This is a preliminary version of an article published in
Proc. of Content Generation and Coding for 3D-Television, High Tech Campus, Eindhoven, The Netherlands, June 2006
by Thomas Haenselmann, Marcel Busse, Stephan Kopf, Thomas King, Wolfgang Effelsberg

Multicamera Video-Stitching

Thomas Haenselmann, Marcel Busse, Stephan Kopf,
Thomas King and Wolfgang Effelsberg
University of Mannheim
Applied Computer Science IV
68161 Mannheim, Germany
Email: {haenselmann,busse,kopf,king,effelsberg}
@informatik.uni-mannheim.de

***Abstract*— We propose a stitching algorithm for multi-camera environments which allows to concatenate views with differing centers of projection into a single panoramic image. A common trajectory is defined in two source images to be merged. It serves as a cut that allows to stitch them together. Usually, the layout of the cuts does not allow to stitch both images together naively. Thus, two convex combinations of a warped and a canonic coordinate are applied so that both source images fit together at the cutting edge while the inevitable distortion decreases towards the borders of the image to obtain a natural appearance.**

In this work, we will particularly investigate the side effects of using multiple perspectives for moving images.

I. Introduction

Cameras and image sensors have recently become almost as cheap and available as scalar sensor which are used for temperature or light measurements. The *Stanford Multi-Camera Array* project is an early example for the simultaneous usage of more than 100 cheap CCD cameras [1], [2], [3], [4].

Other than scalar values which can be displayed on a virtual map or which can simply be aggregated, it is not obvious how to display a massive amount of (possibly uncalibrated) images, particularly in a way that makes sense for a human observer.

Consolidating all images into a single one could be a possible solution. Similar attempts have been made in the field of panoramic images, in which a series of pictures are stitched to one another to produce a continuous view. For a long time, panoramic images have been considered feasible only, if all images have the same focal point respectively if the camera does not alter its location. In this paper we devise a novel method for creating panoramic views from images with varying focal points. The specific problems which arise here are described in the following section along with prior attempts to create panoramic images from movies. Section III suggests a basic warping scheme as a solution. In Section IV we evaluate the basic algorithm and identify a couple of shortcomings which are solved by an extended warping scheme.

In Section V, we hint at problems which are unique only to multiperspective panoramic videos. The outlook in Section VI sketches future improvements to reduce the amount of human interaction.

II. Related Work

In the context of our paper we will distinguish between truly monoperspective panoramic images and multiperspective imaging.

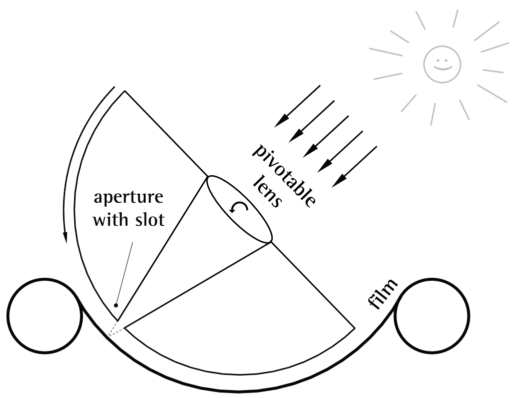

Fig. 1. *Short Rotation Cameras* produce panoramic images by continuously moving a slice-shaped beam over the photographic material thus capturing fields of view of almost 180.

A. Monoperspective Panoramic Images

Panoramic images have been known for more than a century with early applications in war photography, e.g., during the American Civil War in 1860 [5]. Here, photos were captured while rotating a tripod-mounted camera around its optical axis. The panoramic image was simply obtained by setting up the photos next to one another.

In the 20th century the so-called *rotating lens cameras* have been engineered. The mode of operation is depicted in Figure 1.

A lens is mounted pivotable in front of the film. While taking a picture, it rotates between a starting- and finishing-angle. A likewise rotatable aperture with a very narrow slot close to the film prevents the photographic material from being exposed to light more than once while the lens is rotating. During the rotation, a long and narrow light beam will move over the conically bend photographic material like on a photocopier. The entire procedure reminds of a magnetic tape being recorded. The difference is that the photographic material stands still while the light beam exposes it. When these images are laid on a flat surface, straight lines like a horizon line seem to be bent. The image only appears natural if the photo is viewed bend in the same way as was done when it was captured.

