

## Multiview Panoramic Cameras Using Mirror Pyramids

Kar-Han Tan,  
Hong Hua, *Member, IEEE Computer Society,*  
and Narendra Ahuja, *Fellow, IEEE*

**Abstract**—A mirror pyramid consists of a set of planar mirror faces arranged around an axis of symmetry and inclined to form a pyramid. By strategically positioning a number of conventional cameras around a mirror pyramid, the viewpoints of the cameras' mirror images can be located at a single point within the pyramid and their optical axes pointed in different directions to effectively form a virtual camera with a panoramic field of view. Mirror pyramid-based panoramic cameras have a number of attractive properties, including single-viewpoint imaging, high resolution, and video rate capture. It is also possible to place multiple viewpoints within a single mirror pyramid, yielding compact designs for simultaneous multiview panoramic video rate imaging. Nalwa [4] first described some of the basic ideas behind mirror pyramid cameras. In this paper, we analyze the general class of multiview panoramic cameras, provide a method for designing these cameras, and present experimental results using a prototype we have developed to validate single-pyramid multiview designs. We first give a description of mirror pyramid cameras, including the imaging geometry, and investigate the relationship between the placement of viewpoints within the pyramid and the cameras' field of view (FOV), using simulations to illustrate the concepts. A method for maximizing sensor utilization in a mirror pyramid-based multiview panoramic camera is also presented. Images acquired using the experimental prototype for two viewpoints are shown.

**Index Terms**—Panoramic cameras, mirror pyramids, catadioptric systems, omnidirectional imaging and video capture, multiview panoramic imaging, stereoscopic cameras.

### 1 INTRODUCTION

PANORAMIC images and video are useful in many applications such as special effects, immersive virtual environments, remote telepresence, and video games. In recent years, the subject has been actively investigated by a number of researchers [1]. Among the numerous devices proposed for capturing panoramas, mirror pyramid-based camera systems [2], [3] are a promising approach for video rate capture, as they offer single-viewpoint imaging and use only planar mirrors that are easier to produce and introduce less optical aberration than curved mirrors [4]. It is also possible to design single mirror pyramid cameras to capture panoramas from multiple viewpoints simultaneously. Capturing panoramas from multiple viewpoints using single-viewpoint designs would require either relocating a single camera system to the different viewpoints or employing multiple systems located at all viewpoints which could operate in parallel. Obviously, the former sequential solution captures inconsistent panoramas when the scene is not static. On the other hand, the parallel solution results in bulky designs as there would need to be one mirror pyramid per viewpoint, and the

large size would require the adjacent viewpoints to be separated by a sufficiently large distance.

In spite of the many attractive properties of mirror pyramid camera systems, there has not been much attention paid to the study of their properties. We are also not aware of any prototypes and experimental validation of multiview mirror pyramid systems. In this paper, we analyze the general class of multiview mirror pyramid cameras and illustrate their properties through simulations. We also describe a method for designing and maximizing the sensor utilization in mirror pyramid camera systems. We also constructed a multiview prototype that helps validate multiview mirror pyramid designs which complement and extend the basic ideas and schematics given in Nalwa's patents [2], [3].

### 2 PREVIOUS WORK

Techniques for constructing panoramic cameras can be classified into two categories: dioptric methods, where only refractive elements (such as lenses) are employed, and catadioptric methods, where reflective components (such as mirrors) are used in combination with refractive elements. Dioptric systems include camera clusters [5], [6], fish eye lens-based systems [7], [8], [9], and rotating cameras [10], [11], [12], [13], [14], [15], [16], [17]. Catadioptric systems include sensors that use curved mirrors and a single camera [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], and sensors that employ planar mirrors and multiple cameras [4], [28], [29], [30], [31].

Dioptric camera clusters [5], [6], in which multiple cameras point in different directions to achieve a large FOV, are capable of achieving high resolution panoramic video rate capture. However, cameras in these clusters typically do not share a unique viewpoint due to physical constraints, which makes it impossible to mosaic individual images to form a true panoramic view. Although apparent continuity across images may be achieved by ad hoc image blending, panoramas produced in this manner are not suitable for tasks that require images to be captured from a single viewpoint. Systems using a fisheye lens [7], [8], [9] are able to deliver large FOV images at video rate, but have limited sensor resolution as the entire FOV is covered by a single sensor. Fisheye lenses also introduce irreversible distortion for close-by objects and may have different viewpoints for different portions of the FOV. Rotating cameras [10], [11], [12], [13], [14], [15], [16], [17], in which a conventional camera rotates about its viewpoint to acquire panoramic images, deliver high-resolution wide FOV, but are not capable of video rate panoramic capture.

