. The proposed six-fold camera setup is easily integrated into the monitor frame, and is used to interpolate an image as if its virtual camera captured the image through a transparent screen. The peer user has a large freedom of movement, resulting in system specifications that enable genuine practical usage. Our software framework thereby harnesses the powerful computational resources inside graphics hardware, to achieve real-time performance up to 30 frames per second for 800×600 resolution images. Furthermore, an optimal set of finetuned parameters are presented, that optimizes the end-to-end performance of the application, and therefore is still able to achieve high subjective visual quality.

1   INTRODUCTION

Peer-to-peer interactive video chat is becoming increasingly popular, but still has some major drawbacks which prevent a definitive breakthrough in the common public. One of the most important reasons is that the participant is not able to simultaneously look at the screen and the camera, leading to the loss of eye contact (Gemmell et al., 2000). We therefore present a fully functional prototype (see Fig. 1) that corrects the eye gaze of the peers by using multiple cameras, solving typical problems for genuine practical usage. Previous solutions such as (Criminisi et al., 2003) either implement their framework on commodity CPUs, resulting in a very low framerate when sufficient visual quality is required. On the opposite side, solutions such as (Schreer et al., 2001; Baker et al., 2002) involve the use of expensive dedicated hardware, or have an unpractical camera setup. Others optimize parts of the application, such as multi-camera video coding (Chien et al., 2003; Guo et al., 2005) for efficient data communication and real-time view synthesis (Yang and Pollefeys, 2003; Geys and Van Gool, 2004; Nozick et al., 2006) on graphics hardware, but neither of them integrate and optimize the end-to-end performance for eye gaze-corrected video chat. The end-to-end performance for camera interpolation is optimized in (Rogmans et al., 2008), but they assume a rectified two-fold camera setup and no

Figure 1: Peer setup of our prototype.

illumination changes due to the use of the Middlebury dataset (Middlebury, 2001).

Our prototype uses a practical camera setup that can be easily integrated in the monitor frame, and its framework harnesses the powerful computational resources inside the Graphics Processing Unit (GPU) to achieve real-time performance on commodity PCs. Section 2 of the paper describes the system architecture in detail, which proposes the use of several carefully selected and slightly adapted algorithms that are appropriate for implementation on graphics hardware. The end-to-end system thereby achieves both high speed and quality with only few constraints. We

Figure 2: Data flow and overview of our system architecture.

also provide a number of GPU specific optimizations in Section 3 to ensure real-time application performance. Section 4 discusses the results of the prototype, Section 5 ultimately concludes the paper, and Section 6 deals with future work.

2 SYSTEM ARCHITECTURE

As depicted in Fig. 2, the main functionality of our system consists of four consecutive processing modules that are completely running on a GPU. In an initial step, images $I_1, \ldots, I_N$ are fetched from N cameras $C_1, \ldots, C_N$ that are closely aligned along the screen. The first module performs lens correction and image segmentation, as a form of preprocessing, to enhance both the quality and speed of the consecutive view interpolation.

This second module interpolates an image I_v, as seen with a virtual camera C_v that is positioned behind the screen and produces a consistent depth map Z_v. The image I_v is computed as if camera C_v captured the image through a completely transparent screen.

However, the synthesized image still has a number of noticeable artifacts in the form of erroneous patches and speckle noise. The third refinement module is therefore specifically designed to tackle these problems by detecting photometric outliers based on the accompaning depth map.

In a final step, the depth map Z_v is also analyzed to dynamically adjust the system and thereby avoids heavy constraints on the user's movements.

Next to the main processing on graphics hardware that synthesizes I_v, the camera C_v needs to be correctly positioned to restore eye contact between the participants. An eye tracking module thereby concurrently runs on CPU and determines the user's eye position that will be used for correct placement of the virtual camera at the other peer.

By sending the eye coordinates to the other peer, the input images $I_1, \ldots, I_N$ do not have to be sent over the network, but can be processed at the local peer.

This results in a minimum amount of required data communication – i.e. the eye coordinates and the interpolated image – between the two participants.

2.1 Preprocessing

Camera lenses, certainly when targeting the low-budget range, induce a radial distortion that is best corrected. Our system relies on the use of the *Brown-Conrady* distortion model (Brown, 1966) to easily undistort the input images pixel-based on the GPU.