With the advent of digital cameras, panoramic imaging became popular with a larger audience. Here, views with a wide angle are produced by stitching together images of a normal aspect ratio of 4:3. Ideally, the images are produces with a tripod-mounted camera. This ensures a fixed focal point also known as *center of projection*. By rotating the camera around its vertical axis, only its viewing direction is altered. This means, that the projection of the three-dimensional world onto the CCD-chip will never change. The rotation itself only makes one part of the image disappear while another moves in. Unfortunately, this does not mean that images can be put together by a simple concatenation because the rotation changes the vanishing points within the images. This is most obvious in architectural photos, in which lines being parallel in the real world converge against a common vanishing point in the projection.

Prior to stitching two images together, a perspective dewarping of one of them or preferably even of both at the same time has to be carried out. This process must be applied in a common image space. Mapping the images into such a space is usually done by applying either a tubular or a spherical projection as shown in Figure 2 [6].

It can be seen from the left side of the figure that a coordinate (x, y) in the image plane maps to an angle θ and the height h which are defined as

$$\theta = arctan\left(\frac{x}{f}\right)$$
$$h = \frac{y}{\sqrt{x^2 + f^2}}$$

The focal length is denoted with f. In both cases, he tubular and the spherical projection, the perspective of an object will not change. Intuitively, a part of an object which is hidden will never be revealed by the rotation of the camera. This changes in the case of *Multiperspective Panoramic Images*.

B. Multiperspective Panoramic Images

Multiperspective means that the panoramic image consists of patches which do not have a common projection center but which are taken from changing viewpoints. This makes stitching particularly difficult or impossible in a naive way since overlapping parts of neighboring images which may in principal show the same content can not be aligned. This is due to the fact, that changing the viewpoint corresponds to a rotation of the objects in the real world. This may hide parts which had been seen before the rotation and reveal new insights afterwards. As a consequence, none of two neighboring images will exhibit simple cuts where one image can be aligned with its neighbor.

One of the earliest instances of multiperspective imaging is the animated cartoon *Pinocchio* by Walt Disney Productions which was made in 1940. The film starts with a virtual camera flying over a small village. In contrast to conventional techniques used at that time, the movement of the camera is no mere pan over a scene. In fact the camera seems to perform a rotation at the same time which alters its viewing direction continuously. The effect was produced by drawing a panoramic image (of ratio 3:1) showing an overview of the village. However, when viewing the image from one end to the other, the perspective seems to change gradually from house to house which creates an impression of a strangely warped scene. The actual shot in the movie was simply made by panning a focused view over the panoramic drawing thus showing only a small clipping at a time. The artists who had been in charge for this scene must have had a very sophisticated spatial sense and it is also reported that producing the scene consumed a large share of the budget.

Many decades later, Wood et. al proposed to create similar hand-made animations with the help of a computer in a reverse engineered fashion [7]. The process starts with the construction of a 3D-scene in a modeling application. Then, a film is captured by a moving virtual camera. The resulting digital movie is played back afterwards. Each image is reduced to a column of pixels in the middle. By concatenating each of these columns next to one another for each frame, the animation is "unrolled" into a panorama. We could also say that the X-axis is exchanges by the time-axis. An artist paints the scene on top of the artificial panorama in greater detail. In the final step, an animation is produced as described for the Pinocchio movie. A panning and a rotating camera can well be generated in the 3D-animation whereas zooms into a scene must be done by zooming into the final drawing made by the artist. Both approaches, the one used by Disney and the one proposed by Wood create an artificial panoramic image with the aim of extracting a realistic video from it.

The complementary method would be to produce a panorama from existing real world images. Among many others, Kim et. al evaluate the generation of multiperspective panoramas from videos showing real scenes [8]. Again, the idea is to reduce every image to a single column of pixels, preferably in the middle of the image. Thus, every frame of a captured video contributes a column of the panorama image which is growing from one side to the other as long as the movie shows a continuous camera operation. The greatest challenge is to move the camera as continuously as possible both in space and time. Even small accelerations result in a warped appearance or complete discontinuities.