Catadioptric systems that use a curved mirror to map a panoramic view onto a single sensor [18], [19], [20], [21], [22], [23], [24], [25], [26], [27] are able to achieve a single viewpoint at video rate, but have the same limitation on sensor resolution as fisheye lens-based systems. Furthermore, the resolution varies significantly with the viewing direction across the FOV. Similar to the dioptric case, this resolution limitation can be alleviated partially at the expense of the video rate capture capability by panning the camera system [32], [33], [34].

A mirror pyramid camera system, first described in [4], consists of a number of flat mirror surfaces arranged in the form of a pyramid together with a set of conventional cameras each associated with a face on the mirror pyramid. These cameras are strategically positioned such that the mirror images of their viewpoints are located at a single point within the mirror pyramid. Effectively, this creates a virtual camera with a wide FOV that is capable of capturing panoramas at video rates. Fig. 1 illustrates the evolution of mirror pyramid cameras. The first mirror pyramid camera design locates the viewpoints for the conventional cameras at a point on the main axis of the pyramid, between the apex and the base plane. Another recent prototype camera system [29] uses a

- K.-H. Tan and N. Ahuja are with Computer Vision and Robotics Laboratory, Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, 405 N Mathews Avenue, Urbana IL 61801. E-mail: {tankh, ahuja}@vision.ai.uiuc.edu.
- H. Hua is with the Optical Sciences Center, The University of Arizona, 1630 E. University Blvd., Tucson, AZ 85721. E-mail: hhua@optics.arizona.edu.

Manuscript received 4 Sept. 2002; revised 20 Sept. 2003 accepted 16 Nov. 2003.

Recommended for acceptance by L. Quan.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 117240.




Fig. 1. The evolution of mirror pyramid cameras. (a) Original design by Nalwa [4]. (b) Stereo design by Kawanishi [28], vertically stacking two of the cameras from (a). (c) Double vertical FOV design by Hua and Ahuja [29]. (d) Generalized multiview design, shown with three viewpoints [2], [3], [30]. In general, the design is able to accommodate an arbitrary number of viewpoints placed in arbitrary configurations.

double mirror pyramid (two mirror pyramids sharing a common base plane), and locates the viewpoints at the intersection point of the main axis and the base plane. This doubles the vertical FOV. A stereoscopic (two-view) mirror pyramid camera is described in [28], which vertically stacks two mirror pyramid cameras of the type presented in [4], thus realizing two vertically displaced viewpoints, one in each pyramid. Although this creates two panoramic viewpoints, the stereo displacement is in a direction orthogonal to the panoramic strip. In many applications, it is more useful to have the camera displacement aligned with the direction of the panoramic strip, conforming to the commonly encountered mode of stereo vision. This type of configuration would be necessary, for example, when the stereo video stream captured is meant to be viewed by a human user.

One possible approach to a horizontal-baseline, panoramic stereo camera design would be to have two single-viewpoint mirror pyramid cameras located side-by-side. Alternatively, a single mirror pyramid camera could be relocated to sequentially capture the panoramas at each viewpoint if the scene is stationary. Obviously, the second solution will not be capable of video rate capture, and the first would result in bulkier designs since there would be two mirror pyramids next to each other. More importantly, each camera system would occlude a large part of the other system's FOV. This problem with occlusion is compounded when the multiple viewpoints need to be within close proximity of each other, for example, when attempting to setup two viewpoints with the same spacing as that between a pair of human eyes. The size of the mirror pyramids themselves may prevent the two viewpoints from being placed that close, but even if it were possible, the closer together the two viewpoints are, the more the pyramids will occlude each other's FOV.

