A True Omni-Directional Viewer†

Vishvjit S. Nalwa
AT&T Bell Laboratories
Holmdel, NJ 07733, U.S.A.

Abstract

I describe here the design of what I believe is the first true omni-directional camera: a camera that can capture, at any instant, a 360° view of a scene effectively from a single viewpoint. My design is simple and robust, and extends easily to accommodate the following: projectors, rather than imagers, viewers that image/project most of the visual sphere, rather than just a visual cylinder, and viewers that provide high-resolution views using low-resolution viewing devices. An instance of my design that has been implemented and I shall outline provides a live seamless 360° view of a scene using four individual CCD cameras, each camera with a 90°+ field of view and viewing the scene off a face of a mirrored pyramid.

† The commercial interests of AT&T Bell Laboratories in this work are protected by numerous patents filed in the U.S.A. and other countries.

Contents

1. Introduction

I propose here what I believe is the first design of a true omni-directional camera: a camera that can capture, at any instant, a 360° view of a scene effectively from a single viewpoint. I shall describe this design and explain why other previously proposed designs do not provide true omni-directional views. Then, I shall outline an implementation of my design and describe this implementation's performance.

My design is truly simple: It is based on the principles of two well-known optical devices, the pinhole camera and the planar mirror. The principles of these devices have been well understood for several centuries—probably since Kepler in the early seventeenth century—and their operation has been observed for at least another 2000 years, at least since the fifth century B.C., when Mo Ti recorded his observations in China on the operation of a pinhole camera (see [Rosenblum 1989] and [Ronchi 1957]).

In addition to its conceptual simplicity, my design is robust, requires minimal hardware and computation, and extends easily to accommodate the following: projectors, rather than imagers, viewers that image/project most of the visual sphere, rather than just a visual cylinder, and viewers that provide high-resolution views using low-resolution viewing devices.

An instance of my design that has been implemented provides a live seamless 360° view of a scene using four individual CCD cameras, each camera with a 90°+ field of view and viewing the scene off a face of a mirrored pyramid. Competing designs fall broadly into one of four categories, each with its own one or more shortcomings: They comprise multiple ordinary cameras looking out into the scene directly in different directions, and, hence, they work on only scenes with no object close to the omni-directional camera; their lenses rotate, and, hence, they work on still scenes only; and, they use either single nonplanar mirrors or single fish-eye type lenses, and, hence, do not view the scene effectively from a single viewpoint and typically offer lower-than-normal image resolution.

Potential applications of an omni-directional camera include use in interactive television, in surveillance systems, in video conferencing, and in home shopping. **A key benefit of an omni-directional camera is that several users can simultaneously, independently, and perhaps remotely view the complete 360° panoramic image and individually choose from it a narrower view in any direction in greater detail.** This view, in effect, is provided by a virtual ordinary camera whose pan—and to a limited extent, tilt and zoom—is under each user's complete simultaneous control. Recall that **pan** refers to camera rotation about the vertical axis, **tilt** to its rotation about a horizontal axis, and **zoom** to image magnification by narrowing the camera's field of view.

2. Fundamentals

As I mentioned in the introduction, the design I have in mind is very simple. However, before I can explain it to you, we must recall a few key fundamentals of imaging—described at length, for instance, in [Nalwa 1993]:

1. *A typical camera with a lens is an attempt simply to replicate the geometry of a pinhole camera.* Figure 1 illustrates a pinhole camera. The reason for using a lens in a typical camera is to keep the image sharp and unblurred while we increase the image brightness by employing an aperture much larger than a pinhole.

2. *The geometry of image formation in a pinhole camera follows perspective projection of the scene onto a plane up to an inversion, with the center of projection at the pinhole.* Figure 2 illustrates perspective projection. Then, for all purposes of reasoning and analysis, we may consider the plane on which the image is formed in a pinhole camera to be in front of the pinhole, rather than behind the pinhole.
3. *The surface in a pinhole camera on which an image is formed, or equivalently, the surface on which we consider the scene to be perspectively projected, need not be planar at all.* It could be any non-self-occluding surface of known shape, as illustrated in Figure 3, and, we would, in principle, have exactly the same information as we would in a planar image.

In other words, **all the information in a pinhole image can be represented as a continuum—or, in the case of a discrete image, as a collection—of light rays toward a single viewpoint, each ray with a particular orientation, brightness, and color.** This completes our recollection of some fundamentals of imaging.

3. The Task

The task we face is designing a true omni-directional camera. **I define a true omni-directional camera as one that can provide an image of the surrounding world, all 360° around, at any given instant from a single viewpoint.** As shall soon become apparent, the challenge in this task is conceptually equivalent to that in designing a projector that can project an image all 360° around at any given instant from a single viewpoint.

A typical camera can image the world around only within a circular cone that is no more than about 60° at its vertex, as shown in Figure 4(a). What we are seeking is to extend this typical camera's field of view to encompass all 360° around in one dimension, as illustrated in Figure 4(b).

In my definition of a true omni-directional camera, I have included two criteria over and above that the view should encompass 360° in one dimension. One, the complete 360° view should be available at any given instant; this criteria is necessary to ensure that moving objects in the world appear coherent and unblurred in the image. Two, the 360° view should be from a single viewpoint, or, in the terminology of the preceding section, the complete 360° view should have a common center of projection. This criteria allows us to reconstruct from each omni-directional view, typical narrower-angle views that are perspective projections of the scene, which we are accustomed to seeing—with straight lines in the scene appearing as straight lines in planar images for instance. As far as I can tell, no existing design of an omni-directional camera satisfies both these criteria simultaneously.

4. The Problem

Now, assuming the pinhole geometry for each camera, **the problem we face in designing an omni-directional camera from multiple pinhole cameras is that we must colocate the multiple pinholes, each pinhole directed to capture a different set of light rays, without one camera obstructing the view of another.** This problem is illustrated in Figure 5. As is well known—see, for instance, [Nalwa 1993]—the farther away the pinholes are from each other, the greater shall be the disparity between their images within the common fields of view of the pinholes. This disparity, of course, is the basis of stereo vision. And the greater this disparity, the more difficult it shall be to combine the individual images into a single coherent omni-directional image. I am assuming here that we do indeed want to use multiple cameras to accomplish our task at hand. I shall address the problems we face when trying to use a single camera for omni-directional viewing when I outline alternatives to my design.

A simple calculation illustrating the disparity in the views of multiple pinhole cameras whose pinholes are not colocated is instructive here. Consider Figure 6, which shows two pinholes P_1 and P_2, located a distance D apart. Now, consider two rays r_1 and r_2 from a point at infinity toward the two pinholes P_1 and P_2, respectively, the point at infinity located along the perpendicular bisector of the line connecting the two pinholes. These two rays, with identical orientations, but displaced a distance D apart, provide a reference direction for the two views from the two pinholes. Now, consider a point P in the scene, this point along the perpendicular bisector of the line connecting the two pinholes and at a distance d from that line, as shown. We have located our scene point P as described only to make our argument simple, but not at the expense of generality. Now, it is clear that the clockwise angle θ between the ray from the point P toward the pinhole P_1 with respect to the ray r_1, is the same as the anticlockwise angle ϕ between the ray from the point P toward the pinhole P_2 with respect to the ray r_2. In other words, the angular disparity between the images of point P in the pinhole cameras P_1 and P_2 is simply 2θ, which a simple calculation reveals to be $2\tan^{-1}(D/2d)$. As highlighted in Section 2, it is the orientation, color, and brightness of each ray that contain all the image information pertaining to the source of the ray, and we have just shown that **the disparity in the orientations of the two (rays**

forming the) images of a single point in two pinhole cameras increases as we increase the displacement between the two pinholes and decreases as we increase the distance of the point from the two pinholes. This observation is quite general, even though the geometry from which we arrived at it here is not so general.

5. The Solution

The solution to colocating multiple pinholes, without one pinhole camera obstructing another's view, is to use one or more planar mirrors. Figure 7 illustrates the image *I* of a point *S* in a planar mirror, as explained by Kepler. Whereas this geometry of reflection in a planar mirror is so familiar to us today, if we consider that it was first elucidated by Kepler in 1604 "after so many centuries of bafflement and confusion, Kepler's explanation seems simply marvelous" [Ronchi 1957].

Now, returning to the geometry of reflection in a planar mirror shown in Figure 7, if a pinhole of a pinhole camera is located at point *S*, it shall be mapped to point *I* by the planar mirror *M*. Then, the image of the world—in this case, of the two eyes shown—seen from *S* off *M* shall be the same as the image that would be seen directly from *I*, except for the familiar left–right reversion. Thus, **a planar mirror effectively maps a pinhole's view to that from the pinhole's mirror image, except for a left–right reversion**.