Figure 3: The preprocessing module segments the input camera image.

Each input image I_i with $i \in \{1, \ldots, N\}$ is consequently segmented into a binary foreground silhouette S_i (see Fig. 3), to allow the consecutive view interpolation to adequately lever the speed and quality of the synthesis process. Two methods of segmentation are supported; Greenscreening according to (1), where R_{I_i}, G_{I_i} and B_{I_i} are the red, green and blue components of I_i. For clarity the pixel location (x, y) has been omitted.

$$S_i = \begin{cases} 1: & G_{I_i} > \tau_g \cdot (R_{I_i} + G_{I_i} + B_{I_i}) \\ 0: & G_{I_i} \leq \tau_g \cdot (R_{I_i} + G_{I_i} + B_{I_i}) \end{cases} \quad (1)$$

The second method is able to subtract a real-life background (Magnor et al., 2005) according to (2), where I_{B_i} is the static background picture and $\tau_g, \tau_f, \tau_b, \tau_a$ are experimentally determined thresholds which are subjected to parameter finetuning. For shadow removal, the cosine of the angle $\widehat{I_i I_{B_i}}$ between the color component vectors of the image pixel $I_i(x, y)$ and the static background pixel $I_{B_i}(x, y)$ is determined. As a

Figure 4: The view interpolation module generates a virtual image and joint depth map.
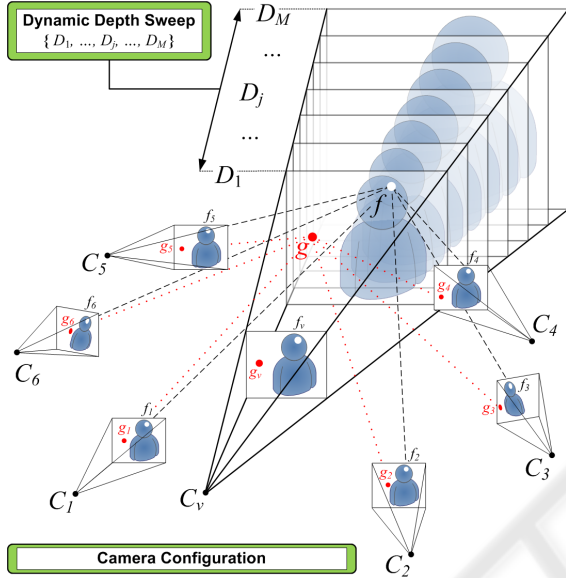

Figure 5: Concept of the plane sweep algorithm.

final step, the silhouette is further enhanced by a single erosion and dilation (Yang and Welch, 2002).

$$
S_i = \begin{cases} 1: & \|I_i - I_{B_i}\| > \tau_f \text{ or} \\ & \|I_i - I_{B_i}\| \geq \tau_b \text{ and } \cos(\widehat{I_i I_{B_i}}) \leq \tau_a \\ 0: & \|I_i - I_{B_i}\| < \tau_b \text{ or} \\ & \|I_i - I_{B_i}\| \leq \tau_f \text{ and } \cos(\widehat{I_i I_{B_i}}) > \tau_a \end{cases} \tag{2}
$$

Both methods are evaluated on a pixel basis and require very little processing power, while still being robust against moderate illumination changes.

2.2 View Interpolation

To interpolate the desired viewpoint (see Fig. 4) we adopt and slightly modify a plane sweeping approach based on (Yang et al., 2002). As depicted in Fig. 5, the 3D space is discretized into M planes $\{D_1, \ldots, D_M\}$ parallel to the image plane of the virtual camera C_v. For each plane D_j, every pixel f_v of the virtual camera image I_v is back-projected on the plane D_j by (3), and reprojected to the input images I_i according to (4). Here $\mathbf{T}_j$ is a translation and scaling matrix that defines the depth and extent of the plane D_j in world space.


Figure 6: Reprojection from virtual to input images.