Rademacher produces similar panoramas [9]. The difference to prior approaches is that he does not aim at producing a result which can easily be interpreted by a human viewer but which enables the rendering of new perspectives.

Vallance and Calder create raytracing images with continuously varying viewpoints. The position of each pixel on the projection surface serves as parameter of a function which changes the viewpoint slightly. As a consequence, each pixel has an individual center of projection [10]. The benefit of the unnaturally appearing results is that opposing object-surfaces can be seen at the same time.

Today, all approaches based on real images assume a slit

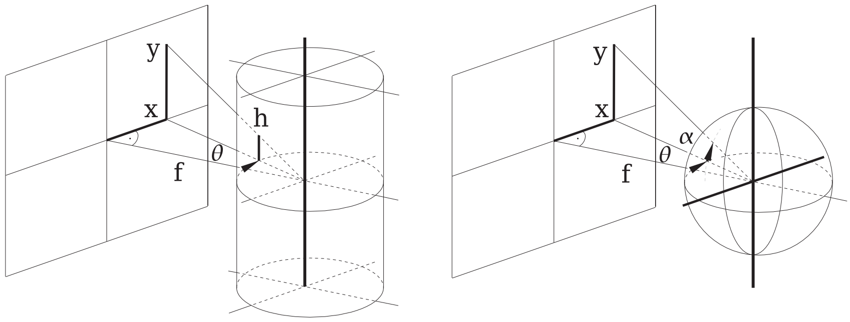

Fig. 2. Stitching panoramic images is done in a common space which can either be tubular or a spherical. The figure exemplifies the mapping of (x, y) coordinates from the image plane to angles and/or lengths.

Fig. 3. Images taken from different perspectives will usually not match, no matter how they are moved or distorted (see semi-transparent images on the left). The middle and right image is cut along a common trajectory. Theoretically, both images would fit together semantically but the layout of the cut allows no concatenation.

camera which produces a sequence of images with only marginally changing viewpoints between successive frames. In the following section we present a new approach to consolidate images with significantly varying viewpoints and viewing directions.

III. Warping for panoramic stitching

In this section we will show that panoramic images are possible, even if the focal point of the camera changes significantly. Of course, the resulting image will imply several changes in perspective and unlike existing approaches, these changes will be continuous by no means. But as we will see, this does not necessarily result in an unnatural output.

Figure 3 shows a building from two different perspectives with a certain overlap. In conventional panoramic image generation, we would overlay both semitransparent images as shown on the left side. By compensating the vanishing points of both images, the overlapping regions could be made congruent. Matching two differing perspectives of the building will fail since moving the center of projection does not only alter vanishing points but also changes the entire content of what can be seen. In an extreme setting, two perspectives could possibly only share a single edge of the building.

Yet, there is a solution to the problem under specific constraints: We have to find a polygonal trajectory (shown as dashed lines in the figure) both in the left and in the right image. When being projected into the real world, both should meet in 3D, theoretically. Vice versa, if a line was drawn onto the objects and surfaces in the real world, it should neither be occluded from the left, nor from the right perspective.

If the images were cut along the trajectories, the right border of the left image would fit to the left border of the neighboring one regarding the semantics of what can be seen. However, concatenating both images is still not possible as the layout of both trajectories is by no means complementary (laying both image next to each other leaves holes). Warping the images to “meet in the middle” could solve the problem. We will now describe a warping approach which tackles the problem.

Figure 4 exemplifies the process. The upper sketch shows the panoramic image which will also be denoted as target image. The lower two images are considered source images. Our warping application iterates over every pixel of the panoramic image in the rendering process. For each pixel, the question has to be answered whether the left or the right source image contributes a color value. If the right one is chosen, the correct source coordinate has to be calculated.

So far, the trajectory was defined for the left and the right source image. A corresponding one has to be defined for the target image as well. In our implementation this can be done manually by the user. Good results were also obtained by simply averaging the left and the right polygonal trajectory and centering the result in the middle of the target image. Whenever the panoramic pixel in question is on the left side of the polygonal trajectory we use the left image, otherwise the right one.

