

360 × 360 Mosaics*

Shree K. Nayar and Amruta Karmarkar

Department of Computer Science
Columbia University, New York
{nayar, amruta}@cs.columbia.edu

Abstract

Current mosaicing methods use narrow field of view cameras to acquire image data. This poses problems when computing a complete spherical mosaic. First, a large number of images are needed to capture a sphere. Second, errors in mosaicing make it difficult to complete the spherical mosaic without seams. Third, with a hand-held camera it is hard for the user to ensure complete coverage of the sphere. This paper presents two approaches to spherical mosaicing. The first is to rotate a 360 degree camera about a single axis to capture a sequence of 360 degree strips. The unknown rotations between the strips are estimated and the strips are blended together to obtain a spherical mosaic. The second approach seeks to significantly enhance the resolution of the computed mosaic by capturing 360 degree slices rather than strips. A variety of slice cameras are proposed that map a thin 360 degree sheet of rays onto a large image area. This results in the capture of high resolution slices despite the use of a low resolution video camera. A slice camera is rotated using a motorized turntable to obtain regular as well as stereoscopic spherical mosaics.

1 Spherical Mosaics

A mosaic is constructed by stitching¹ together multiple images, where the individual images correspond to different views of the scene captured from approximately the same viewpoint. Several methods for image mosaicing have been proposed (for examples, see [Burt and Adelson, 1983], [Mann and Picard, 1994], [Zheng and Tsuji, 1992], [Chen, 1995], [Irani *et al.*, 1995], [Kang and Szeliski, 1996] [Peleg and Herman, 1997], [Rouso *et al.*, 1997], [Szeliski, 1996], [Sawhney *et al.*, 1995], [Krishnan and Ahuja, 1996], [McMillan and Bishop, 1995], [Szeliski and Shum, 1997]). These techniques use a conventional imaging lens to capture the image sequence. Since such lenses have limited fields of view, the computation of a complete spherical mosaic requires the capture and processing of a large number of images. In addition, errors in the image projection model and errors in the estimation of motion between images makes it difficult to complete the sphere

*This work was supported in parts by DARPA's Image Understanding Program and an ONR/DARPA MURI grant under ONR contract No. N00014-97-1-0553.

¹In our definition of mosaicing, we will include both image based as well as slit (a slice through the image) based techniques. In the case of slits, the slices are not really stitched but rather concatenated together to form the mosaic.

without undesirable seams in the final mosaic. Further, in the case of a hand-held camera, it is hard for the user to ensure that the complete sphere has been scanned during the capture process.

An alternative approach is to use a wide-angle imaging system such as a fish-eye lens (see [Kuban *et al.*, 1994], [Xiong and Turkowski, 1997]) or a catadioptric imaging system (see [Nayar, 1997], [Yagi, 1999] for surveys). In both cases, a hemispherical field of view can be captured within a single image. Hence, a small number of such images can be stitched together to obtain a spherical mosaic². However, this approach typically results in inadequate resolution due to the inherent trade-off between field of view and image resolution; as the field of view increases, the resolution decreases, causing the computed spherical image to be of lower quality than in the case of a conventional imaging system.

This paper presents two efficient approaches for capturing high resolution spherical mosaics. In the first approach, a wide-angle imaging system is used to capture a sequence of 360 degree strips on the sphere by a single rotation of the capture device. For this, we suggest the use of a catadioptric imaging system since such a system typically produces higher resolution in the periphery of the hemispherical field of view than a fish-eye lens. The unknown rotations between the strips are estimated and used to blend the multiple strips into a single spherical mosaic. Our second approach seeks to further enhance the resolution of the computed mosaic. This is done by designing new catadioptric sensors that capture a single 360 degree slice of the scene³. Mirror shapes are derived that enable the projection of a thin slice onto a large image area. This results in the capture of high resolution slices despite the use of a low resolution (640x480 pixel) image detector. Such a slice camera is rotated on a turntable and the captured slices are concatenated to obtain a high resolution spherical mosaic. Though a large number of images (slices) are needed to obtain a high resolution mosaic, the processing of each image is minimal and is easily done in real time.