The relationship between these coordinate spaces is represented in Fig. 6.

$$
f = \mathbf{V}_v^{-1} \times \mathbf{P}_v^{-1} \times \mathbf{T}_j \times f_v \tag{3}
$$
$$
f_i = \mathbf{P}_i \times \mathbf{V}_i \times f \tag{4}
$$

Points on the plane D_j that project outside a foreground silhouette in at least one of the input images, are immediately rejected – e.g. point g in Fig. 5 – and all further operations are automatically discarded by the GPU hardware. This provides a means to lever both speed and quality because segmentation noise will, with a high probability, not be available in all N cameras. Otherwise, the mean (i.e. interpolated) color ψ and a jointly defined custom matching cost κ are computed as in (5).

$$
\psi = \sum_{i=1}^{N} \frac{I_i}{N}, \quad \kappa = \sum_{i=1}^{N} \frac{\|\psi - I_i\|^2}{3N} \tag{5}
$$

As opposed to (Yang et al., 2002), we propose the use of all input cameras to compute the matching cost. The plane is swept for the entire search range $\{D_1, \ldots, D_M\}$, and the minimum cost – together with the corresponding interpolated color – is per pixel selected on a *Winner-Takes-All* basis, resulting in the virtual image I_v and a joint depth map Z_v (see Fig. 4).

2.3 Joint View/Depth Refinement

The interpolated image calculated in the previous section still contains erroneous patches (see Fig. 4, magnified in Fig. 7) and speckle noise due to illumination changes, partially occluded areas and natural homogeneous texturing of the human face. These errors are even more apparent in the depth map Z_v and we therefore propose a photometric outlier detection algorithm that detects and restores the patches in Z_v.

To suppress the spatial high frequency speckle noise, we consequently run a low-pass Gaussian filter over the depth map.

In a final step, the refined depth map is used to recolor the interpolated image with the updated depth

Figure 7: Joint view/depth refinement module concept.

Figure 8: (a) The proposed filter kernel, and (b–d) the outlier detection concept.

values. As opposed to other geometrically correct approaches (Lei and Hendriks, 2002), we thereby significantly enhance the subjective visual quality.

2.3.1 Erroneous Patch Filtering

To detect erroneous patches, we propose a filter kernel as depicted in Fig. 8a. For every pixel z_v of depth map Z_v, a two dimensional depth consistency check is performed for its neighbourhood λ according to (6), where ε is a very small constant to represent the depth consistency. λ thereby defines the radius of the filter kernel, and the maximum size of patches that can be detected.

$$
\begin{aligned}
\|Z_v(x-\lambda,y) - Z_v(x+\lambda,y)\| < \varepsilon \text{ or} \\
\|Z_v(x,y-\lambda) - Z_v(x,y+\lambda)\| < \varepsilon
\end{aligned}
\quad (6)
$$

If the area passes the consistency check in one of the dimensions, the depth pixel z_v – and therefore the joint image pixel f_v – is flagged as an outlier if z_v does not exhibit the same consistency by exceeding a given threshold τ_o. Eq. (7) shows the outlier test when a depth consistency is noticed in the X-dimension, an analogous test is used in case of consistency in the Y-dimension.

$$
\left\| Z_v(x,y) - \frac{Z_v(x-\lambda,y) + Z_v(x+\lambda,y)}{2} \right\| > \tau_o \quad (7)
$$

After performing the proposed filter kernel, the center of patches – as conceptually represented in Fig. 8b and Fig. 8c – are detected. Consistently, a standard morphological grow algorithm is executed in a loop, which causes the detected center to grow only if the neighbouring pixels exhibit the same depth consistency as the initial outliers. As depicted in Fig. 8d, the complete patch is thereby detected. As a final step for the patch filtering, the morphological grow is reversed and the detected patch is filled with reliable depth values from its neighbourhood. Since all of these operations are implemented on a pixel basis, they are inherently appropriate for implementation on a GPU, achieving a tremendous speedup compared to a generic CPU algorithm.

2.3.2 Speckle Noise Filtering

Due to the nature of the human face, a significant amount of large homogeneous texture regions are present. As indicated by (Scharstein and Szeliski, 2002) these areas cause the depth map to contain spatial high frequency speckle noise. The noise is most effectively filtered by a low-pass filter, but eliminates the geometrical correctness of the depth map. Therefore as opposed to methods such as (Lei and Hendriks, 2002), we rather enhance the subjective visual quality instead of geometrical accuracy.