6. The Design

Now, if we had two planar mirrors, we could, in effect, colocate two pinholes using the mirror arrangement shown in Figure 8. This brings us to my design: **My design of an omni-directional camera consists of a mirrored *n*-sided pyramid with a pinhole camera looking off each side of the pyramid such that the mirror image of every pinhole (in the mirror the pinhole looks off) lies at the same location in space.** Figure 9 illustrates this design with a 4-sided pyramid. By an ***n*-sided pyramid** I mean a polyhedron, with $n+1$ faces in all, whose base is an n-gon and whose n sides extending from its base to its vertex are triangles. As we are concerned only with the surface of the pyramid off which our cameras see the world, our pyramid here could be partial—perhaps, with its base and vertex missing.

I have intentionally left the above statement of the design general to emphasize the variations possible within its framework. However, let us now pin down several aspects of the design to simplify both its realizability and our discussion of it. First, let us assume that our mirrored pyramid is a **right regular pyramid**—that is, a pyramid whose base is a regular n-gon (a polygon with n identical sides and n identical corners) and whose vertex is perpendicularly above the center of its regular base. Then, each side of the pyramid extending from the base of the pyramid to its vertex shall be an identical isosceles triangle. Let us further assume that each triangular side of the pyramid forms a 45° angle with a plane parallel to the pyramid's base. Now, let us stand such a pyramid on its vertex with its base horizontal, as shown in Figure 9. Further, as also shown in the figure, let us position all the pinholes in the horizontal plane that contains the pyramid's vertex such that each pinhole is equidistant from the vertex and along the bisector of the horizontal projection of the vertex angle of the associated pyramid side. Then, as each pyramid side is at 45° to the horizontal plane, it is easy to see that the mirror images of all the pinholes shall coincide at that point in space on the axis of the pyramid that is the same distance from the vertex as the pinholes are from the vertex. If our cameras are now pointed vertically upward, they shall effectively view the world horizontally outward from a single point in space up to a lateral reversion, which is easy to undo.

Now, let us also assume that the imaging surface of each pinhole camera is a plane that is horizontal. Then, each triangular face of the mirrored pyramid shall be imaged as a rectangle of infinite extent, with the vertex of the pyramid being imaged at infinity in each pinhole camera, and the rectangular image oriented such that it is symmetric about the vertical plane containing the vertex and the pinhole. Can you see this? It has been a cause of considerable confusion to several individuals who have examined my design. An easy way to see this is to consider an object that is a rectangle of the type described, and let the plane of the associated pyramid side be the image plane, with the pinhole's position unchanged. Then, the vanishing point of the two infinite sides of the rectangle is easily seen to be at the vertex of the pyramid—see, for instance, [Nalwa 1993]—and because projection from one planar surface to another is completely reversible, we have our sought result.

Finally, yes, we have limited ourselves only to pinhole cameras to this point, not because we plan to use such cameras, but because a pinhole camera provides a very useful approximation to the geometry of most cameras with lenses. Under this approximation, the pinhole corresponds to the **optical center** of the lens if the camera lens is thin. More generally, as illustrated in Figure 10, the pinhole corresponds to the **first nodal point** or

forward nodal point of the camera lens as far as rays from the world to the lens are concerned, and to the **second nodal point** or **rear nodal point** of the camera lens as far as rays from the lens to the imaging surface are concerned. Figure 10 also illustrates the **optical axis** of the lens, which is the axis of rotation of the lens; the two nodal points lie on this axis. Thus, from a geometrical point of view, a typical lens can be approximated by two pinholes, an entrance pinhole, which corresponds to the forward nodal point, and an exit pinhole, which corresponds to the rear nodal point. When we refer simply to a pinhole of a camera with a lens, without further qualification, we shall be referring to its entrance pinhole. Of course, there are rays emanating from any object point that enter the lens through a path other than that toward the forward nodal point, such rays finally reuniting with the ray directed toward the forward nodal point at the image of the object point. It is the collection of such rays that makes the use of lenses desirable to begin with. The complete cone of light rays emanating from an object point that is gathered by a lens to from an image of the point is determined by the **entrance pupil** of the lens, which, in effect, is the imaging aperture of the lens. This entrance pupil, illustrated in Figure 10 in front of the forward nodal point of the lens, could as well be located at or behind the forward nodal point. For more on lenses, see the books by Cox [Cox 1974] and Hecht [Hecht 1987], especially the book by Cox, *Photographic Optics*, which provides the most lucid description of nodal points and other lens characteristics I have seen.

7. Alternate Designs

I promised you in the introduction that I would establish that other designs proposed for omni-directional cameras do not provide true omni-directional views, where a true omni-directional view is a 360° view available at any given instant from a single viewpoint—mathematically, from a single center of projection. As we shall see now, at worst, other designs do not provide an instantaneous 360° view, and, at best, they provide an instantaneous 360° view from a continuum of viewpoints, and with limited image resolution.

The obvious way to design an omni-directional camera is to pack together a collection of small cameras pointing in different directions, perhaps with their pinholes centered on the faces of a regular Platonic solid; see [Coxeter 1969] for an account of the five Platonic solids. There is a commercial product available based on this design (see [Anderson 1995]). When I first set out to design an omni-directional camera, I began by considering such a strategy, using optical fibers to pack together lenses pointing in different directions, each lens at one end of an optical fiber at whose other end is an imaging surface. But a simple analysis, of the type in Section 4, quickly revealed that even if each lens was just half an inch in diameter, the disparity between adjoining images of an object several feet away from the lens packing would be too large to allow us to assemble directly the individual images into a seamless composite image. Under such circumstances, we could attempt to blend the images together at the seams, perhaps using sophisticated techniques, such as in [Burt and Adelson 1983], but blending is time consuming and guaranteed to degrade the image fidelity. Hence, I summarily rejected such an approach.

Then, three possibilities remain for the design of an omni-directional camera. We could either rotate our camera or lens about its pinhole, looking at different directions at different instants, or we could alter the optical paths of light rays from the world toward the pinhole—either by reflection or by refraction—such that we are able to image the world all around. My design, of course, alters the paths of light rays through reflection, but it is not the only such design. Let us now discuss, in turn, each of these three strategies to create omni-directional images.

7.1 Motion Based

Panoramic views have held a fascination for both photographers and the public at large since the very invention of photography, and rotating a camera or lens to create a panoramic view is the oldest approach toward this goal.

Photography was invented in the summer of 1827 by Joseph Nicéphore Niépce, who exposed a pewter plate coated with bitumen for eight hours in a *camera obscura* to obtain an image of a view from a window in his estate at Le Gras in France. Within 25 years of this invention, photographers were arranging, in contiguous order, series of individual photographs obtained by rotating a camera to depict a wider view than that possible with a single photograph. Before long, a 360° panorama had been created in this fashion—the first, probably a view of Chicago in 1859 by Alexander Hesler. See [Rosenblum 1989] for a comprehensive history of photography, and [Meehan 1990] for more on panoramic photography.

Similar efforts—to create 360° views from discrete images obtained by rotating a camera—have been reported more recently too, the emphasis of these efforts being on the creation of composite images that are visually seamless. See, for instance, [Focal 1969] under the topic *Panorama*, [Dixon

1989], [Szeliski 1994], [McMillan and Bishop 1995], and [Chen 1995], the last of these efforts leading to a commercial product (see [Halfhill 1995]).

It is not hard to see that if we want an omni-directional view from a single viewpoint using such an approach, we must rotate the camera about its forward nodal point. Else, the resulting view shall not be from a single viewpoint, and also we shall have to expend considerable effort to produce a seamless result from the individual images.

An alternative to acquiring discrete views by rotating a camera, and then merging these views into a single 360° view, is to build up a 360° cylindrical view by sweeping over the imaging surface an image through a slit while rotating the camera about an axis parallel to this slit. This approach, which too has its origins in the mid-nineteenth century, has traditionally been the most popular method of creating a 360° image, with the venerable Kodak No. 10 Cirkut camera being its workhorse (see [Meehan 1990]). In this approach, the camera must rotate about its rear nodal point if we wish to keep the images of distant points sharp (see [Cox 1974]). But, then, the images of points close to the camera will not be as sharp, and, further, the complete image will not be acquired from a single viewpoint, but from a circular collection/continuum of viewpoints, the radius of this circle being the distance between the forward and rear nodal points. More recent efforts based on this approach, employing electronic rather than film cameras, include [Zheng and Tsuji 1990] and [Krishnan and Ahuja 1993].