The next question to solve for a given pixel in the panoramic image is: Which is the corresponding pixel in the chosen source image? This is shown in detail in Figure 5. The polygonal trajectory is piecewise linear. Each line segment can be considered as the y-axis of a local coordinate system. We will refer to this axis as *ordinate*. In 2D, a base is entirely defined by calculating the orthogonal x-axis, the *abscissa*. As a result, the left side of the left source image is segmented into distinct areas by those local bases. The same is done with the right side of the right source image and the panoramic target image, accordingly.

The actual mapping works as follows: A coordinate in the (panoramic) target image is expressed by two linear

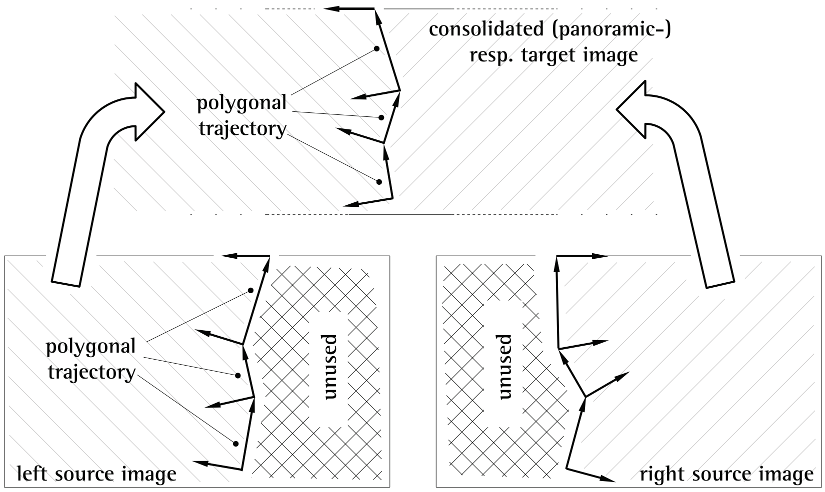

Fig. 4. The sketch shows which parts of the left and the right source image in the lower part of the figure contribute to the rendering of the panoramic image.

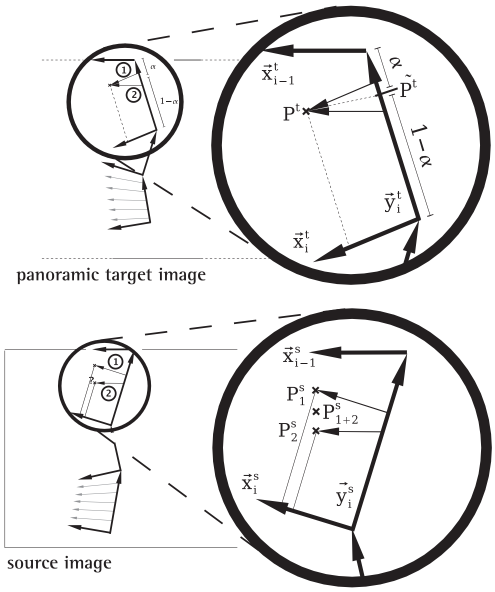

Fig. 5. Each pixel in the panoramic image is linearly combined by two bases, $span(\vec{x}^t_{i-1}, \vec{y}^t_i)$ and $span(\vec{x}^t_i, \vec{y}^t_i)$. The two linear combinations yield different pixels P^s_1 and P^s_2 in the source image as the spanning vectors point into different directions. They have to be merge into a single weighted average P^s_{1+2}.

combinations, one consisting of the base $span(\vec{x}^t_{i-1}, \vec{y}^t_i)$ and one by $span(\vec{x}^t_i, \vec{y}^t_i)$ where ^{t,s} in the exponents stand for *target* and *source*, symbolically. Both linear combinations can be considered as two interpretations of the same location P^t. Both are based on the same ordinate $\vec{y}^t_i$ but they consist of different abscissas, namely $\vec{x}^t_{i-1}$ at the tip of the y-vector and $\vec{x}^t_i$ at y's base.