A standard 2D isotropic Gaussian filter is applied on the depth map and thanks to its separable convolution properties, it can even be highly optimized on graphics hardware.

2.3.3 Recoloring

All of the previous steps involve changing the depth map Z_v, which is normally – due to the plane sweep – jointly linked to the image color in I_v. To restore this link, the image I_v is recomputed with an updated $\mathbf{T}_j$ matrix, according to the filtered depth information. Since erroneous patches and speckle noise are now filled or leveled with consistent – rather than geometrically correct – depth values, the recolored image is interpreted as plausible and thereby subjectively regarded as higher quality.

2.4 Movement Analysis

To avoid heavy constraints on the participant's movement, a large depth range has to be scanned. This actually infers a lot of redundant computations, since the head of the user only spans a small range. We therefore propose to dynamically limit the effective depth range to $\{D_{min}, \ldots, D_{max}\}$ similar to (Geys et al., 2004), through a movement analysis on the normalized depth map histogram. This implicitly causes a quality increase of the plane sweep, as the probability of a mismatch due to homogeneous texture regions is significantly reduced. Moreover, all M depth planes can be focused as $\{D_1 = D_{min}, \ldots, D_M = D_{max}\}$,

which leverages the dynamic range and thereby significantly increases the accuracy of the depth scan. Three separate cases can be distinguished, as the user moves in front of the screen:

- *Forward*: If the user moves forward, he will exit the active scanning range. Therefore, the histogram will indicate an exceptionally large number of detected depth pixels towards D_{min}.

- *Stable*: The histogram indicates a clear peak in the middle, this resolves to the fact that the user's head remains in the same depth range.

- *Backward*: Analog to forward movement of the user, the depth histogram will indicate a peak towards D_{max}.

As depicted in Fig. 9a to f, we fit a Gaussian distribution function $G(\mu, \sigma)$ with center μ and standard deviation σ on the histogram. The effective depth range is updated according to (8) and (9), where b_1, b_2 are constant forward and backward bias factors that can be adopted to the inherent geometry of the scanned object. D'_{min} represents the previous minimal depth, and $\Delta = D'_{max} - D'_{min}$ for denormalization.

$$D_1 = D_{min} = D'_{min} + (\mu - b_1 \cdot \sigma)\Delta \quad (8)$$

$$D_M = D_{max} = D'_{min} + (\mu + b_2 \cdot \sigma)\Delta \quad (9)$$

As the user performs forward or backward movement, the center μ of the Gaussian fit changes and dynamically adapts the effective scan range of the system. The image will briefly distort in this unstable case (see extreme cases in Fig. 9a and c), but will quickly recover as the depth scan is adapted for every image iteration. A real-time high framerate therefore increases the responsiveness of the system, and is able to achieve fast restabilization. Normal moderate speed movement will thereby not be visually noticed by the participants.

2.5 Concurrent Eye Tracking

To restore the eye contact between the video chat participants, the camera C_v needs to be correctly positioned. Eye tracking can be performed more robust and efficiently on CPU, and is therefore executed concurrently with the main processing of the system. Afterall, the N input images are implicitly available on system memory before they are sent to the video memory of the GPU.

Face and eye candidates are detected in every input image as described in (Hsu et al., 2002), but we optimized the detection process by predicting the eye positions based on previous frames and using a hierarchical approach to search for eye pixels. An epipolar cross check is consequently used to estimate the most
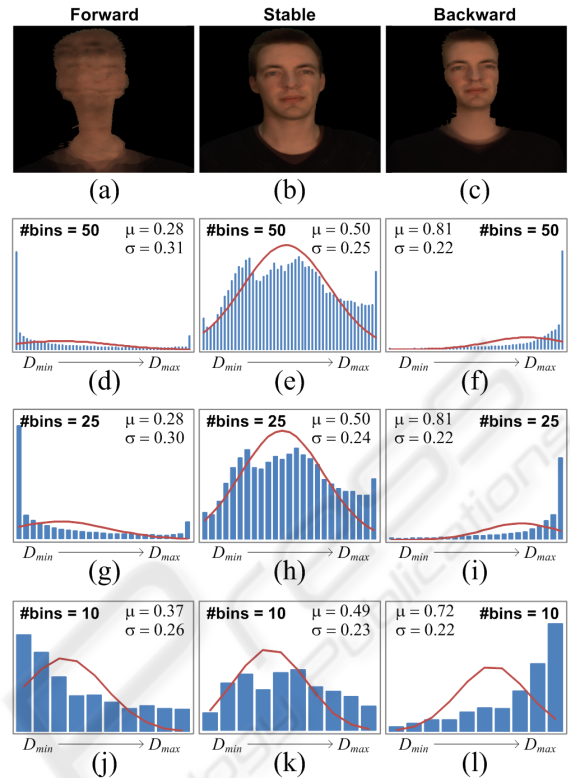