The primary disadvantage of a motion-based approach to creating an omni-directional view is, of course, that the view is not instantaneous, and, hence, incapable of capturing moving objects that are arbitrarily located sharply and coherently. The total time required to acquire a 360° view in a scanning-slit camera can be anywhere from a second to a few minutes, several seconds being typical (see [Meehan 1990]). Two additional causes of concern in a scanning-slit camera are uniformity of mechanical motion, lack of which leads to uneven exposure, and synchronization of camera motion with image capture, lack of which leads to image distortion even for stationary objects. Further, the strategy of combining discrete individual images typically requires considerable post processing of images, and the strategy of sweeping a cylindrical image provides a view from a circular collection/continuum of viewpoints, as already explained.

7.2 Reflection Based

This approach has again apparently been around for over a century. A passing reference is made to it in [Meehan 1990] without any details, where it is mentioned that designers had tried conical mirrors to extend the field of view of a lens. A direct way to use a conical mirror to create a panoramic view is to place the mirror right above the lens, as shown in Figure 11, with the vertex of the conical mirror—the cone of this mirror assumed to have a right circular cross section—facing the lens, and its axis aligned with the optical axis of the lens. Such an approach has been outlined recently in [Jarvis and Byrne 1988] and [Yagi and Yachida 1991]. An alternative to using a conical mirror is to use a mirror of another shape, such as a sphere.

The fundamental problem with viewing the world off a nonplanar mirror is that the rays of light that eventually form the image do not all share the same viewpoint. To establish this, consider the pinhole of the lens in Figure 11—that is, its forward nodal point—and then consider the paths of individual rays from the world, off the mirror, and toward the pinhole. What you will observe from the figure is that the rays that come off the left and right edges of the conical mirror have different effective (virtual) pinholes: The effective pinhole for the rays on the right is at the reflection of the actual pinhole in the tangent plane to the right edge of the cone, and, similarly, the effective pinhole for the rays on the left is at the reflection of the actual pinhole in the tangent plane to the left edge of the cone. Then, it is clear that the world viewed off a conical mirror whose vertex is above the pinhole, as in all previously proposed designs, is viewed (virtually) from a continuum of viewpoints that lie on a circle, as illustrated in the figure. Similar arguments apply to every nonplanar mirror, except that, in general, the viewpoints shall lie not on a curve, but on a surface, as would be the case for a sphere for instance. Returning to the conical mirror, it is now clear that if we wish to reduce the circle of viewpoints to a single point, we must colocate the (physically absent) vertex of a truncated conical mirror with the pinhole itself, but then we might not have much of a useful view of the world.

In our discussion here to this point, we have assumed a pinhole aperture implicitly. Now, as we discussed in Section 2, in practice, to increase the image brightness we must use an aperture larger than a pinhole. But as we increase the aperture size here, not only shall we increase the image brightness, but we shall also increasingly blur the image due to our use of a nonplanar mirror. To see this, reconsider Figure 7, replacing an eye with a lens, and then noting the impact of the mirror *M* being nonplanar on light rays emanating from the object point S, off M, toward the lens. Clearly, these rays will not appear to originate from a single (virtual) point in space, and thus, they will be imaged as originating from various points in space. As a result, the image of the object point will be blurred, the extent of this blurring depending on the total curvature of the intersection of the mirror with the cone of rays collected by the lens to form an image of the object point. This curvature depends not only on the shape and location of the mirror—the farther away the mirror from the lens, and the flatter its local shape, the less this curvature—but also on the aperture size, increasing which increases this curvature. Thus, when we use a nonplanar mirror, our maximum usable aperture size is limited by the mirror. To ease this limit, we may make our mirror have a low curvature and be distant from the lens, but then our effective viewpoints might have a locus with a greater spatial extent and our useful field of view might be further limited.

In addition to the lack of a single viewpoint with nonplanar mirrors, and limitations on the maximum usable aperture size, a practical problem with using a single camera is the resolution of the captured image: Capturing a 360° view on a single piece of standard film, or with an electronic imager—which nowadays is typically a charge-coupled device (CCD) array—provides inadequate image detail. Whereas, in principle, this limitation can be overcome if we are using film—by resorting to a larger than normal format—we are somewhat limited when imaging electronically. An electronic imager typically consists of an array of discrete sensors arranged in a rectangular grid, which is about 500 wide along each dimension. The size of this array can currently be increased to about a 1000 or so along each dimension, at a substantial price, but even that would not provide adequate resolution in the useful part of the captured image here. Note here also that, although large-format film might give us the resolution we are seeking, film cannot provide us live images, as can electronic imagers, and, in many applications, such as broadcast television, live images are what we seek.

Before we move on to the third approach, I would like to establish further connections between my work, which clearly belongs to the current category, and previous work. The use of a planar mirror to map the center of projection **from which** light rays are emanating is widespread and well-established in optics; see, for instance, the common overhead transparency projector. On the other hand, the use of a planar mirror to map the center of projection **toward which** light rays are directed seems to be less common—perhaps, because of an inadequate appreciation of the connection between a camera with a typical lens and a pinhole camera. However, more than one example of such a use of a planar mirror is described in [Judd and Smoot 1989], where it is sought to increase the horizontal resolution of standard-resolution electronic cameras by combining images from multiple cameras.

7.3 Refraction Based

This is a fascinating approach, even though it is unable to provide omni-directional views from a single viewpoint. This approach is based on using so-called fish-eye lenses, which provide extremely large fields of view, even greater than 180°. This extremely large field of view is accomplished in a fish-eye lens by introducing one or more meniscus-shaped lens elements at the front of the lens to compress the desired field of view into a much smaller field that can be imaged by the rest of the lens. See Figure 12(a). In fact, there are fish-eye adaptors available, which when attached to the front of a normal lens, convert it to a fish-eye lens. It is clear that a fish-eye lens can provide us an omni-directional view at the fringe of its image. In fact, it does more: It provides a view of the top too. See [Kingslake 1989] for a brief introduction to fish-eye lenses.

Fish-eye lenses are not typical: They have large intentional image distortion, which is necessary to form an extremely wide-angle image on a limited-size flat surface. Further, fish-eye lenses do not obey the pinhole geometry at all—a geometry normal lenses try so hard to duplicate. For the arrangement shown in Figure 12(a), where the inner surface of the single meniscus is part of a sphere whose center at the center of the physical aperture of the lens below the meniscus, the locus of the effective viewpoints for light rays within the plane of the figure that are directed toward the center of the physical aperture after passing through the meniscus is a cusped curve of the type indicated in the figure. Mathematically, this curve is a *diacaustic* of the outer curve of the meniscus cross section; see [Yates 1974] for a description of diacaustics. But, as every cross section of the lens through its optical axis is identical, the locus of the viewpoints for all the light rays directed toward the center of the physical aperture is a surface of revolution of the type shown in Figure 12(b), whose every cross section is identically the diacaustic sketched in Figure 12(a). Now, if we assume that the lens below the aperture in Figure 12(a) is a typical lens, with forward and rear nodal points, then the locus of the effective viewpoints of light rays directed toward the forward nodal point of this lens after passing through the frontal meniscus will also have a shape similar to that illustrated in Figure 12(b). Further, although not all fish-eye lenses have a single outer meniscus element as in Figure 12(a), the general nature of the frontal elements of all fish-eye lenses is the same, and, hence, every fish-eye lens effectively has a continuum of viewpoints on a surface of the type shown in Figure 12(b). The distribution in space of these viewpoints would not be a cause for concern if this distribution had a small extent—say, a maximum dimension across of a small fraction of an inch. But such a small extent is not practical, as a large frontal element is necessary in a fish-eye lens to make our usable aperture size sufficiently large—the reason for this requirement being similar to that for wanting a locally flat and distant nonplanar mirror in Section 7.2. For example, in a well-known fish-eye lens, whose angular field of view of 220°, the outer diameter of the lens is 9.3 inches [Nikon 1995].

Thus, we see that the fundamental problems with using fish-eye lenses are similar to those with using nonplanar mirrors, which we discussed in Section 7.2: one, the lack of a single viewpoint, and, two, a limit on the maximum usable aperture size without resorting to a large frontal lens element. Further, the limits on the resolution of a single captured image described in Section 7.2 also hold here, and the useful portion of a captured image might typically be limited to the periphery of the image.