In a number of recently-granted patents [2], [3], Nalwa described designs for multiview mirror pyramid cameras and projection systems where multiple viewpoints can be placed within one mirror pyramid. However, these patents did not provide any analysis of the cameras. Further, although implementations of a single-pyramid single-viewpoint camera is documented [35], we are not aware of any experimental validation of single-pyramid multiple-viewpoint designs, or even designs with a single viewpoint, which is not on the axis of symmetry of the mirror pyramid.

Our goal in this paper is to analyze the properties of these cameras and provide a general method for their design. The analysis and design method was first presented in [30]. We first describe the imaging geometry of mirror pyramid cameras and investigate the relationship between the placement of viewpoints within the pyramid and the cameras' field of view (FOV), using simulations to illustrate the concepts. We also describe how the focal lengths for the cameras can be chosen given a particular orientation, and how each physical camera's orientation can be optimized to maximize sensor utilization. We show how the viewpoint position impacts the physical camera configuration. Finally, we present experimental results from a single-pyramid two-viewpoint prototype we constructed, as well as the raw




Fig. 2. The geometry of a mirror pyramid.

images from the individual cameras and the final mosaiced images. The prototype helps validate the aforementioned ideas underlying mirror pyramid cameras as well as some of those mentioned by Nalwa [2], [3].

### 3 DESIGNING MULTIVIEW MIRROR PYRAMID CAMERAS

In this section, we describe a mirror pyramid camera design that allows two or more horizontally displaced viewpoints to be located within one mirror pyramid. If desired, the viewpoints can also be placed in arbitrary spatial configurations within the mirror pyramid so that, for example, three viewpoints lie in a plane inclined at an arbitrary angle to the base plane, or four viewpoints lie at the vertices of an irregular tetrahedron with arbitrary orientation. Essentially, each viewpoint within the mirror pyramid dictates the positions of a set of conventional cameras around the pyramid. A designer can thus start with the desired spatial configuration of the viewpoints and work out the required configuration of the set of conventional cameras. Video rate multiview imaging is achieved since all cameras can acquire images at video rate simultaneously.

We start with a description of mirror pyramids. We then examine the relation between a desired viewpoint inside a mirror pyramid and the positions of the corresponding set of conventional cameras around the pyramid. Subsequent sections show how, for each conventional camera, the focal length and orientation can be chosen to maximize the utilization of each camera's optical sensor area. We then discuss the limitations and design trade offs of the proposed cameras and show the results obtained from an experimental prototype that uses four conventional cameras to realize two viewpoints.

#### 3.1 Properties of Mirror Pyramids

We now describe the class of symmetric mirror pyramids used in the designs described in this paper. A camera design may use one such pyramid, or two pyramids stacked base-to-base to form a *double pyramid*. Any such mirror pyramid can be fully characterized by the following parameters: radius, tilt angle, height, and the number of faces. Radius refers to the perpendicular distance from the main axis to the line of intersection of each planar mirror face with the base of the pyramid. Tilt angle refers to the angle between each mirror face plane and the base plane. If the pyramid is not truncated, all the mirror faces will intersect at the apex of the pyramid. The distance between the apex and the base plane is called the height of the pyramid. If the pyramid is truncated, then the distance between the truncation plane and the base plane is called its height. Finally, the number of faces refers to the number of mirror faces in a single pyramid (twice as many in a corresponding double pyramid). Fig. 2 illustrates the geometry involved.

#### 3.2 Individual Viewpoint Placement

As mentioned earlier, the early designs of mirror pyramid cameras locate the viewpoint on the axis of symmetry of the pyramid. This viewpoint is placed at the base of the pyramid in a double mirror pyramid, and at a distance away from the base in a single mirror