On the lower side of Figure 5, the two linear combinations are applied to the corresponding spanning vectors in the source image. Since almost all involved vectors point into different directions (as compared to the target image) it is not surprising that the linear combinations do not yield the same coordinate. That is why the resulting points P^s_1 and P^s_2 have to be merged into a single coordinate P^s_{1+2} using an simple convex combination.

$$P^s_{1+2} = (1-\alpha)P^s_1 + \alpha P^s_2 \tag{1}$$

The upper side of the figure indicates, how the value α and $(1-\alpha)$ can be obtained. First, the point P^t in the target image is projected onto the ordinate. The projected point $\tilde{P}^t$ splits the ordinate $\vec{y}^t_i$ into two parts where $\alpha \in [0,1]$ denotes the ratio the split point defines. The effect of the convex weighting scheme is that the abscissa $\vec{X}^s_i$ becomes more important if P^t converges against it. If P^t moves into the opposite direction, the influence of $\vec{X}^s_i$ is diminished in favor of $\vec{X}^s_{i-1}$. The effect of the weighting scheme can be interpreted as a continuous warping from one abscissa to the next. In Figure 5, the resulting intermediate axes are drawn in gray at the lower part of the trajectory. The basic warping scheme described above seems to solve the alignment problems in multi-view panoramic images so we implemented it and ran the software on a couple of image sequences.

IV. Evaluation and Extended Warping

A typical result can be seen in Figure 6 which is obviously not yet a satisfactory outcome. We identified four different kinds of artifacts, namely *expansions*, *contractions*, *undefined areas*, and *mirroring*.

The emergence of expansions and contractions are most obvious. If two neighboring abscissas exhibit a large opening angle in the source image but a smaller angle in the target image this means that a large image patch in the source area will be squeezed into a small patch within the panoramic

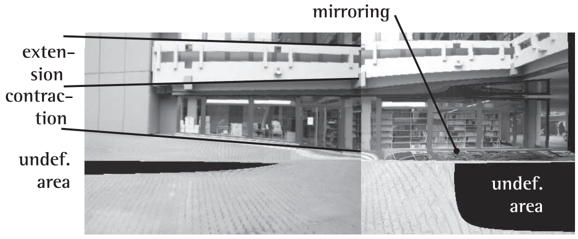

Fig. 6. The naive warping approach described in Section III create four classes of artifacts, namely *expansions*, *contractions*, *mirroring*, and *undefined areas*.

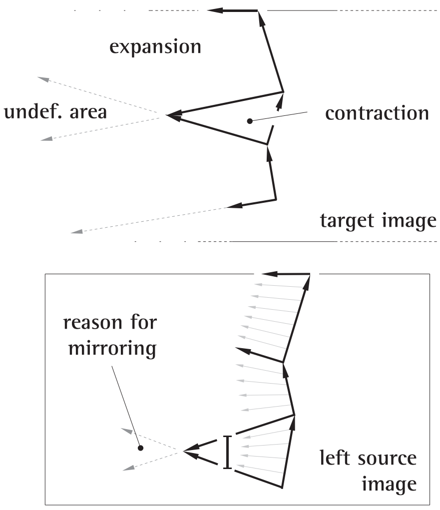

Fig. 7. Expansions originate from the fact that spanning vectors are converging in the source image and diverging in the target image. The complementary case is responsible for contractions. Crossing abscissas in the target image result in undefined areas while they produce mirrored image parts in the source image.

image. This will obviously lead to a contraction. In the opposite case, few source pixels have to contribute to many target pixels which results in an expansion which is even getting worse near the image borders.

The most disturbing artifacts are caused by undefined and mirrored regions. Undefined areas can be found whenever two neighboring abscissas intersect in the target image. Beyond the intersection, the order of the two vectors changes. The one that used to be above its neighbor will be below afterwards and vice versa. In the rendering process, a pixel has to find out which line of the polygonal trajectory (or which ordinate) is responsible for it. This is true for the one with the first abscissa (with a smaller index) being above the pixel and the second (with a higher index) being below. After the intersection this does no more hold true for any of the line segments. Thus, no linear combination can be obtained and as a consequence, no source pixel can be chosen.