Fish-eye lenses get their name from their origin, which lies in an attempt by Wood [Wood 1906] to duplicate the monocular view of a fish of the world above water from below its surface. This view encompasses the complete hemisphere bounded below by the water surface. See [Wood 1906] for several interesting fish-eye views. From Wood's original idea, two lens designs followed: the first by Bond [Bond 1922], and the second by Hill (see [Hill 1924] or [Beck 1925]). Hill's design, on which Figure 12(a) is based, is the forerunner of currently prevalent designs, such as the one described in [Miyamoto 1964].

So far, we have discussed only the fish-eye lens itself. Now, let us turn to reports of the creation and use of omni-directional views provided by such lenses. Amazingly, Hill [Hill 1924] not only proposed a lens design, but also realized and described the use of such a lens first to capture highly distorted fish-eye images, and then to provide a host of user-chosen undistorted narrow-angle images of the original scene from each fish-eye image. More recently, Lippman [Lippman 1980] reported using an interesting variation of a fish-eye lens that deemphasized the central portion of its field of view to devote a larger fraction of its image to the periphery of its field of view; Lippman also reported using a conical mirror, interestingly not to capture an omni-directional image, but to display it. Two other interesting and relevant lenses I would like to bring to your attention here are the Sutton lens (see [Kingslake 1989]) and a nameless lens that I have seen mentioned and diagramed only in [Focal 1969] under the topic *Panoramic Camera*. The use of a fish-eye lens for omni-directional imaging has also been reported in [Ripley 1989] and [Oh and Hall 1987]. Only the last effort, of all I have mentioned, uses an electronic camera for imaging, rather than film, and, not surprisingly, the aim of this effort was to acquire some useful data for robotic tasks, rather than to produce images with adequate quality for human viewing. Finally, I mention that there is a well-known commercial product based on this approach to omni-directional imaging (see [Greene and Heckbert 1986]), this commercial product using film that has a larger-than-normal format for reasons we discussed in Section 7.2.

8. Practice

Our discussion of my design in Section 6 was primarily conceptual. However, to turn concept into reality, we must address several practical issues. Let us make these issues, then, our first order of business here. Then, we shall turn our attention to an implementation of my design.

The first issue we shall address here is this: What about the seams of the mirrored pyramid? That is, given that the edges of the mirrored pyramid have nonzero thickness, the images from neighboring cameras shall exhibit some artifacts at their common boundaries, and the question is this: Given these image artifacts at the shared image boundaries, how do we composite the individual images into a single visually seamless omni-directional image. The next issue we shall address is the effect of using cameras with lenses that have non-pinhole apertures, instead of using cameras with pinholes, on which we based our design. As we shall see, the two preceding issues are interrelated. Next, we shall discuss the calibration of the geometry, radiometry, and colorimetry of each camera. Such calibration is essential whenever we seek to combine images from multiple cameras into a single coherent image. Having attended to camera calibration, we shall address some miscellaneous practical issues that deserve mention even though they are relatively minor. Finally, I shall describe a first-cut implementation of my design, providing images to illustrate the performance of this implementation.

Throughout our discussion here, without any loss of generality, we shall limit ourselves to the geometry of Figure 9, assuming a right regular mirrored pyramid with a square base and four isosceles sides, each of these sides at 45° to the horizontal.

8.1 Pyramid Seams

Figure 13(a) illustrates the particular horizontal cross section of the mirrored pyramid in Figure 9 that contains the mirror images of the four pinholes. As exaggerated in the figure, in practice, this cross section will have non-pointed corners, these corners lying along the seams of the pyramid in three-dimensional space. Now, a simple strategy we can adopt to avoid viewing the world off the pyramid seams completely is to raise all the actual pinholes simultaneously slightly above their theoretically desirable position we discussed in Section 6. If we do so, the horizontal cross section of the pyramid that contains the mirror reflections of the pinholes shall remain the same as before, but the locations of the mirror images of the pinholes shall no longer be coincident. Rather, they shall be at the corners of a square, as illustrated in Figure 13(b), half the diagonal of this square being equal to the distance by which we raised the pinholes. Thus, we can easily make this square large enough such that each individual camera's useful horizontal field of view off the mirrored pyramid is 90°, excluding the pyramid seams altogether, the four horizontal fields adding up to 360°, as we desire.

For the geometry of Figure 9, then, we shall have four images that are uncorrupted by boundary artifacts possibly requiring complex postprocessing. We shall pay a dual price for these uncorrupted images: One, we shall not have an omni-directional view from a single viewpoint, and, two, we shall not see the world within two orthogonal sheets, each of whose thickness is of the order of the thickness of the pyramid seams we are avoiding. We are assuming pinhole apertures here implicitly, variation from which we shall address in Section 8.2.

At this point, you might think, “Well, so we don't have a single viewpoint after all, just as we did not have it with a fish-eye lens.” The answer is, “Yes, if we adopt the above strategy.” However, as we shall see in the next section, we are not bound to adopt this strategy. Further, note that there is a difference in the orders of magnitude of the spread of the viewpoints here and the spread of the viewpoints when we use a typical fish-eye lens: The pyramid seams can easily be made much less than a tenth of an inch thick each, whereas a fish-eye lens typically has a diameter of several inches. Thus, whereas with my design we shall be able to produce wide-angle images with effectively a pinhole perspective—which ensures, for instance, that straight lines remain straight in a planar image—with a fish-eye lens, we shall be able to provide only very narrow-angle images with effectively a pinhole perspective.

Then, the principal drawback of adopting the strategy described here to avoid the seams is that we shall not view the world within sheets each of whose thickness is of the order of the thickness of the pyramid seams. What we shall gain in return is that we shall avoid **all** image processing at the shared image boundaries. In practice, owing to our ability to make the non-visible sheets a tenth of an inch or less thick quite easily, and the highly redundant nature of most images, I do not expect these non-visible zones to pose a problem in most cases. This claim is substantiated by my experiences with the first-cut implementation I shall describe in Section 8.5.

8.2 Finite Aperture

In the discussion of my design so far, we have limited ourselves to pinhole apertures. In practice, of course, our aperture must be much larger than a pinhole and we must use a lens. Non-pinhole apertures are not necessarily a drawback for my design: They could, in fact, prove very beneficial. As we shall see now, they offer us, at least in principle, a strategy to seam together the individual images while viewing the world from a single viewpoint and without any non-visible zone within a 360° field of view, unlike the strategy of Section 8.1.

As we discussed in Section 2, lenses allow us to use non-pinhole apertures that increase the brightness of the image. Now, we already saw at the end of Section 6 and in Figure 10 that the pinhole geometry is duplicated by a lens as far as rays toward the forward nodal point of the lens are concerned. The question here is, What about other rays that enter the entrance pupil of the lens? As illustrated in Figure 14, depending on the sizes and locations of the mirror images of the entrance pupils of our camera lenses relative to the sizes and locations of our pyramid seams, our strategy of the preceding section to avoid looking at the seams would either work unchanged, or might require the cameras to be moved up slightly higher than originally planned, thus increasing the thicknesses of the non-visible sheets. This higher location would be necessary so that light off the mirrored pyramid from each viewed point can enter the complete entrance pupil, rather than reach just the forward nodal point and a part of the pupil. To give you a feel for numbers, the focal length of each of our lenses might typically be about a seventh of an inch, and the diameter of its entrance pupil, assumed circular—the ratio of this diameter to the focal length determining the image brightness—might typically be about half that, which is about a fourteenth of an inch. Then, if we assume that the entrance pupil and the forward nodal point are coplanar, to accommodate our non-pinhole apertures we would, in principle, have to increase the thickness of each of our non-visible sheets by $1/\sqrt{2}$ times the diameter of the entrance pupil, which turns out here to be about a twentieth of an inch. Of course, we can always design or choose lenses with entrance pupils whose shapes (perhaps, elliptical) and locations (as far in front of the forward nodal point as possible) minimize the extent to which they increase the thicknesses of the non-visible sheets.

What is interesting here is that even though light from an object point might not be able to make it off the mirror to a lens's forward nodal point, light from that object point might make it through the entrance pupil anyway. Then, in principle, we could adjust the sizes, locations, and shapes of our entrance pupils relative to the forward nodal points such that even when all the forward nodal points are virtually colocated, we collect enough light from points within our formerly non-visible zones collectively in adjacent cameras so as to effectively have no non-visible zone. I believe that this can be done, but will be difficult in practice owing to the various unavoidable imperfections and variations in the reflectance and geometry of the mirrored pyramid, especially along its seams. I have not had a chance to investigate this possibility further yet. If we do succeed, the price we shall pay for the avoidance of all non-visible zones in this fashion—and for simultaneously viewing the world effectively from a single viewpoint—is that we shall now have to combine our four images along their shared boundaries to produce a composite image, rather than just slap these images together.