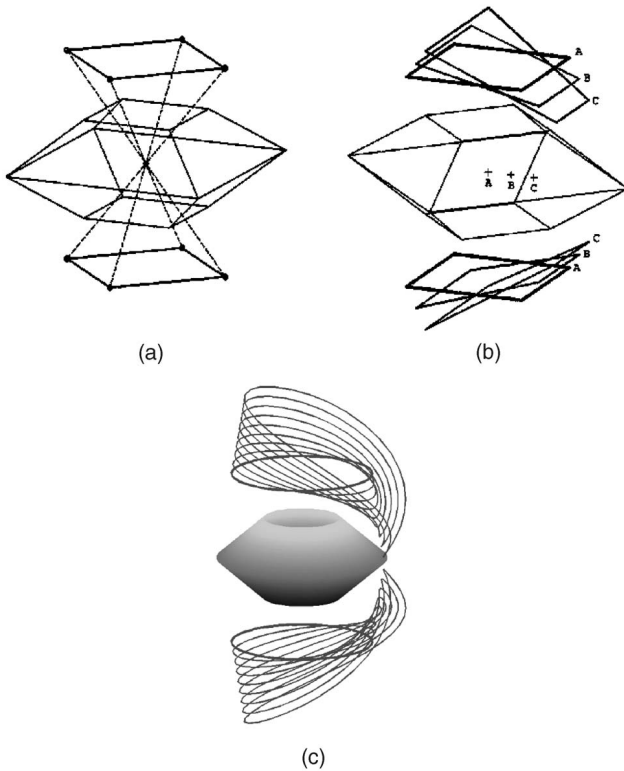


Fig. 3. Variation in the physical camera position with viewpoint position. (a) Viewpoint is centered in a four-sided pyramid. There are eight mirror faces and, thus, a ring of four cameras each for the upper and lower pyramids. (b) When the viewpoint shifts from the center, the geometric configuration of the physical cameras changes accordingly. The figure shows camera positions for viewpoints positioned at points A, B, and C. (c) Same as (b), but with a mirror pyramid with a large number of faces to show how the shape changes as the viewpoint translates from the center toward the edge of the mirror pyramid.

pyramid. More recently, Nalwa [2], [3] discussed the possibility of placing multiple viewpoints at arbitrary locations in the same pyramid. In this section, we discuss in detail the placement of arbitrary viewpoints and selection of associated imaging parameters. We start with a viewpoint located inside the mirror pyramid and projecting its image into the physical world by finding the reflections of the viewpoint in the planes containing the mirror faces. Each such projection is the location of the viewpoint of the physical camera associated with the corresponding pyramid face. An example is shown in Fig. 3a, in which a four-sided double mirror pyramid is used to create a viewpoint at its center. In the figure, dotted lines join the viewpoint and its corresponding physical camera positions for each mirror face.




Fig. 4. The geometric relationship between mirror face images and optical sensors. (a) and (b) show the side view of a sensor and the pyramid faces. The dashed lines illustrate how a mirror face edge projects onto a sensor through a given optical center. The optical sensor is shown as a thick line. In (a), the mirror face image is projected onto a larger portion of the optical sensor than in (b), where the optical axis and sensor have different orientations. (c) and (d) illustrate the variation in the mirror face image as the camera orientation changes. The quadrilateral drawn with thick lines represents the image of a mirror face on the image plane, and the small cross is the point at which the optical axis intersects the image plane. The enclosing rectangle is drawn to show the smallest possible sensor area that contains the mirror face image. Clearly, the sensor utilization is much higher in (c) than in (d).

When the viewpoint is on the main axis of the pyramid, the positions of the cameras for each of the upper and lower pyramids form the vertices of a regular polygon due to symmetry of the pyramid about the axis. As the viewpoint is shifted away from the center, the polygonal shape changes. Fig. 3b illustrates this effect for the mirror pyramid of Fig. 3a, and Fig. 3c illustrates the same effect for a pyramid with a very large number of faces; they show how the shape deforms as the viewpoint approaches the outer edge of the pyramid. It can be seen from this last diagram that the initial, almost circular (approaching a circle for a mirror pyramid with an arbitrarily large number of faces) shape smoothly deforms into an irregular nonplanar shape as the viewpoint shifts away from the center. The practical implication of this observation for camera designers is that if it is necessary for a mirror pyramid camera to change the position of a viewpoint on-the-fly, the camera mounting mechanism would have to take into account this irregular deformation.

### 3.3 Physical Camera FOV Determination

After placing the physical cameras at the locations dictated by the viewpoint in a given mirror pyramid, we need to determine the orientation and focal length of each physical camera, which together with the size of the camera optical sensor determine the FOV of each camera. While the positioning of each of the cameras are dictated by the placement of the virtual viewpoint and the shape of the mirror pyramid, the cameras are free to rotate about their optical centers. The range of orientation is however constrained by the FOV of each camera: Since the mirror face images captured by the cameras need to be mosaiced to form a seamless panoramic image, we require that each camera capture a complete image of its corresponding mirror face.