Recently, several investigators have explored the capture of stereoscopic panoramas. Ishiguro *et al.* [Ishiguro *et al.*, 1992] were the first to use stereo panoramas for computing structure. Then, Huang and Hung [Huang and

²See [IPIX, 1999] for results on the stitching of two fish-eye images to obtain a spherical mosaic.

³In [PanoScan, 1999], a 180 degree slice is captured by using a fish-eye lens and a high resolution line detector.

Hung, 1998] used a rotating stereo head to show that two panoramic images are sufficient to generate stereo views for any direction within the panoramas. Subsequently, Peleg and Ben-Ezra [Peleg and Ben-Ezra, 1999] showed that the rotation of a single camera provides all the information needed to obtain a stereo panorama. More recently, Shum et al. [Shum *et al.*, 1999] extended these ideas to capture omnivergent stereo data, using the rotation of a camera. Shum et al. also showed synthetic examples of spherical stereo mosaics but did not present techniques for obtaining such data in practice. We show that our strip and slice cameras can be used to easily capture stereoscopic spherical mosaics by displacing the viewpoint of the imaging system from the axis of rotation. We conclude with examples of high resolution stereoscopic spherical mosaics, that enable a user to freely pan and tilt while perceiving the depths of objects in the scene.

2 360 Degree Strips

A spherical mosaic can be represented in several ways. Without loss of generality, we will use the spherical panorama shown in Figure 1(a) as our representation of choice. This representation is convenient as it is linear in the polar angle θ and the azimuth angle ϕ . As shown in Figure 1(a), if we use only half the strip (180 degrees), a single 360 degree rotation along the azimuth angle ϕ is sufficient to cover the entire sphere. If the strip is cylindrical, it maps to a bow-shaped band in the spherical panorama. Alternatively, a 180 degree rotation of the sensor is sufficient if full 360 degree strips are used.

A 360 degree strip can be captured using a fish-eye lens with a field of view slightly greater than a hemisphere, or a catadioptric sensor that uses a curved mirror and an imaging lens. As shown in [Baker and Nayar, 1998], parabolic, hyperbolic and ellipsoidal mirrors will produce strips while maintaining a single viewpoint. Catadioptric sensors have a clear advantage over fish-eye lenses here, since they tend to have greater resolution in the periphery of the field of view (and hence, within the strip). It was shown in [Baker and Nayar, 1998] that if a conic mirror with profile $z(r)$ is used with a perspective lens located at $z = c$, the resolution of the catadioptric system is:

$$\frac{dA}{d\Omega} = \frac{r^2 + z^2}{(c - z)^2 + r^2} \frac{dA}{d\omega} \quad (1)$$

where, dA is the area of a single pixel on the image detector, $d\omega$ is the solid angle subtended by the pixel through the imaging lens and $d\Omega$ is the solid angle subtended by the pixel after reflection by the curved mirror. In the case of a parabolic system, the resolution perpendicular to the optical axis (in the middle of the strip) is approximately 4 times the resolution in the direction of the optical axis. This increase in resolution (with respect to a fish-eye lens) is of course only in the ϕ dimension of the mosaic, as the maximum achievable resolution along θ is determined by the number of pixels on the image detector and is similar for all 360 degree imaging systems.

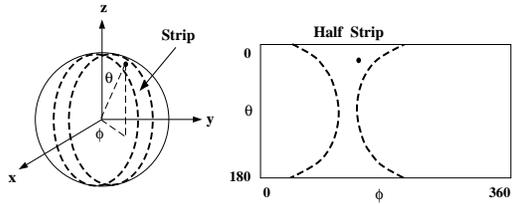
Let us consider the case where the sensor is freely rotated

by a human and not a controlled turntable. In this case, the rotations between consecutive frames are unknown and need to be estimated, so that all captured strips can be mapped to the coordinate frame of the spherical mosaic. To this end, the rotation matrices $\mathbf{R}_{k-1,k}$ between consecutive images k and $k-1$ are computed using a set of corresponding features, that are found using a feature tracking algorithm. We can assume that the frame of reference of the mosaic is defined with respect to the initial orientation ($k = 0$) of the sensor. Then, each strip is mapped to the mosaic by using the rotation matrix:

$$\mathbf{R}_k = \mathbf{R}_{0,1}\mathbf{R}_{1,2}\dots\mathbf{R}_{k-1,k} \quad (2)$$