8.3 Camera Calibration

Recall from Section 2, that an image from a particular viewpoint is a representation of light rays toward that point, each ray with a particular orientation, brightness, and color. For our purposes, a convenient representation of such rays forming an omni-directional image is on a cylinder centered at the image's viewpoint, as shown in Figure 15. Note that, without any loss of generality, we are once again reverting to a strictly pinhole geometry for our discussion.

Now, we would like to know what is the brightness and color of each point imaged onto our imaging cylinder from the viewpoint at the cylinder's center. To do so, we need to discover where in our four images does a ray of known brightness and color and in a known direction toward our viewpoint end up being imaged, and, then, with what brightness and color. In other words, we wish to calibrate the geometry, radiometry, and colorimetry of our omni-directional camera.

One effective strategy to calibrate our omni-directional camera is to rotate the camera in small discrete equal steps about the axis of its pyramid while imaging a fixed vertical column of equispaced white dots on a black background—at each step of the rotation, locating the image positions, brightnesses, and colors of the dots. Actually, we note the brightnesses and colors of the black background too as we step along, because two observations in brightness and color space are necessary for the effective calibration of brightness and color. I illustrate this calibration procedure in Figure 16(a). (This calibration strategy of rotating the camera in small discrete steps is equivalent to, but much easier to implement than, immersing the omni-directional camera in a cylindrical arrangement of white dots on a black background, aligning the axis of the pyramid with the axis of the cylinder.) It is advantageous to actually use elliptical dots whose major axes are vertical, rather than circular dots, as then, for a given dot area, a dot will be less likely to intersect a non-visible zone at any stage of the calibration. Thus, column by column, we build up a mapping from our four images to our representational cylinder—in effect, determining the positions, brightnesses, and colors of rays through points on our representational cylinder toward its central viewpoint, as illustrated in Figure 16(b).

8.4 Miscellanea

Some miscellaneous issues remain to be addressed yet, and let us attend to them quickly before we move on to an implementation of my design.

One, each camera can potentially see itself in the mirror, as illustrated in Figure 17(a), and this limits the camera's useful field of view.

Two, if the camera's self image is not the factor limiting the camera's useful field of view, then it is the vertical extent of the accompanying mirror. For any given size of a mirrored pyramid, we can easily work out the distance of each camera—that is, of the camera's forward nodal point—from the pyramid axis that would maximize our vertical field of view. Our total horizontal field of view, of course, is fixed at 360°. Going through the exercise, which I skip here, we shall arrive at the following result: If the square base of the pyramid has sides of length 2γ and its truncated square top has sides of length 2δ, as illustrated in Figure 17(b), then, assuming that it is the extent of the mirrors that limits our cameras' fields of view, the optimum distance of each camera from the pyramid axis is mathematically $\sqrt{2\delta\gamma}$. This optimum distance does not guarantee us a field of view that is symmetrical about the optical axis of each camera if the camera is pointed vertically upward. If we wish to maximize a symmetrical field of view while keeping our cameras pointed vertically upward, the optimum camera distance from the pyramid axis is $2/(1/\delta + 1/\gamma)$, which is the harmonic mean of δ and γ. In any event, it is clear that we must pay close attention to the usability of the vertex end of the mirrored pyramid if we wish to keep the size of the pyramid small for a given vertical field of view.

Three, our cameras can catch bright overhead or other light from beyond the edges of the mirrored pyramid, and so, in practice, it is necessary that we shade our lenses from extraneous light—just as we shade lenses in normal photography. Such shading is conveniently accomplished here by placing above the pyramid a flat piece of black cardboard that extends from the base of the pyramid without obstructing the useful fields of view of the four cameras off the mirrors, as shown in Figure 17(c).

8.5 An Implementation

A first implementation of my design is up and running. This implementation, although not optimal in any respect, provides an affirmation that my design is practical. In reality, this implementation performs much better than I anticipated: I expected a fair amount of image processing to be necessary to combine the individual images seamlessly, never anticipating that merely slapping together data from the four images would suffice for most purposes, as it seems to.

A mirrored pyramid, of the type sketched in Figure 9, was constructed. A close-up photograph of this pyramid is shown in Figure 18. The individual mirrors of the pyramid are of polished stainless steel, and the seams of the pyramid are easily each less than a millimeter thick, but I am unsure as to how accurate the various angles are. The four cameras used, all visible in Figure 18, are each a Sony XC-999 with a lens of focal length 3.5 millimeters and a nominal horizontal field of view of 100°+. As is to be expected, the lenses all exhibit noticeable barrel distortion (see [Hecht 1987]), but that is not a problem owing to our calibration of the geometries of image capture by the four cameras.

The system architecture is as in Figure 19. Each camera's NTSC output is fed to a frame grabber, currently a Snapper-24, that continuously multiplexes through its four inputs sequentially. The brightness and color of each image are corrected according to a red, a green, and a blue (RGB) look up table (LUT) determined during the calibration phase. The output of the RGB LUT is fed to a digitizer that converts its analog input to a digital output. All the cameras and the digitizer are synchronized by an external source of synchronization (Sync). Each output image from the digitizer is read into the memory of an ordinary personal computer (PC) running a 100 MHz Pentium processor, and, then, this image is mapped onto a cylinder in accordance with the mapping determined during calibration. This cylinder onto which the four image streams are mapped independently is unwrapped and displayed on a regular monitor.

Figure 20 shows sample images before and after their mapping onto a 360° cylindrical image. Figure 20(a) shows four raw (uncorrected) images acquired by the four cameras. Figure 20(b) shows these images assembled together into a 360° view after each image has been corrected individually and independently for geometric, radiometric, and colorimetric distortions. In Figure 20(b), I have marked the boundaries of the individual images for your attention. Also, I have included in the imaged scene a meter rule that I placed less than 1 foot away from the axis of the pyramid to affirm that this design is truly capable of working at all distances, not just for objects that are far away. The camera apertures for this 360° view of my laboratory were set at f/2.0.

I show only one omni-directional image to keep my paper short, but this image is typical. Using completely off-the-shelf components, we have been able to provide a 360° x 50° panoramic-image video stream that we can display at 7.5 panoramic frames per second for a display size of 780-pixels x 130-pixels on a relatively inexpensive PC. Our frame rate is currently limited by our frame grabber, which sequences through the four input video streams. To accomplish even the current panoramic frame rate, given the processing power of a PC, we are forced to cut corners in our mapping of the four input image streams to the single output image stream in two ways. One, we map the brightnesses and colors of all pixels in an image identically (through the RGB LUT), rather than individually. Two, we perform only a closest-pixel integral approximation in our position mapping, rather than interpolate the result from the real positions of observed pixels on the representational cylinder. This integral approximation of positions results in some apparent aliasing. But, on the other hand, the image seams are typically hard to locate, as is evident from Figure 20. Typically, one has to look for the seams, often by waving an object and relying on the sequential frame capture to reveal the seams. The non-visible zones of this implementation are two orthogonal sheets, each about an eighth of an inch thick. Non-visible zones of even this thickness, which can definitely be reduced by a factor of two or more by a more careful implementation, are typically not apparent to viewers unless specifically pointed out with the aid of a meter rule or some other such device.

9. Conclusion

I have presented here the design of what I believe is the first true omni-directional camera. A first implementation of this design suggests that it is practical. Far more sophisticated and superior implementations of this design are possible.

Further, several extensions of the design are possible. I would like to mention three. One, we could create an omni-directional projector, rather than a camera, by replacing each camera with a projector. Two, we could extend the field of view of our omni-directional viewer to encompass a view of the world above by using a hollow mirrored pyramid with an upwardly directed camera sitting within this pyramid, this camera's forward nodal point colocated with the virtual locations of the forward nodal points of the other cameras that are all outside the pyramid. Three, almost the same design as in Figure 9 can allow us to double both the horizontal and vertical resolutions of a camera. In the design of such a high-resolution viewer, we could locate four cameras as in Figure 9, but point these cameras not upward, but toward the vertex of the pyramid so as to view the world below the pyramid. Then, without getting into details, each camera would capture a quarter of the total field of view looking downward. This resolution doubling strategy of dividing a field of view into four quadrants by a mirrored pyramid can, in principle, be applied recursively.