As the image of a mirror face on a sensor is in general smaller than the sensor, the sensor area is usually not fully utilized. However, given a camera orientation, it is possible to maximize the sensor usage by making the camera FOV as small as possible. This corresponds to choosing a focal length that is as large as possible for a given orientation. In this section, we show how the maximum focal length can be found for each given camera orientation. We also describe a method for optimizing sensor usage by finding the optimal camera orientation.

Before we proceed, it is important to note that in this section we are concerned with the shape and size of the mirrors, and the geometry of their projections onto camera sensors under perspective viewing transformation. We are not considering the photometric contents of the images.

#### 3.3.1 Maximal Focal Length Determination

Figs. 4a and 4b illustrate the projection geometry of a mirror face onto the sensor plane of a physical camera. In commercially available cameras, the sensor typically covers a rectangular region that is approximately centered at the point where the optical axis




Fig. 5. Maximizing sensor utilization and enforcing the uniform resolution constraint. (a) Mirror pyramid and viewpoint. (b) Ensuring maximum sensor utilization by using the maximum possible focal length or the minimal possible FOV to image the entire mirror face associated with each camera. Different cameras capture images at different resolutions. (c) Enforcing uniform resolution by using a constant FOV for all cameras and, thus, reducing overall sensor utilization.

intersects the image plane. The focal length, together with the size of the sensor, then determines the effective field of view of the camera. Given a particular orientation and position of a camera and the size of its sensor, we can then find the largest focal length such that the mirror face image is still contained in the sensor's capture area. Assuming that the sensor is rectangular and its sides are aligned with the axes of the frame of reference, the procedure to determine the focal length is as follows:

1. Given a camera orientation, find the image of the corresponding mirror face on the image plane by projecting the vertices of the mirror face.
2. Find the smallest axis-aligned rectangle on the image plane with the same aspect ratio as the sensor, centered at the optical axis, that contains the four projected vertices.
3. Find the focal length that makes the rectangle of Step 2 coincide with the sensor.

Fig. 4c and 4d show an example of mirror face images projected on the sensor plane with two different camera orientation settings. The figure also shows the bounding rectangles corresponding to the maximum allowed focal length for the two mirror face images. It can be seen that changes in orientation affect the shape, size, and location of the face image within the sensor capture area. It should be noted that this solution for focal length is unique for each given orientation and position of the camera, and the size of the sensor.

### 3.3.2 Camera Orientation Optimization

For each orientation and position setting, we now know how to determine the maximum possible focal length and the associated utilizations of the sensor area. We, therefore, search for the camera orientation that will give the maximum possible utilization of the sensor in a given camera. We define  $Utilization = \frac{F}{S}$ , where  $F$  = area of mirror face image on sensor, and  $S$  = area of sensor. We find the orientation for each camera that maximizes its utilization using the Levenberg-Marquardt method. The results are shown in Fig. 5. In Fig. 5a, a mirror pyramid is shown with an off-center viewpoint. Fig. 5b shows the results of optimizing the FOV for each camera: each rectangle represents the effective sensor capture area, and contains a quadrilateral which is the image of the corresponding mirror face.

### 3.3.3 The Uniform Resolution Constraint

When the viewpoint is not placed on the axis of symmetry of the pyramid, the maximal focal length obtained with the method described above would, in general, be different for each camera. A drawback of this situation is that the sensor resolution per unit solid angle now varies among cameras. A simple way of obtaining a uniform resolution is to find the minimum among all the optimal focal lengths found, and use this minimal focal length for all cameras. The result of imposing this uniform-resolution constraint on sensor utilization is shown in Fig. 5c. It can be seen that some of the sensors are not fully utilized. However, the per unit solid angle resolution is now constant across the entire panoramic image captured. The vertical FOV of a camera system which uses the




Fig. 6. Impact of focal length choice on vertical FOV. A mirror pyramid with a virtual viewpoint shown as a small cross inside, and the corresponding two rings of cameras. The large outer cylinder depicts the vertical FOV of the camera system. (a) Choosing minimum focal length results in variable vertical FOV. (b) Choosing maximum focal length results in uniform vertical FOV. In both cases, the FOV is delimited above and below by piecewise linear curves.