When strips are mapped to the spherical panorama, they are expected to overlap with previously accumulated data. Two steps are taken to ensure a seamless mosaic. First, we note that the computed rotations between strips are expected to include small errors. Therefore, the computed rotation for a strip is only used as an initial estimate of the location of the strip with respect to the mosaic. The Euler angles $(\alpha_k, \beta_k, \gamma_k)$ corresponding to the rotation matrix \mathbf{R}_k are then varied within a small search range to find the strip location on the mosaic that minimizes the sum-of-squared difference in brightness between the strip and the mosaic. Once the rotation matrix has been refined in this manner, a blending process is used to merge the strip data with the mosaic. During blending, the brightness of a point in the region of overlap between the mosaic and the strip is computed as a weighted sum of its brightnesses in the mosaic and the strip. The weights are proportional to the shortest distances of the point from the boundaries of the mosaic and the strip.

Figure 1(b) shows the catadioptric imaging system we have used for strip mosaicing. This system was described in [Nayar, 1997] and includes a parabolic mirror and a telecentric lens. This optics is attached to a Canon Optura video camera. The imaging system is mounted on a tripod and rotated by hand, and the captured video is processed using the above mosaicing algorithm. Since the sensor was rotated by 360 degrees, only half the strip (180 degrees) was used for the mosaic computation. In our experiment, a strip width of 30 degrees was used. The computed spherical mosaic (4000x2000 pixels in size) is shown in Figure 1(c). Using one of several image rendering softwares, one can create perspective images and freely navigate around the spherical field of view. Three examples of computed perspective images are shown in Figure 1(d), which reveal the resolution of the computed mosaic. While the resolution is reasonable for objects at short distances from the imaging system, it is not sufficient for distant object. Several factors contribute to the lack of resolution. First, the alignment of a strip with the mosaic is never exact and therefore the blending process low-pass filters the mosaic. More importantly, the video camera has only 640x480 pixels. Hence, though there are 2000 pixels in the θ dimension of the mosaic, these pixels are interpolated from just $\pi 240 = 753$ pixels (measurements). We will now present ways to enhance mosaic resolution, without increasing the resolution of our image detector.



(a) 360 degree strip.



(b) Rotation of 360 degree camera.



(c) Complete spherical mosaic.



(d) Perspective views.

Figure 1: (a) A single rotation of a 360 degree camera results in strips on the sphere that together cover the entire sphere. Half of a strip (180 degrees) maps to a bow-shaped band in the spherical panorama. The unknown rotations between the strips are estimated and the strips are blended together to obtain a seamless mosaic. (b) A 360 degree camera can be rotated on a tripod or rotated by hand, to capture the strips. For this experiment, the sensor described in [Nayar, 1997] that includes a telecentric lens and a parabolic mirror was used and a strip width of 30 degrees was chosen. The unknown rotations between strips are computed from corresponding image features. (c) The computed rotations are used to blend the strips together into a single spherical mosaic, shown here as a 4000x2000 pixel spherical panorama, identical to the one illustrated in (a). (d) An interactive viewer is used to generate perspective views from the spherical mosaic. (See [CAVE, 2000] for color figures).

3 360 Degree Slices

We now introduce a new class of sensors that can image a thin 360 degree slice with higher resolution. Before we proceed to present the problem of slice imaging, it is worth comparing the use of slices with the use of strips. The advantage of strips is that for smooth rotation of the sensor, an overlap between successive frames is guaranteed and can be used for computing the rotation between frames. However, as we have seen, the resulting mosaic is limited in resolution. In the case of slices higher resolution is achieved. However, since there is no overlap between slices, the rotation of the sensor must be known. As illustrated in Figure 2, a 360 degree slice corresponds to a greater circle on the sphere and half the slice maps to a vertical line in the spherical panorama.