The primary disadvantage of the omni-directional camera I have designed, as is true of alternate designs that are not motion based, is that we cannot gradually zoom the camera into a smaller portion of its field of view. In a normal camera with a zoom lens, we can gradually shrink the field of view of the camera down from its maximum size to its minimum size. In my design, however, if we use zoom lenses on the individual cameras and reduce their fields of view, we shall immediately obtain a disjointed field of view. One possible solution to this drawback is to have available at all times the maximum desired image resolution, perhaps by using an n-sided pyramid where n is large, and then provide a lower-resolution image unless the user asks to zoom in, which then would be possible computationally in a gradual fashion.

Acknowledgments

Without the assistance of Brian Schmult, I would have no implementation to show you at this time. In a span of a few months, Brian engineered the mirror assembly, configured the processing architecture, and wrote a major portion of the accompanying software.

Much of the relevant literature was brought to my attention by others after I began demonstrating the implementation I outlined. Bill Ninke led me to discover the patent by Judd and Smoot that proposes the use of planar mirrors to improve camera resolution. Frank Pirz pointed me to the use of fish-eye lenses for omni-directional electronic imaging, and Freddy Bruckstein pointed me to the use of conical mirrors for the same purpose. Dick Kollarits pointed me to several sources of information on lenses. Kevin Chen brought the article by Halfhill in *Byte* to my attention. Ingrid Carlbom pointed me to the papers by Chen and by McMillan and Bishop in *SIGGRAPH 95*. Mike Naimark brought the book by Meehan to my attention.

Again, after I had a working prototype, several individuals provided useful technical input. Paul Henry corrected a misconception I harbored about the effect of mirror reflection on scene brightness early on, and later suggested several improvements to a draft of this paper. Andy Lippman and Mike Naimark shared with me their experiences with fish-eye lenses for omni-directional viewing. John Denker pointed out to me first that non-pinhole camera apertures could possibly eliminate the non-visible zones extending from the seams of my design.

Finally, I acknowledge with great pleasure the support I received at Bell Laboratories from Paul Henry and Arun Netravali to pursue my ideas unhindered.

References

Anderson, P. I. 1995. "From Telepresence to True Immersive Imaging: Into Real-Life Video—Now!," *Advanced Imaging*, Vol. 10, No. 7 (July), pp. 48, 50, & 95.

Beck, C. 1925. "Apparatus to Photograph the Whole Sky," *Journal of Scientific Instruments*, Vol. 2, No. 4 (January), pp. 135–139.

Bond, W. N. 1922. "A Wide Angle Lens for Cloud Recording," *Philosophical Magazine*, Series 6, Vol. 44, pp. 999–1001.

Burt, P. J., and Adelson, E. H. 1983. "A Multiresolution Spline with Application to Image Mosaics," *ACM Transactions on Graphics*, Vol. 2, No. 4 (October), pp. 217–236.

Chen, S. E. 1995. "QuickTime VR—An Image-Based Approach to Virtual Environment Navigation," in *Computer Graphics: Proceedings of SIGGRAPH 95*, Los Angeles, August, pp. 29–38.

Cox, A. 1974. *Photographic Optics*, Fifteenth Revised Edition, Focal Press, London.

Coxeter, H. S. M. 1969. *Introduction to Geometry*, Second Edition, John Wiley & Sons, New York.

Dixon, D. F. 1989. "Life Before the Chips: Simulating Digital Video Interactive Technology," *Communications of the ACM*, Vol. 32, No. 7 (July), pp. 824–831.

Focal 1969. *The Focal Encyclopedia of Photography*, Revised Desk Edition, Focal Press, London.

Greene, N., and Heckbert, P. S. 1986. "Creating Raster Omnimax Images from Multiple Perspective Views Using the Elliptical Weighted Average Filter," *IEEE Computer Graphics and Applications*, Vol. 6, No. 6 (June), pp. 21–27.

Halfhill, T. R. 1995. "See You Around," *Byte*, Vol. 20, No. 5 (May), pp. 85–90.

Hecht, E. 1987. *Optics*, Second Edition, Addison-Wesley Publishing Co., Reading, Massachusetts.

Hill, R. 1924. "A Lens for Whole Sky Photographs," in *Quarterly Journal of the Royal Meteorological Society*, Vol. 50, pp. 227–235.

Jarvis, R. A., and Byrne, J. C. 1988. "An Automated Guided Vehicle with Map Building and Path Finding Capabilities," in *Robotics Research: The Fourth International Symposium*, R. C. Bolles and B. Roth, Eds., MIT Press, Cambridge, Massachusetts, pp. 497–504.

Judd, T. H., and Smoot, L. S. 1989. "Teleconference Facility with High Resolution Video Display," United States Patent, Patent Number: 4,890,314, Date of Patent: Dec. 26, 1989.

Kingslake, R. 1989. *A History of the Photographic Lens*, Academic Press, San Diego, California.

Krishnan, A., and Ahuja, N. 1993. "Range Estimation from Focus Using a Non-Frontal Imaging Camera," in *Proceedings of the Eleventh National Conference on Artificial Intelligence, AAAI-93*, Washington, D.C., July, pp. 830–835.

Lippman, A. 1980. "Movie-Maps: An Application of the Optical Videodisc to Computer Graphics," in *Computer Graphics: Proceedings of SIGGRAPH'80*, Seattle, Washington, July, pp. 32–42.

McMillan, L., and Bishop, G. 1995. "Plenoptic Modeling: An Image-Based Rendering System," in *Computer Graphics: Proceedings of SIGGRAPH 95*, Los Angeles, August, pp. 39–46.

Meehan, J. 1990. *Panoramic Photography*, Amphoto, New York.

Miyamoto, K. 1964. "Fish Eye Lens," *Journal of the Optical Society of America*, Vol. 54, No. 8 (August), pp. 1060–1061.

Nalwa, V. S. 1993. *A Guided Tour of Computer Vision*, Addison-Wesley Publishing Co., Reading, Massachusetts.

Nikon 1995. *Nikkor Lenses*, Nikon, Melville, New York.

Oh, S. J., and Hall, E. L. 1987. "Guidance of a Mobile Robot Using an Omnidirectional Vision Navigation System," in *Proceedings of the Society of Photo-Optical Instrumentation Engineers, SPIE* Vol. 852, Cambridge, Massachusetts, November, pp. 288–300.

Ripley, G. D. 1989. "DVI—A Digital Multimedia Technology," *Communications of the ACM*, Vol. 32, No. 7 (July), pp. 811–822.

Ronchi, V. 1957. *Optics: The Science of Vision*, translated into English from the Italian and revised by E. Rosen, New York University Press, New York. Republished by Dover Publications, New York.

Rosenblum, N. 1989. *A World History of Photography*, Revised Edition, Abbeville Press, New York.

Szeliski, R. 1994. "Image Mosaicing for Tele-Reality Applications," in *Proceedings of the Second IEEE Workshop on Applications of Computer Vision*, Sarasota, Florida, December, pp. 44–53.

Wood, R. W. 1906. "Fish-Eye Views, and Vision Under Water," *Philosophical Magazine*, Series 6, Vol. 12, pp. 159–162.

Yagi, Y., and Yachida, M. 1991. "Real-Time Generation of Environmental Map and Obstacle Avoidance Using Omnidirectional Image Sensor with Conic Mirror," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR'91*, Maui, Hawaii, June, pp. 160–165.

Yates, R. C. 1974. *Curves and Their Properties*, The National Council of Teachers of Mathematics, Reston, Virginia.

Zheng, J. Y., and Tsuji, S. 1990. "Panoramic Representation of Scenes for Route Understanding," in *Proceedings of the Tenth International Conference on Pattern Recognition*, Vol. 1, Atlantic City, New Jersey, June, pp. 161–167.

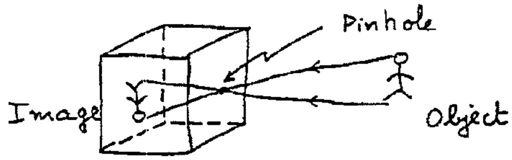

Figure 1 Pinhole camera.

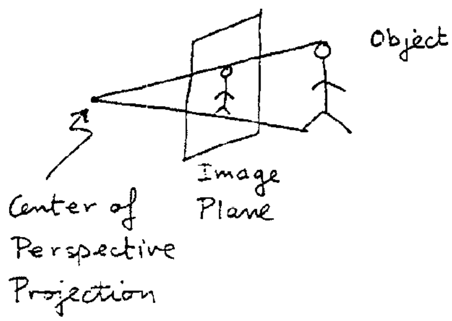

Figure 2 Perspective Projection.