minimal focal length is shown in Fig. 6a. As can be seen in the figure, the panoramic image has a variable vertical FOV delimited above and below by a piecewise-linear curve. One can also choose the maximum focal length among all the optimum focal lengths found. This makes it impossible for some of the cameras to capture the entire mirror face, creating gaps in the visual field between adjacent mirror faces. However, the cameras can be reoriented so that the mirror face images are only clipped at the top (bottom) for cameras in the upper (lower) ring. This ensures a continuous visual field, uniform resolution, and also a uniform vertical FOV, as shown in Fig. 6b.

## 3.4 Multiview Setup Considerations

Having described the method for placing individual viewpoints at arbitrary locations within a mirror pyramid, we now discuss the issues involved in the placement of multiple viewpoints within the same pyramid. Each additional viewpoint adds a new set of physical cameras configured by the method discussed in the previous section. A fact of interest to the camera designer is that the locations of the physical cameras associated with a given mirror face corresponding to different viewpoints are simply the mirror images of the locations of the viewpoints in the face. This means that the camera configuration, i.e., relative distances and angles among the physical cameras, are the same for all mirror faces. Further, if a set of viewpoints within the pyramid undergoes a rigid transformation, the corresponding physical camera configuration also undergoes a rigid transformation which is given by reflections of the viewpoints into the corresponding faces. This is illustrated in Fig. 7. It should be noted, however, that this invariance property applies only to the camera positions and not the FOV-maximizing orientations.

The number of viewpoints and their spatial configuration will also have another constraint arising from the need to place the physical cameras around the mirror pyramid, which, of course, will depend on the physical sizes, shapes, orientations and locations of the cameras, and the size and shape of the mirror pyramid.




Fig. 7. Configuration of groups of cameras corresponding to each mirror face remains rigid as the set of virtual viewpoints undergo rigid transformations. The three images show the effect of translating a group of viewpoints on the positions of the set of physical cameras.




Fig. 8. (a) A design showing a 6-face pyramid and only two pairs of cameras ( $AA'$ ,  $BB'$ ). Each pair is associated with one face and the two viewpoints are shown as crosses (+). (b) The experimental setup implementing the design in (a). Note that the left virtual camera (inside the pyramid) is a combination of the mirror images of two cameras from two adjacent mirror faces.

#### 4 IMPLEMENTATION

We implemented a stereo (two-view) mirror pyramid camera system that utilizes two conventional monochrome cameras for each viewpoint, as shown in Fig. 8. We used an experimental setup that does not allow all the degrees of freedom required for optimized performance, as discussed in previous sections. We had limitations on the achievable camera orientations, and also employed lenses with equal focal lengths on all the cameras. The most significant implication of these limitations is that the usage of the individual sensors may not be optimized, as described in the previous sections. However, the setup validates the basic design experimentally, namely, that it is possible to construct a mirror pyramid camera with more than a single viewpoint, each located at an arbitrary position within the mirror pyramid.

We estimated the intrinsic parameters and the radial distortion are estimated and compensated for using the camera calibration software described in [36]. Fig. 9 shows the images captured by each individual camera (after radial distortion compensation). The mosaiced images are shown in Fig. 10. It may appear from the experimental results that the setup shown makes suboptimal utilization of the sensors and it might even be possible to obtain the same results using a pair of conventional cameras. However, as noted earlier in this section, the experimental setup utilizes only two faces of the mirror pyramid and does not fully maximize the use of sensor area. If a full set of cameras were used, even this suboptimal setup would still be able to capture 360-degree panoramas, which is beyond the capabilities of a conventional camera. This experimental prototype, as far as we know, is the first physical realization of a single-pyramid, multiple-viewpoint, mirror pyramid camera.

#### 5 CONCLUSION AND FUTURE WORK

We have investigated the class of mirror pyramid cameras that place multiple conventional cameras around a mirror pyramid and provide panoramic views from multiple viewpoints. We have studied the impact of changing viewpoint on the placement of the conventional cameras around the pyramid, and experimentally demonstrated the feasibility of a two-view mirror pyramid camera.