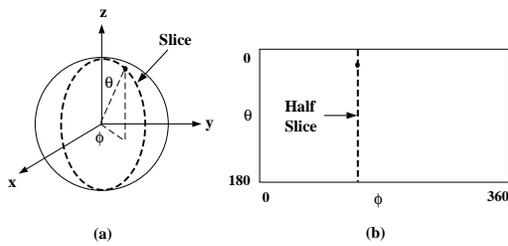


Figure 2: A single rotation of a 360 degree slice sensor produces a higher resolution spherical mosaic. (a) Each slice corresponds to a greater circle on the sphere. (b) Half the slice maps to a line in the spherical panorama.

As shown in Figure 3(a), a slice camera is rotated using a motorized turntable. The turntable is spun at constant velocity as the images are captured, keeping the rotation between consecutive slices constant. As noted in [Peleg and Ben-Ezra, 1999] and [Shum *et al.*, 1999], if the center of projection of the image sensor is displaced from the axis of rotation, a stereoscopic panorama can be captured. In the case of spherical mosaics, complete 360 degree slices must be captured during rotation. If this is feasible, one half (180 degree field of view) of the slice is used to construct a left-view mosaic while the other half can be used to obtain a right-view mosaic. Clearly, we could also use a 360 degree strip sensor to accomplish this. However, as we have seen, the resolution is not of sufficient quality and is expected to degrade when the strip sensor is placed off-axis since this introduces a translation in addition to the rotation between frames. Both these problems can be overcome using a slice sensor.

We define a slice camera as one that projects a very thin 360 degree sheet of rays onto a large number of image pixels as shown in Figure 4(a). While a fish-eye lens or a wide-angle catadioptric sensor projects a slice of the scene onto a circle (say, of radius R_1), the slice camera projects the same slice onto a wide disc (of thickness $R_1 - R_2$). Since this disc is imaged using a rectangular grid of pixels, additional measurements along θ are obtained within the disc. Therefore, in place of the small number of mea-

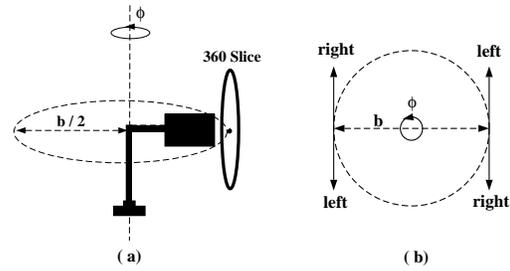


Figure 3: (a) Rotation of a slice camera with the center of projection away from the axis of rotation results in a stereoscopic spherical mosaic. (b) The two halves of the 360 degree slice are used to construct left and right spherical panoramas such that all scene points are seen with approximately the same baseline b . The idea of off-axis rotation is described in detail in [Peleg and Ben-Ezra, 1999] and [Shum *et al.*, 1999].

surements on a circle shown in Figure 4(b), we obtain a larger number of measurements as shown in Figure 4(c). These measurements are easily interpolated to obtain a dense, uniform sampling of the brightness function $I(\theta)$ along the slice. It may appear that the resolution in this case is proportional to the width of the image disc. Due to finite (non-zero) pixel size this is not the case. A detailed analysis of exactly how resolution varies with disc width is beyond the scope of this paper. For now, we will simply note that a significant improvement in resolution is achievable.

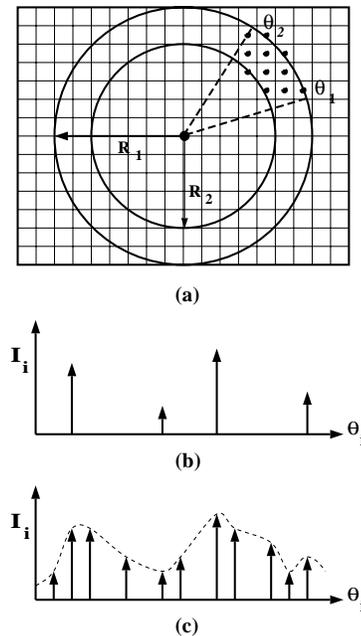


Figure 4: (a) A slice camera projects a thin 360 degree sheet of parallel rays onto a large disc in the image. The number of measurements within the sheet increases from (b) pixels on a circle to (c) pixels within a disc. These measurements can be interpolated to obtain a uniformly sampled high resolution slice.