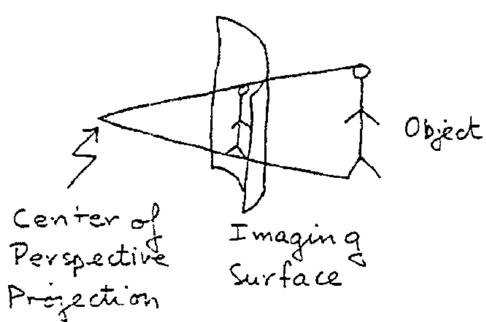

Figure 3 Perspective projection onto a nonplanar surface.

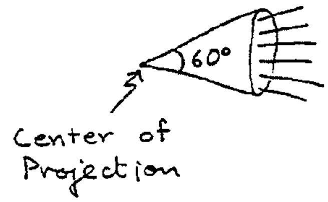

Figure 4(a) Typical camera's field of view.

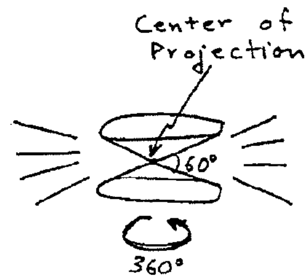

Figure 4(b) Omni-directional camera's field of view.

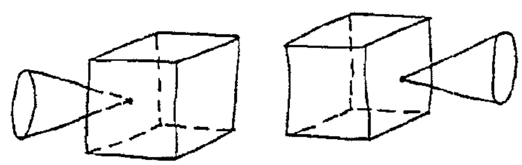

Figure 5 The problem: We must colocate multiple pinholes.

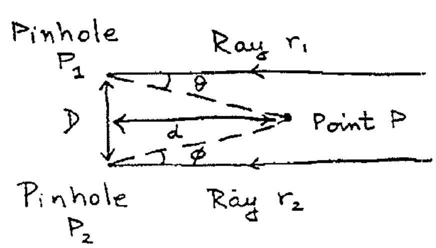

Figure 6 Geometry to compute angular disparity between two pinhole-camera images.

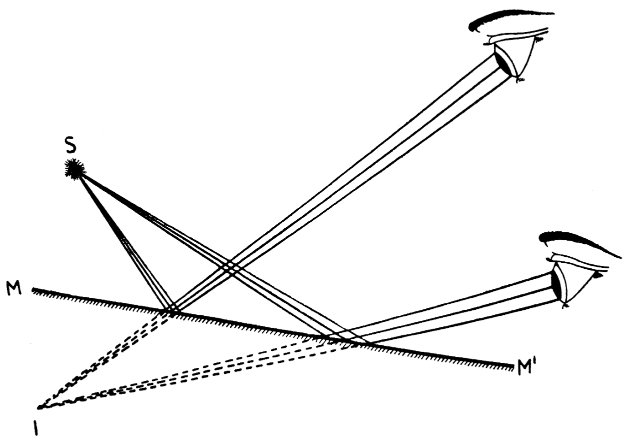

Figure 7 Reflection in a planar mirror, as explained by Kepler (after [Ronchi 1957]).

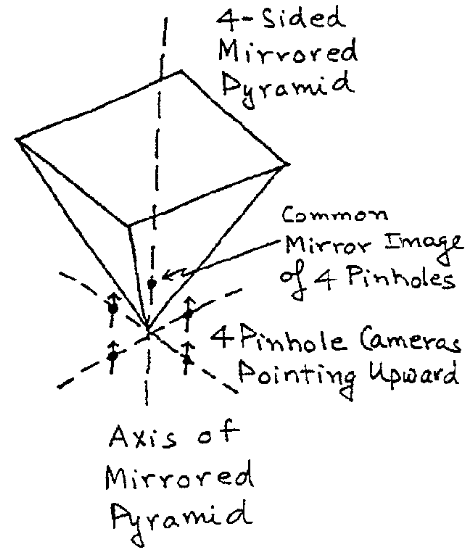

Figure 9 Design of an omni-directional camera.

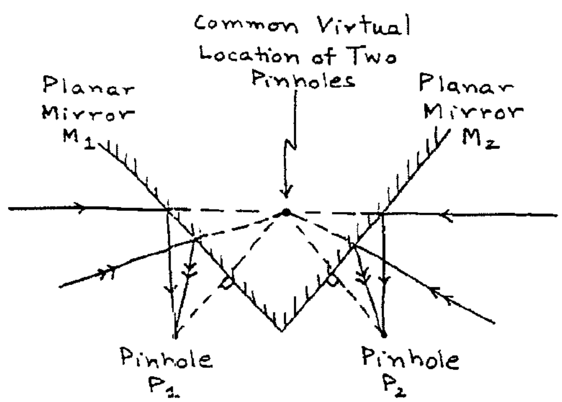

Figure 8 Colocation of two pinholes by two planar mirrors.

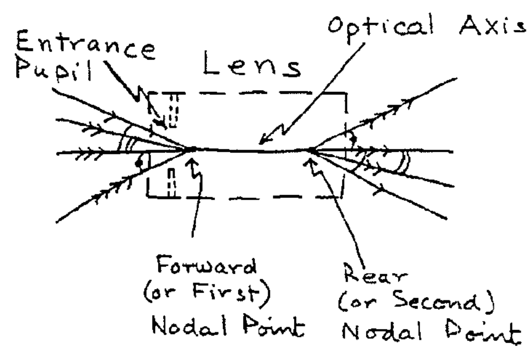

Figure 10 Forward and rear nodal points.

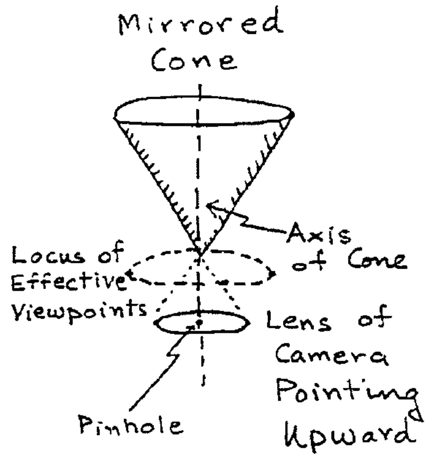

Figure 11 Omni-directional camera using a mirrored cone.

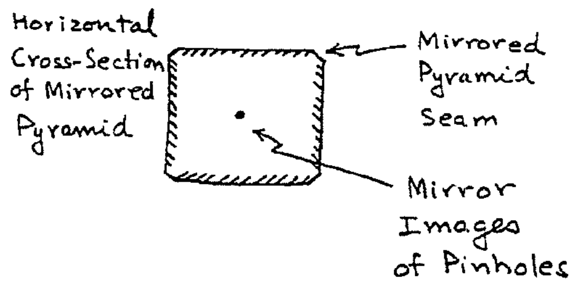

Figure 13(a) Colocated mirror images of pinholes.

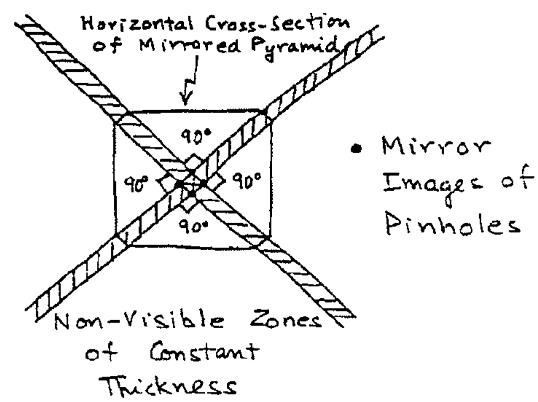

Figure 13(b) Displaced mirror images of pinholes.

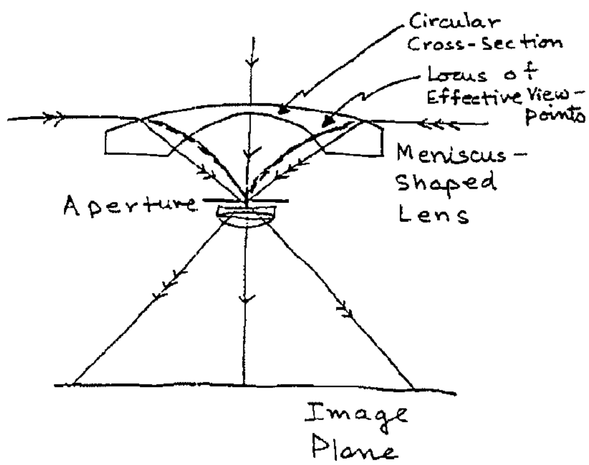

Figure 12(a) Typical fish-eye lens (after [Hill 1924]).