Figure 9: Overview of the depth histogram-based movement analysis in normalized coordinates.

reliable eye candidates in all images. The two best eye candidates are ultimately used to triangulate their 3D positions.

The 3D eye position is then mirrored towards the screen, resulting in the correct virtual viewpoint that is needed to restore the eye contact between the system users. The coordinates are adjusted in a manner that places the two screens in a common space, as if the two screens were pasted against each other. Hence, this creates the immersive effect of a virtual window into the world of the other participant.

2.6 Networking

Our prototype system sends the eye coordinates over the network, and therefore the requested image I_v can be computed locally at the peer that captures the relevant images. These cross computations bring the required network communication to a minimum, by avoiding the transfer of N input images. The total peer-to-peer communication thereby exists out of the synthesized images and the eye coordinates.

Real-time speed can therefore easily be achieved over various types of networks without the need of any compression. Due to the segmentation result S_v, the system can be efficiently equipped with *Region-*

of-Interest coding, to enable the use over low-data rate networks as well.

3 OPTIMIZATIONS

Our framework harnesses the powerful computational resources of the graphics hardware to ensure real-time performance. The use of carefully selected and adapted algorithms allows us to exploit the GPU for general-purpose computations, a technique that is often referred to as General-Purpose GPU (GPGPU).

GPGPU is traditionally only capable of utilizing a part of the GPU resources due to its computer graphics nature (Owens et al., 2007). We optimized the preprocessing to bypass this constraint, which leverages the GPU utilization and execution speed without any noticeable loss of quality.

Being an additional advantage, the histogram for movement analysis is approximated by reducing its number of bins, since only the correct Gaussian distribution fit is required.

As this approximation infers dynamic looping inside the GPU, we have built-in support for dynamic GPU programming to be able to optimize and unroll loops at run-time.

3.1 Improved Utilization

Standard lens distortion is generally corrected on a pixel-basis level, but can be approximated by applying an equivalent geometrical undistortion to small image tiles. Since a GPU pipeline exists out of a geometry and pixel processing stage, the lens correction can hence be ported from the pixel to the geometry stage. Initially, the computational complexity of the undistortion is inverse proportional to the granularity of the tessellation, resulting in a speedup without any visual quality loss if the tile size is chosen correctly. Next to this fact, the pixel processing stage is clear to perform the consecutive segmentation processing in a single pipeline pass, which significantly leverages the utilization of the GPU.

3.2 Accelerated Histogramming

For the movement analysis, the essential part is deriving the parameters μ and σ to adjust the dynamic range of the depth scan. As depicted in Fig. 9d to l, we are able to approximate the histogram by reducing the number of bins, without a large impact on the Gaussian parameters. Heavily reducing the number of bins (see Fig. 9j to l) causes the center μ to become less accurate, as it is shifted towards the center of the effective scan range. An optimal trade-off point
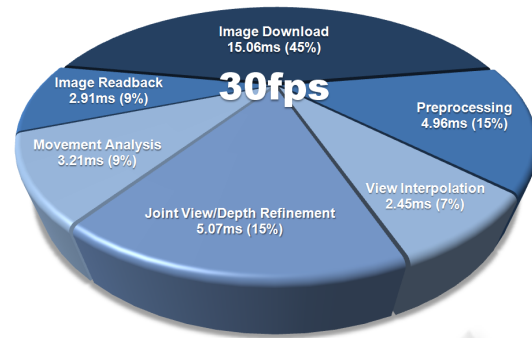

Figure 11: Workload profiling.

can therefore be defined, since the accuracy loss will cause the responsiveness of the system to decrease.