4 Slice Cameras

We are now left with the problem of designing a slice sensor. Once again, catadioptric imaging provides the flexibility needed for such a design. As shown in Figure 5, the problem can be formulated as one of deriving the mirror shape that reflects only a set of parallel incoming rays towards the center of projection (pinhole) O of the imaging lens. Since this is a rotational symmetric imaging system, the problem is reduced to finding the profile $z(r)$ of the mirror. Let the pinhole of the perspective imaging lens be at $z = c$. For any point on the mirror, we denote the angle made by the normal with respect to the optical (vertical) axis by β and the angle made by the reflected ray with the horizontal axis by α . Here, we will assume that the slice is perpendicular to the optical axis, and hence $\beta + \gamma = \pi/2$. However, the formulation is general in that $\beta + \gamma$ can be set to any other angle, which would result in the imaging of a conical sheet of rays with its axis aligned with the optical axis, rather than a flat sheet.

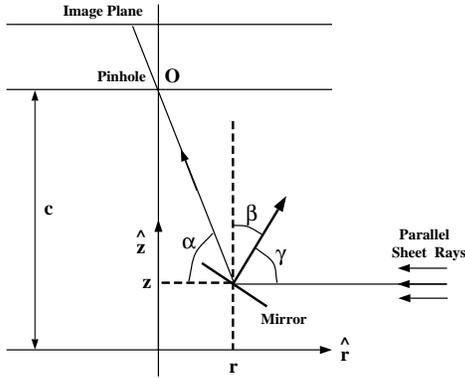


Figure 5: The mirror shape that produces a compact 360 degree slice is defined as one that reflects a thin sheet of parallel rays through the effective pinhole of the imaging lens.

Note that the reflecting point (z, r) is related to the angle of reflection α as:

$$\tan \alpha = \frac{c - z}{r}. \quad (3)$$

Since the incoming ray is specularly reflected, it is easy to show that

$$\alpha = 2\beta. \quad (4)$$

Also, we know that the slope of the mirror at the point of reflection is related to the angle β of the normal as

$$-\tan \beta = \frac{dz}{dr}. \quad (5)$$

Using the above expressions in the well-known identity

$$\tan 2\beta = \frac{2 \tan \beta}{1 - \tan^2 \beta}, \quad (6)$$

we get the first-order quadratic differential equation:

$$\frac{-2 \frac{dz}{dr}}{1 - \frac{dz}{dr}^2} = \frac{c - z}{r}. \quad (7)$$

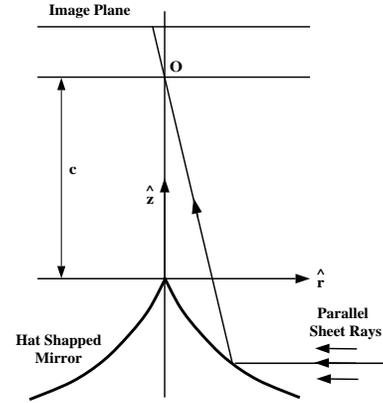


Figure 6: For perspective projection, the mirror shape that images a thin 360 degree sheet of parallel rays is hat-shaped and given by equation (8).

This equation is solved to obtain the mirror profile

$$z = c - 2\sqrt{k} \sqrt{|r| + k}, \quad (8)$$

where k is the constant of integration. If we set $z = 0$ at $r = 0$, we get $k = c/2$. The resulting mirror is shown in Figure 6 and has a hat-like shape. An even simpler mirror shape is obtained by using orthographic image projection. In this case, $\beta = \pi/2$ and hence equation 5 reduces to $dz/dr = -1$. The resulting mirror is a cone with a 90 degree angle at the apex, as shown in Figure 7. In [Nayar, 1997], orthographic projection was used with a parabolic mirror and in [Yagi and Yachida, 1991] a cone was used with perspective lens. In both cases, a wide field of view was sought in both dimensions (ϕ and θ). In our case, the combination of orthographic projection and a conical mirror results in the desired image projection model for parallel rays. It is worth noting that the thickness Δz of the sheet of rays can be made arbitrarily small. For instance, if a cone with a 1 cm outer diameter is projected onto a 500x500 pixel detector, a 0.5 mm thick sheet of rays is projected onto an image disc that is 25 pixels wide.