Mirrored image patches emerge in the same intersection

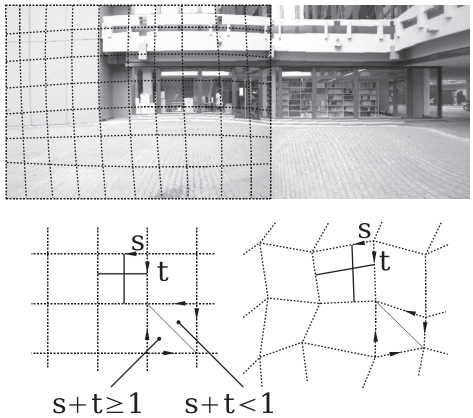

Fig. 9. Even the extended warping scheme suffers from some small artifacts which are emphasized by the dashed lines (upper part). Mapping this distorted grid which is symbolized on the lower right to the one on the left of the figure, diminishes the residual errors.

scenario, however this time the intersection takes place in the source image. The source pixel will be derived as usual, but beyond the intersection, once again, the order of the vectors change which causes horizontal mirroring.

We will now consider how these problems can be overcome. On one hand, the distortion of the source images can not be avoided as this allows to concatenate them in the first place. On the other hand, the distortion becomes increasingly worse towards the borders of the image where they are actually not needed anymore. This led to the idea of using another convex combination in a similar fashion as described above. So far, the warping in the image was done vertically by weighting neighboring abscissas as described in the previous section. Now we will perform the same process in the horizontal direction. Figure 8 depicts the proceeding. Again, we have two points to merge into one by a factor β. The warped linear combination of a point is denoted with A in the figure. Another is labeled with B. B is a trivial linear combination by the image borders (which span the natural coordinate system of the image). We will refer to this as the natural base. Once again, both linear combinations point to the same coordinate in the target image but to different locations in the source image. So the question arises how to weight both coordinates. The weighing factor is denoted with β. Let us consider the scan line through P starting at the left border and ending at the trajectory in the target image. The line is naturally split by P in the ratio $(1-\beta):\beta$. Eventually, a final location P_F is obtained in the source image by

$$P_F = (1-\beta)B + \beta A \tag{2}$$

The more point P converges against the left image border, the more dominant becomes the canonical coordinate B. Otherwise, the warped coordinate A is gaining a larger weight. The result of this weighting scheme is that the zig-zag shape of the trajectory is getting straighter against the left side thus converging against the vertical border of the image.

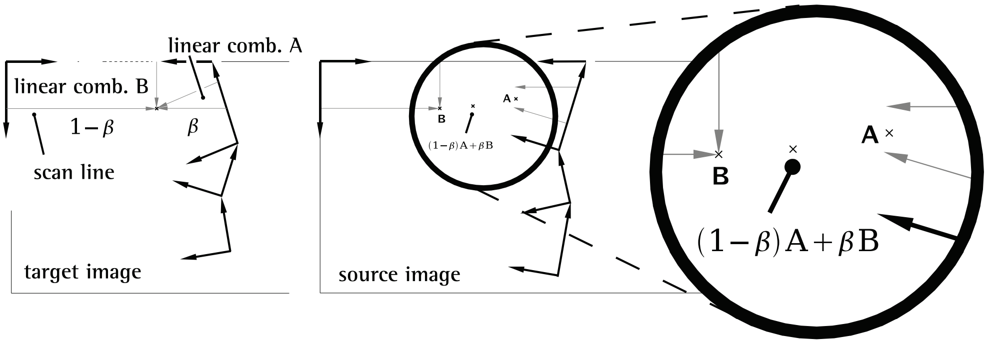

Fig. 8. Here, in the target image every pixel is linearly combined once by the bent coordinate system (A) originating from the trajectory and by a canonical coordinate system (B) spanned by the natural borders of the image. In the source image, these two interpretations can again be weighted by a factor β. Decreasing values for β result in increasingly undistorted image coordinates.

Figure 9 shows an improved version of Figure 6. Most parts of the image appear far more natural but some warping effects remain. They have been outlined with dashed lines on the left side of the figure. The naturalness increases smoothly from the middle of the image towards to borders. The speed of convergence can be adjusted by manipulating β. If the parameter is e.g. squared, the warped image converges much faster against the natural one. Therefor, a stronger local bending takes place near the trajectory. Depending on the image content, this can at times be more disturbing than distributing the distortion over a larger area in the image.