3.3 Dynamic GPU Programming

Similar to the dynamic looping needed to variate the number of bins in the histogramming, looping inside the GPU is very often required. To enhance the execution speed, our framework has built-in support to generate and compile the GPU programming code at run-time. Hereby, the loops can always be unrolled and internally fully optimized.

4 RESULTS

Our optimal prototype setup is built with $N = 6$ auto-synchronized Point Grey Research Grasshopper cameras mounted on an aluminum frame, so they can be closely aligned to the screen (See Fig. 1). The presented camera setup avoids large occlusions, and has the potential to generate high quality views since no image extrapolation is necessary. We have used the *Multi-Camera Self-Calibration* toolbox (Svoboda et al., 2005) to calibrate the camera setup offline, but a built-in camera setup into the screen would avoid this procedure due to fixed calibration parameters. Our software framework runs on an Intel Xeon 2.8GHz, equipped with 2GB system memory and an NVIDIA GeForce 8800GTX graphics card. Communication with the GPU is done through OpenGL, and it is programmed with the high level language Cg.

Final quality results are shown in Fig. 10, under moderate variable illumination conditions, but with a fixed set of finetuned practical system parameters grouped in Table 1. Some small artifacts along the ears and chin, together with minor ghosting around the neck, can still be noticed due to limitations of the joint view and depth refinement. Nonetheless, the images maintain their integrity and are regarded

Figure 10: Eye-gaze corrected images with variable illumination and fixed optimal parameters.

as high subjective visual quality, while they convincingly seem to be making eye contact with the reader.

Table 1: Set of optimized system parameters.

Module	Parameter	Value
Preprocessing	τ_g	0.355
	τ_f	0.010
	τ_b	0.002
	τ_a	0.998
View Interpolation	N	6
	M	35
Joint View/Depth Refinement	λ	20
	ε	0.2
	τ_o	0.3
Movement Analysis	b_1	2.0
	b_2	2.0
	#*bins*	15

A detailed workload profiling for the main processing modules can be seen in Fig. 11, with image and camera resolutions of 800×600 pixels. A large amount of time is needed to perform the six input image download and synthesized image readback to and from the GPU. Hence, the typical data locality importance is illustrated, and the core reason to implement all main processing steps on graphics hardware is motivated. The rather large weight for the preprocessing is mainly due to the $N = 6$ times execution on all input images. Furthermore, a significant amount of processing is concentrated in the refinement and movement analysis, as it levers the quality independent of the amount of input images.

Summing up the different timings of the individual modules, we can reach a confident speed of 30fps, but in our experimental setup is limited by 15Hz support in the cameras and Firewire controller hardware. We foresee these specifications as for genuine practical usage, since the end-to-end system is optimized for minimum practical constraints.

5 CONCLUSIONS

We have presented a system prototype for practical eye-gaze correction between two video chat participants, with a minimal amount of constraints. A sixfold camera setup that is closely aligned along the screen, is proposed for possible integration into the monitor frame. Besides the convenient camera placement, the presented setup also avoids large occlusions, and has the potential to generate higher quality views compared to camera setups that require constant image extrapolation.

Our software framework harnesses the computational resources of the graphics hardware through GPGPU, and is able to achieve real-time performance. A framerate of 30fps can be achieved for 800×600 image resolutions, but our experimental setup is limited by 15 Hz cameras. Thanks to a movement analysis, the system provides the user with a large freedom of movement.

We have improved the end-to-end performance of the system, by introducing optimizations that have no noticeable loss of visual quality. A fixed set of finetuned parameters is able to generate interpolated images with high subjective visual quality – rather than being geometrically correct – under variable conditions. Moreover, the system specifications thereby enable genuine practical usage of convincing eye-gaze corrected video chat.

6 FUTURE WORK

Future work will focus on improving the movement analysis. Furthermore, multiple objects can be intelligently distinguished by combing silhouette information, enabling eye-gaze corrected multi-party video conferencing.