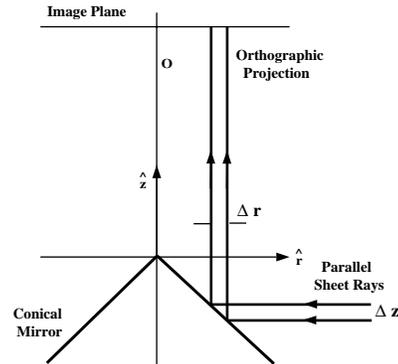


Figure 7: For orthographic projection, a conical mirror is used to image a 360 degree sheet of parallel rays. The thickness Δz of the sheet that maps to a given image disc can be made arbitrarily small by increasing the magnification of the orthographic lens and reducing the size of the cone.

5 Experimental Results

We have implemented the slice camera illustrated in Figure 7. A telecentric lens is used to ensure orthographic projection and a cone with a 65 mm outer diameter is attached to the lens using a transparent acrylic cylinder. This optical system is attached to a Canon Optura digital video camera, as shown in Figure 8. The complete imaging system is rotated using a Daedal motorized turntable. The center of projection of the imaging system is offset from the axis of rotation by 3 inches to get a stereo baseline of 6 inches. In most of our experiments, a complete 360 degree rotation of the sensor was done in 3 minutes. This results in a sequence of approximately 5000 images.

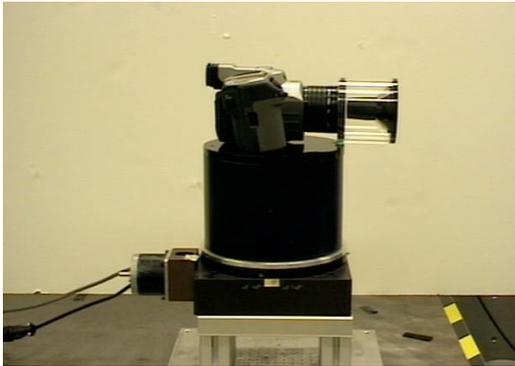


Figure 8: A slice imaging system based on the design shown in Figure 7. A telecentric lens and a conical mirror are attached to each other using a transparent acrylic tube. This optical system is attached to a Canon Optura video camera. The complete imaging system is rotated (off-axis for stereo) on a motorized turntable.

An example of an image produced by the slice camera is shown in Figure 9. As expected, scene features are mapped to radial strips in the image. A 4-pixel wide disc was used to obtain approximately 3000 color measurements within a 180 degree (half) slice. These measurements are interpolated to obtain 2000 uniformly distributed samples. The right half and the left half of the slice were mapped to vertical lines in the left and right spherical panoramas, respectively. Figures 11(a) and (b) show the left and right spherical mosaics computed for an indoor scene. Each mosaic is 4000x2000 pixels in size. It is easy to see that all scene points produce disparity in the left and right mosaics, the disparity varying with distance from the imaging system. Figures 11(c) and (d) show stereo pairs of perspective views for two regions of the scene. We have used a perspective viewer to enable a user to freely roam around the captured sphere, while perceiving the depths of objects using red-green eye glasses.