Fig. 9. The raw images captured by the four cameras, after correcting for radial distortion. (a) and (c) are from one viewpoint and need to be mosaiced to form a continuous visual field. (b) and (d) are from the second viewpoint and, likewise, need to be mosaiced.




Fig. 10. The stereo pair of panoramic views captured with the experimental setup. The top view is the result of mosaicing Figs. 9a and 9c, and the bottom view is the result of mosaicing Figs. 9b and 9d.

While the basic ideas for mirror pyramids were first described by Nalwa [2], [3], [4], this paper makes the following contributions to the study and understanding of mirror pyramid cameras:

1. **Vertical FOV Analysis.** We show how the shape of the mirror pyramid and the choice of focal lengths for the cameras affect the vertical FOV.
2. **Camera Focal Length Selection.** We show how the focal length for each individual camera can be chosen to maximize sensor utilization given a fixed orientation, and how the focal lengths can also be chosen to impose a uniform-resolution constraint.
3. **Camera Orientation Optimization.** We present a method for optimizing the camera orientation (and, consequently, their focal lengths) to maximize overall sensor utilization in the camera system.
4. **Impact of Viewpoint Position on Camera Configuration.** We show how the camera configuration varies with the position of the viewpoint, and also show that while the configuration of each set of camera corresponding to each viewpoint undergoes a nonrigid transformation, the configuration of the set of cameras corresponding to each face of the mirror pyramid remains rigid if the set of multiple viewpoints undergoes a rigid transformation.
5. **Experimental Validation.** We have constructed an experimental prototype that helps validate the basic design ideas of single-pyramid multiple-viewpoint mirror pyramid cameras.

In our ongoing work, we are investigating the use of these multiview mirror pyramid cameras in areas such as robot navigation and immersive telepresence. The process of mosaicing images of adjacent FOVs acquired by multiple cameras also needs to be examined since it may lead to optical artifacts such as those noted in [29].

#### ACKNOWLEDGMENTS

The support of the US National Science Foundation under Grants ITR 0083037 and ECS 02-25523 is gratefully acknowledged. Hong Hua is also supported by a Beckman Institute Fellowship. The authors would like to thank Chunyu Gao and John M. Hart for their help and suggestions.

#### REFERENCES

- [1] R. Benosman, S.B. Kang, and O. Faugeras, *Panoramic Vision*. Springer-Verlag, 2001.
- [2] V.S. Nalwa, "Stereo Panoramic Viewing System," US patent, vol. 6,141,145, 2000.
- [3] V.S. Nalwa, "Panoramic Viewing System with Offset Virtual Optical Centers," US patent, vol. 6,219,090, 2001.
- [4] V.S. Nalwa, "A True Omnidirectional Viewer," Bell Laboratories technical report, 1996.
- [5] P.I. Anderson, "From Telepresence to True Immersive Imaging: Into Real-Life Video-Now!," *Advanced Imaging*, vol. 10, no. 7, pp. 48-50, 1995.
- [6] R. Swaminathan and S.K. Nayar, "Nonmetric Calibration of Wide-Angle Lenses and Polycameras," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, Oct. 2000.