Fortunately, the remaining artifacts can be compensated easily. In Figure 9, a regular grid is shown in the lower left. The underlying image will now be calculated in a similar fashion like done before in the panoramic image. Each pixel is contained in a unique grid cell. Thus, it can be linearly combined by means of the spanning cell boundaries which results in two scalars s and t. The scalars are then applied to the boundaries of the corresponding cell of the distorted grid shown on the lower right of the figure. The color value at the resulting coordinate serves the original pixel in the undistorted image as input. Whenever $s+t<1$ holds true, the pixel is in the upper triangular half of the cell. Otherwise, a new linear combination has to be obtained by the lower and the left cell boundary, again.

The intentional bending of the image is only one reason for its waviness. Another reason is the distortion caused by the lens which results in a fish eye-like appearance of the left and the right part[1].

Figure 10 and 11 show examples of panoramic images with changing perspectives. In our evaluations we came across the limitations of the approach which were less dependent on the implementation but on natural image consistency. We will go into detail on that in the next chapter.

V. Panoramic video-artifacts

The warping and stitching of images works best if the object being shown can be unfolded theoretically. This is particularly true for buildings. More generally, stitching works without problems if the camera follows a straight path. The reason is that succeeding images can be concatenated to existing ones if the above mentioned trajectory exists which can be seen unoccludedly from two neighboring viewpoints. We can even generalize feasible scenarios further by assuming that a camera must never be able to see the location of another camera. This situation is shown in Figure 12 on the left side. Here, the same objects appear multiple times from differing perspectives, according to the number of cameras. There is no way to consolidate different views of a transparent or at least none-solid object like a tree into a single view because the tree can not be “unfolded” in the same way like a building. In the latter case, the building itself prevents opposite cameras from viewing each other or from viewing the same impression more than once. Figure 11 shows an example produced by four cameras which do not lay on a common line and which do not view into the same direction. Yet, the unwarped sculpture looks credible as each camera sees the solid object and none of the cameras face each other. The fact that the sculpture could theoretically be flattened into a plane is a sufficient precondition for producing a panoramic view.

In the middle of Figure 12, two distant cameras look at a

[1]The images were taken with a focal length of 18mm

Fig. 10. The above view of the building can not be accomplished using classical panorama techniques since the opposite building prevents the camera from gaining enough distance. Despite the fact that the viewpoint changes four times, the image still appears credible. Notice the two shadows of the same street lamp marked with a black circle. It is an instance of a number of inevitable artifacts described in Section V.

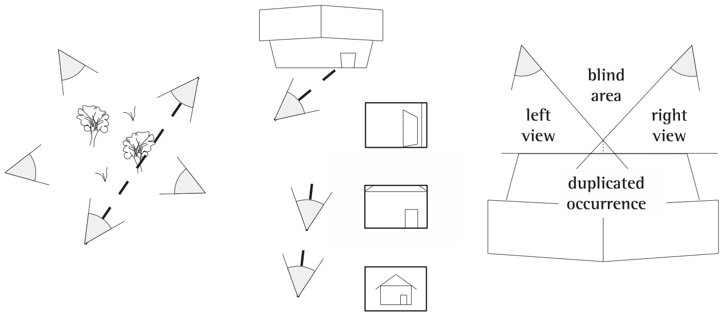

Fig. 12. Different perspectives of a single translucent object can not be merged into a common view (left). In the middle, the uppermost view into the door shows another scenario than the one which can be seen by the frontal-facing cameras. Even two cameras facing different parts of a flat object produce blind areas or areas of duplicated appearance.

Fig. 11. The curved sculpture was unrolled from several camera perspectives with changing locations and viewing directions.

house. Merging the more focused view into the wide angle view would be straight forward. It would simply result in a larger picture with differing resolutions throughout the image. However, the camera with the least distance to the house shows the door and parts of the interior of the inner right side. Here, no easy way can be found to merge the *new* content into a global image, as the interior would in any case have to overwrite existing parts of the global picture.

The above mentioned artifacts apply to still images and moving ones in the same way. An effect which would be more obvious in moving images can be exemplified on the right hand side of Figure 12. Here, two cameras face the front of a house. Their views are cut and merged to a continuous panoramic image like, e.g., the one shown in Figure 10.