The resolution of these mosaics is significantly greater than the one in Figure 1 generated using a strip camera. This increase in resolution is not obvious since all mosaics have been scaled down in size prior to printing on paper. The resolution advantage of the slice camera is illustrated in Figure 10, where a step edge in the scene is reconstructed from the slices using image discs of 1 pixel and 4 pixel widths. In both cases, the same interpolation algorithm (a

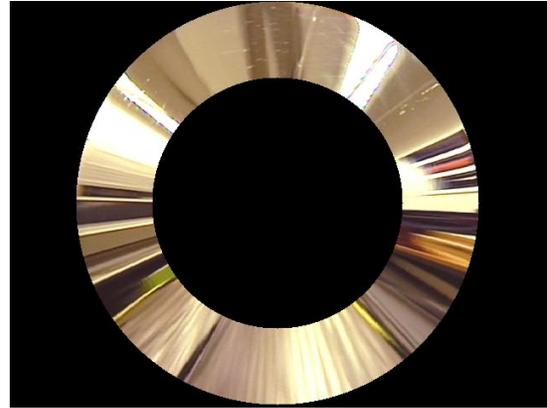


Figure 9: Example of an image produced by the slice camera shown in Figure 8. As expected, scene features are projected to radial strips in the image.

Gaussian filter with σ equal to 0.001 radians) was used. It is clear that the 4-pixel disc provides higher resolution than the 1-pixel disc.



Figure 10: Magnified images of a step edge in the scene reconstructed using (a) an image disc of 1 pixel width and (b) an image disc of 4 pixels width. Note the effects of aliasing in (a).

6 Variants of the Slice Camera

We conclude by mentioning a few variants of the slice camera shown in Figure 8. As always, there exists an inherent trade-off between slice resolution and slice field of view. If a greater resolution is desired at the cost of field of view, the magnification of the imaging lens can be increased to image smaller sections of the cone. In Figure 12, two 90 degree sections of the cone are configured to fill the larger dimension (640 pixels in this case) of the image detector. Consequently, a further increase in resolution by a factor of two is achieved. Finally, as with wide-angle catadioptric systems, multiple mirrors can be used to reduce sensor size while possibly improving image quality (see [Nayar and Peri, 1999]). For instance, a perspective lens and a concave parabolic mirror can be used to achieve orthographic projection of the conical mirror. Similarly, an ellipsoidal or hyperboloidal mirror can be used with a perspective lens to image the hat-shaped mirror. In this case, the near focus of the ellipsoidal or hyperboloidal mirror serves as the pinhole for the projection of the hat-shaped mirror.



(a) Left spherical mosaic.



(b) Right spherical mosaic.



(c) Left and right perspective views.



(d) Left and right perspective views.

Figure 11: The slice camera shown in Figure 8 was rotated with the center of projection 3 inches away from the axis of rotation to obtain the left and right spherical mosaics shown in (a) and (b). Note the horizontal disparity between the image coordinates of scene points in the left and right mosaics. Each mosaic is 4000x2000 pixels in size. (c) and (d) show left and right (stereo) perspective views for two different parts of the scene. (See [CAVE, 2000] for color figures).



Figure 12: Two 90 degree sections of a cone are projected on to the larger dimension of the rectangular image. This results in a panoramic slice camera with a two-fold increase in slice resolution.

7 Summary

In this paper we have presented several results on the capture of regular as well as stereoscopic spherical mosaics. We described two methods, one based on strips and other based on slices, to capture all the required scene information with a single rotation of the sensor. In addition, we derived a class of catadioptric slice cameras that project a thin sheet of parallel rays onto a wide disc in the image. The additional measurements provided by such a sensor were used to construct high resolution stereoscopic spherical mosaics. We are currently in the process of developing compact and portable slice cameras.

References

[Baker and Nayar, 1998] S. Baker and S. K. Nayar. Catadioptric Image Formation. *Proc. of International Conference on Computer Vision*, January 1998.

[Burt and Adelson, 1983] P. J. Burt and E. H. Adelson. A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics*, 2(4):217–236, October 1983.

[CAVE, 2000] S. K. Nayar and A. Karmarkar. 360 × 360 Mosaics. <http://www.cs.columbia.edu/CAVE/>, April 2000.

[Chen, 1995] S. E. Chen. QuickTime VR - An Image Based Approach to Virtual Environment Navigation. *Computer Graphics: Proc. of SIGGRAPH 95*, pages 29–38, August 1995.

[Huang and Hung, 1998] H. C. Huang and Y. P. Hung. Panoramic stereo imaging system with automatic disparity warping and seaming. *Graphical Models and Image Processing*, 60(3):196–208, May 1998.