- [7] K. Miyamoto, "Fish Eye Lens," *J. Optical Soc. of Am.*, vol. 64, pp. 1060-1061, 1964.
- [8] Y. Xiong and K. Turkowski, "Creating Image-Based Virtual Reality Using a Self-Calibrating Fisheye Lens," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 1997.
- [9] S. Zimmerman and D. Kuban, "A Video Pan/Tilt/Magnify/Rotate System with No Moving Parts," *Proc. IEEE/AIAA Digital Avionics Systems Conf.*, 1992.
- [10] A. Castano and N. Ahuja, "Omnifocused 3D Display Using the Nonfrontal Imaging Camera," *Proc. Workshop Computer Vision for Virtual Reality Based Human Comm.*, 1998.
- [11] A. Krishnan and N. Ahuja, "Range Estimation from Focus Using a Non-Frontal Imaging Camera," *Proc. Nat'l Conf. Artificial Intelligence*, 1993.
- [12] A. Krishnan and N. Ahuja, "Range Estimation from Focus Using a Nonfrontal Imaging Camera," *Int'l J. Computer Vision*, vol. 20, no. 3, pp. 169-185, 1996.
- [13] A. Krishnan and N. Ahuja, "Panoramic Image Acquisition," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 1996.
- [14] M. Aggarwal and N. Ahuja, "On Generating Seamless Mosaics with Large Depth of Field," *Proc. Int'l Conf. Pattern Recognition*, 2000.
- [15] M. Aggarwal and N. Ahuja, "A New Imaging Model," *Proc. IEEE Int'l Conf. Computer Vision*, 2001.
- [16] M. Aggarwal and N. Ahuja, "High Dynamic Range Panoramic Imaging," *Proc. IEEE Int'l Conf. Computer Vision*, 2001.
- [17] S. Peleg, "Panoramic Mosaics by Manifold Projection," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 1997.
- [18] J.S. Chahl and M.V. Srinivasan, "Reflective Surfaces for Panoramic Imaging," *Applied Optics*, vol. 36, no. 31, pp. 8275-8285, 1997.
- [19] J.S. Chahl and M.V. Srinivasan, "A Complete Panoramic Vision System, Incorporating Imaging, Ranging, and Three Dimensional Navigation," *Proc. Workshop Omnidirectional Vision*, pp. 104-111, 2000.
- [20] R.A. Hicks, "Reflective Surfaces as Computational Sensors," *Proc. Workshop Perception for Mobile Agents*, 1999.
- [21] R.A. Hicks and R. Bajcsy, "Catadioptric Sensors that Approximate Wide-Angle Perspective Projections," *Proc. Workshop Omnidirectional Vision*, pp. 97-103, 2000.
- [22] H. Ishiguro, M. Yamamoto, and S. Tsuji, "Omni-Directional Stereo," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 257-262, 1992.
- [23] S.K. Nayar, "Catadioptric Omnidirectional Camera," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 1997.
- [24] S.K. Nayar and V. Peri, "Folded Catadioptric Cameras," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 1999.
- [25] K. Yamazawa, Y. Yagi, and M. Yachida, "Omnidirectional Imaging with Hyperboloidal Projection," *Proc. Int'l Conf. Intelligent Robots and Systems*, 1993.
- [26] Z. Zhu, K.D. Rajasekar, E.M. Riseman, and A.R. Hanson, "Panoramic Virtual Stereo Vision of Cooperative Mobile Robots for Localizing 3D Moving Objects," *Proc. Workshop Omnidirectional Vision*, 2000.
- [27] S. Peleg, M. Ben-Ezra, and Y. Pritch, "Omnistere: Panoramic Stereo Imaging," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 3, pp. 279-290, 2001.
- [28] T. Kawanishi, et al. "Generation of High Resolution Stereo Panoramic Images by Omnidirectional Imaging Sensor Using Hexagonal Pyramidal Mirrors," *Proc. Int'l Conf. Pattern Recognition*, 1998.
- [29] H. Hua and N. Ahuja, "A Uniform-Resolution Panoramic Camera," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 2001.
- [30] K.-H. Tan, H. Hua, and N. Ahuja, "Multiview Panoramic Cameras Using a Mirror Pyramid," *Proc. IEEE Workshop Omnidirectional Vision*, 2002.
- [31] M. Aggarwal and N. Ahuja, "Split Aperture Imaging for High Dynamic Range," *Proc. IEEE Int'l Conf. Computer Vision*, 2001.
- [32] S. Coorg, N. Master, and S. Teller, "Acquisition of a Large Pose-Mosaic Dataset," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 1998.
- [33] S.K. Nayar and A. Karmarkar, "360x360 Mosaics," *Proc. IEEE Computer Soc. Conf. Computer Vision and Pattern Recognition*, 2000.
- [34] H.-Y. Shum and R. Szeliski, "Panoramic Image Mosaics," Microsoft Research Technical Report, vol. MSR-TR-97-23, 1997.
- [35] V.S. Nalwa, <http://www.fullview.com>, 1999.
- [36] Z. Zhang, "A Flexible New Technique for Camera Calibration," Microsoft Research Technical Report, vol. MSR-TR-98-71, 1998.