If a person was walking from the left to the right between the building and the cameras, she would first appear in the viewing volume of the left camera. After a while, she would walk inside the blind area. Within the panoramic video, this would mean that the person disappears behind an certain column of pixels without any reason since the house itself appears to be continuous for the viewer. After a while she would enter the viewing volume of the right camera which means that the person would suddenly appear after the same column of pixels mentioned before. For a human observer there is no obvious reason for the disappearance as no obstacle can be seen. Yet, the effect is a natural consequence of merging two differing viewpoints.

Another equally wired effect is the duplication of entities beyond intersecting viewing volumes. This is denoted as *duplicated occurrence* in the figure. The panorama of the building was designed such that its front appears continuous. Likewise, the screens of the building in Figure 10 have preserved their natural size and aspect ratio. However, objects behind the screen can be seen twice next to the intersecting line of the left and right image. Everything behind the intersection point of the boundaries of the viewing volumes can both be seen from the left and the right camera. So by examining Figure 10 in more detail the reader may have noticed that the shadow of the street lamp is projected twice onto the white blinds behind the screen. Theoretically, a person walking behind the screen would appear natural in the beginning and then suddenly be duplicated in the area of *duplicated occurrence* before becoming singular again.

We can conclude that strict optical consistency can only be preserved within a planar surface before the camera.

VI. Conclusion and Outlook

A new approach for producing panoramic images from photos with varying centers of projection was proposed. A trajectory has to be found in neighboring images which serves as cut. A vertical warping scheme distorts neighboring images such that they fit together. Another horizontal mapping makes the distorted parts of the image near the cut converge against the natural image near the left and right border of the panorama. Finally, remaining artifacts originating from the warping and lens distortion are corrected.

In Figure 9 the image was dewarped to compensate the distortion. As we have seen before, the warping was done according to the piecewise defined local bases which originated from the trajectory through the image. If the entire image is overlaid with the vectors spanning those coordinate systems, then a continuous vector field will be visible. We can interpret this vector field as the derivation of the overlaid (primitive) dewarping function shown as dashed lines in the figure. As a consequence, this dewarping function does not necessarily have to be defined by the user but it could as well be obtained numerically. This would not compensate for the lens distortion. But since the latter is a constant function of the lens and the chosen focal length, it could be considered additionally with no dependence on the image content.

REFERENCES

[1] B. Wilburn, N. Joshi, V. Vaish, E. Talvala, E. Antunez, A. Barth, A. Adams, M. Levoy, and M. Horowitz, "High performance imaging using large camera arrays," in *Proc. SIGGRAPH*, 2005.

[2] B. Wilburn, N. Joshi, V. Vaish, M. Levoy, and M. Horowitz, "High speed video using a dense camera array," in *Proc. of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.

[3] V. Vaish, B. Wilburn, N. Joshi, and M. Levoy, "Using plane and parallax for calibrating dense camera arrays," in *Proc. CVPR 2004*, (Washington, DC, USA), June 2004.

[4] B. Wilburn, M. Smulski, H. K. Lee, and M. Horowitz, "The light field video camera," in *Proc. Media Processors 2002, SPIE Electronic Imaging 2002*, (San Jose (CA), USA), January 2002.

[5] G. N. Barnard, *Barnard's Photographic Views of the Sherman Campaign*. New York, NY, USA: Press of Wynkoop and Hallenbeck, 1866.

[6] D. Farin, *Automatic Video Segmentation Employing Object/Camera Modeling Techniques*. Eindhoven, Netherlands: CIP-Data Library Technische Universiteit Eindhoven, 2005. pp. 423–425.

[7] D. N. Wood, A. Finkelstein, J. F. Hughes, C. E. Thayer, and D. H. Salesin, "Multiperspecive panoramas for cel animation," *Proc. SIGGRAPH*, 1997.

[8] J. Kim, K. Henning, S. Seitz, and D. Salesin, "Multiperspective images from videos," tech. rep., Graphics and Imaging Laboratory, University of Washington, 2001. http://grail.cs.washington.edu/projects/videoslice/results.

[9] P. Rademacher and G. Bishop, "Multiple-center-of-projection images," *Proc. SIGGRAPH*, 1998.

[10] S. Vallance and P. Calder, "Multi-perspective images for visualization," in *Pan-Sydney Area Workshop on Visualization Information Processing*, 2002.