In the largest extent, background objects can be recognized and interpolated with correct motion parallax, providing a true immersive experience for the participants.

ACKNOWLEDGEMENTS

REFERENCES

Baker, H. H., Bhatti, N. T., Tanguay, D., Sobel, I., Gelb, D., Goss, M. E., Culbertson, W. B., and Malzbender, T. (2002). The coliseum immersive teleconferencing system. In *Proceedings of the International Workshop on Immersive Telepresence*, Juan-les-Pins, France.

Brown, D. C. (1966). Decentering distortion of lenses. *Photometric Engineering*, 32(3):444–462.

Chien, S.-Y., Yu, S.-H., Ding, L.-F., Huang, Y.-N., and Chen, L.-G. (2003). Efficient stereo video coding system for immersive teleconference with two-stage hybrid disparity estimation algorithm. In *ICIP 2003: Proceedings of the 2003 International Conference on Image Processing*, pages 749–752.

Criminisi, A., Shotton, J., Blake, A., and Torr, P. H. S. (2003). 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. IEEE Computer Society.

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

Geys, I., Koninckx, T. P., and Van Gool, L. (2004). Fast interpolated cameras by combining a gpu based plane sweep with a max-flow regularisation algorithm. In *3DPVT '04: Proceedings of the 3D Data Processing, Visualization, and Transmission, 2nd International Symposium*, pages 534–541, Washington, DC, USA. IEEE Computer Society.

Geys, I. and Van Gool, L. (2004). Extended view interpolation by parallel use of the gpu and the cpu. In *Videometrics VIII: Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, volume 5665, pages 96–107.

Guo, X., Gao, W., and Zhao, D. (2005). Motion vector prediction in multiview video coding. In *ICIP 2005: Proceedings of the 2005 International Conference on Image Processing*.

Hsu, R.-L., Abdel-Mottaleb, M., and Jain, A. K. (2002). Face detection in color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):696–706.

Lei, B. J. and Hendriks, E. A. (2002). Real-time multi-step view reconstruction for a virtual teleconference system. *EURASIP Journal on Applied Signal Processing*, 2002(1):1067–1087.

Magnor, M., Pollefeys, M., Cheung, G., Matusik, W., and Theobalt, C. (2005). Video-based rendering. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Courses*, New York, NY, USA.

Middlebury (2001). Middlebury stereo vision page. www.middlebury.edu/stereo.

Nozick, V., Michelin, S., and Arquès, D. (2006). Real-time plane-sweep with local strategy. In *Journal of WSCG*, volume 14.

Owens, J. D., Luebke, D., Govindaraju, N., Harris, M., Krger, J., Lefohn, A. E., and Purcell, T. J. (2007). A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum*, 26(1):80–113.

Rogmans, S., Lu, J., and Lafruit, G. (2008). A scalable end-to-end optimized real-time image-based rendering framework on graphics hardware. In *Proceedings of 3DTV-CON, The True Vision, Capture, Transmission, and Display of 3D Video*, pages 129–132, Istanbul, Turkey.

Scharstein, D. and Szeliski, R. (2002). A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47(1–3):7–42.

Schreer, O., Brandenburg, N., Askar, S., and Trucco, M. (2001). A virtual 3d video-conference system providing semi-immersive telepresence: A real-time solution in hardware and software. In *Proceedings of the eBusiness-eWork Conference*, pages 184–190, Venice, Italy.

Svoboda, T., Martinec, D., and Pajdla, T. (2005). A convenient multi-camera self-calibration for virtual environments. *PRESENCE: Teleoperators and Virtual Environments*, 14(4):407–422.

Yang, R. and Pollefeys, M. (2003). Multi-resolution real-time stereo on commodity graphics hardware. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 211–220, Madison, WI, USA. IEEE Computer Society.

Yang, R. and Welch, G. (2002). Fast image segmentation and smoothing using commodity graphics hardware. *Journal of Graphics Tools*, 7(4):91–100.

Yang, R., Welch, G., and Bishop, G. (2002). Real-time consensus-based scene reconstruction using commodity graphics hardware. In *PG '02: Proceedings of the 10th Pacific Conference on Computer Graphics and Applications*, page 225, Washington, DC, USA. IEEE Computer Society.