[IPIX, 1999] IPIX. <http://www.ipix.com/>. 1999.

[Irani et al., 1995] M. Irani, P. Anandan, and S. Hsu. Mosaic Based Representations of Video Sequences and their Applications. *Proc. of Fifth International Conference on Computer Vision*, pages 605–611, June 1995.

[Ishiguro et al., 1992] H. Ishiguro, H. Yamamoto, and S. Tsuji. Omnidirectional Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:257–262, 1992.

[Kang and Szeliski, 1996] S. Kang and R. Szeliski. 3-D Scene Data Recovery Using Omnidirectional Multibaseline Stereo. *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pages 364–370, June 1996.

[Krishnan and Ahuja, 1996] A. Krishnan and N. Ahuja. Panoramic Image Acquisition. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR-96)*, pages 379–384, June 1996.

[Kuban et al., 1994] D. P. Kuban, H. L. Martin, S. D. Zimmermann, and N. Busico. Omniview Motionless Camera Surveillance System. *United States Patent*, (5,359,363), October 1994.

[Mann and Picard, 1994] S. Mann and R. Picard. Virtual Bellows: Constructing high quality stills from video. *IEEE International Conference on Image Processing*, pages 363–367, November 1994.

[McMillan and Bishop, 1995] L. McMillan and G. Bishop. Plenoptic Modeling: An Image-Based Rendering System. *Computer Graphics: Proc. of SIGGRAPH 95*, pages 39–46, August 1995.

[Nayar and Peri, 1999] S. K. Nayar and V. Peri. Folded Catadioptric Cameras. *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, June 1999.

[Nayar, 1997] S. K. Nayar. Catadioptric Omnidirectional Camera. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, June 1997.

[PanoScan, 1999] PanoScan. <http://www.panoscan.com/>. 1999.

[Peleg and Ben-Ezra, 1999] S. Peleg and M. Ben-Ezra. Stereo Panorama with a Single Camera. *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pages 395–401, June 1999.

[Peleg and Herman, 1997] S. Peleg and J. Herman. Panoramic Mosaics by Manifold Projection. *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pages 338–343, June 1997.

[Rousso et al., 1997] B. Rousso, S. Peleg, and I. Finci. Mosaicing with Generalized Strips. *Proc. of DARPA Image Understanding Workshop*, pages 255–260, May 1997.

[Sawhney et al., 1995] H. Sawhney, S. Ayer, and M. Gorkani. Model-Based 2D and 3D Dominant Motion Estimation for Mosaicing and Video Representation. *Proc. of Fifth International Conference on Computer Vision*, pages 583–590, June 1995.

[Shum et al., 1999] H-Y. Shum, A. Kalai, and S. M. Seitz. Omnivergent stereo. *Proc. of Seventh International Conference on Computer Vision*, pages 22–29, September 1999.

[Szeliski and Shum, 1997] R. Szeliski and H. Y. Shum. Creating Full View Panoramic Image Mosaics and Environment Maps. *Proc. of ACM SIGGRAPH'97*, pages 251–258, August 1997.

[Szeliski, 1996] R. Szeliski. Video Mosaics for Virtual Environments. *IEEE Computer Graphics and Applications*, 16(2):22–30, 1996.

[Xiong and Turkowski, 1997] Y. Xiong and K. Turkowski. Creating Image-Based VR Using a Self Calibrating Fisheye Lens. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, June 1997.

[Yagi and Yachida, 1991] Y. Yagi and M. Yachida. Real-Time Generation of Environmental Map and Obstacle Avoidance using Omnidirectional Image Sensor with Conic Mirror. *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pages 160–165, June 1991.

[Yagi, 1999] Y. Yagi. Omnidirectional Sensing and Its Applications. *IEEE Transactions*, E82-D(3), March 1999.

[Zheng and Tsuji, 1992] J. Zheng and S. Tsuji. Panoramic Representation for Route Recognition by a Mobile Robot. *International Journal of Computer Vision*, 9:55–76, 1992.