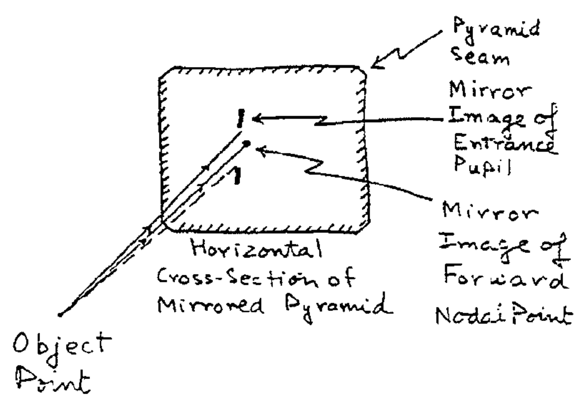

Figure 14 Role of finite aperture in avoidance of pyramid seams.

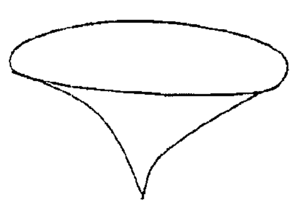

Figure 12(b) Typical locus of effective viewpoints of a fish-eye lens.

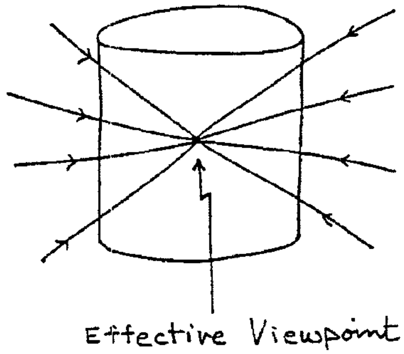

Figure 15 Cylindrical representation of omni-directional image.

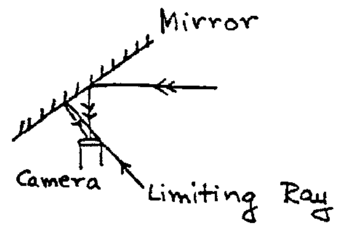

Figure 17(a) Camera's self image limits its field of view.

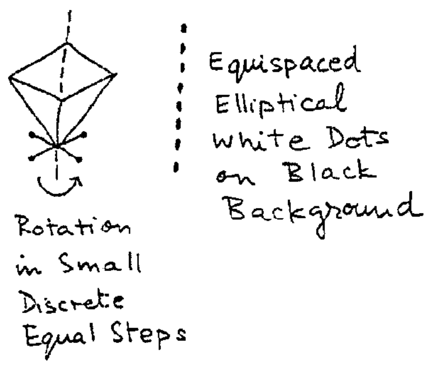

Figure 16(a) Calibration procedure.

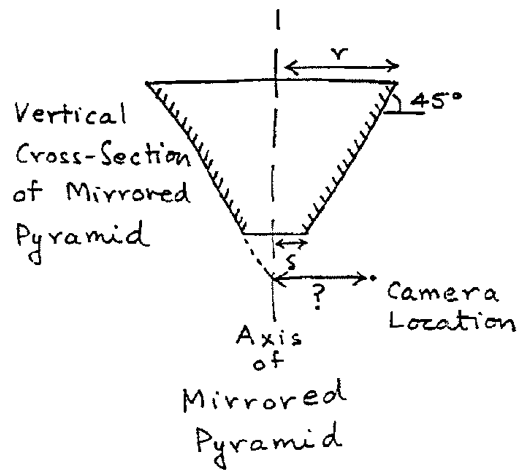

Figure 17(b) Geometry to compute location of camera that maximizes its vertical field of view.

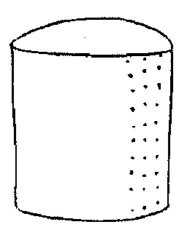

Figure 16(b) Column by column build up of mapping to representational cylinder.

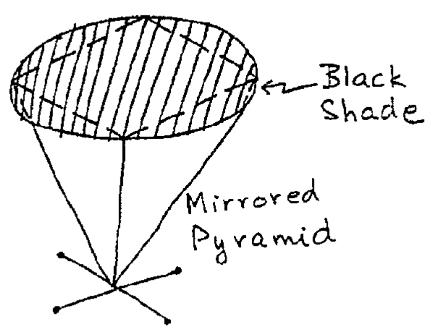

Figure 17(c) Camera shade to block extraneous light.

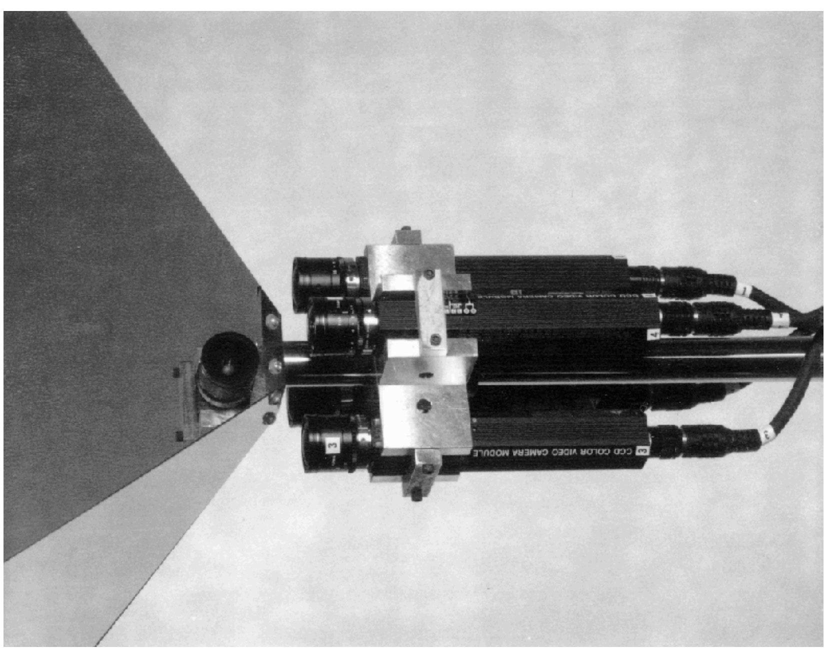

Figure 18 Photograph of implemented mirrored pyramid with attached CCD cameras.

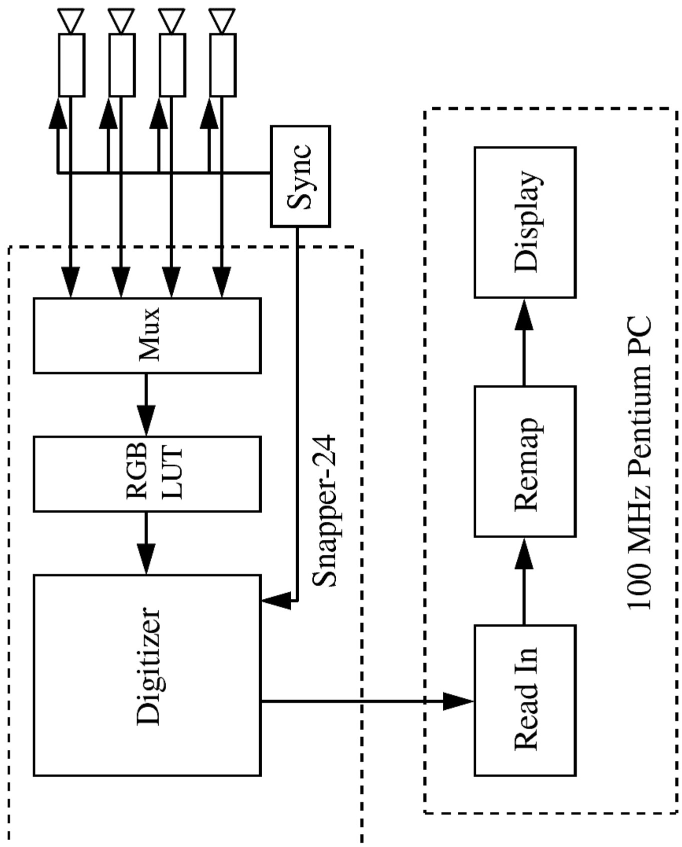

Figure 19 Implemented system architecture.

Figure 20(a) Raw uncorrected images.

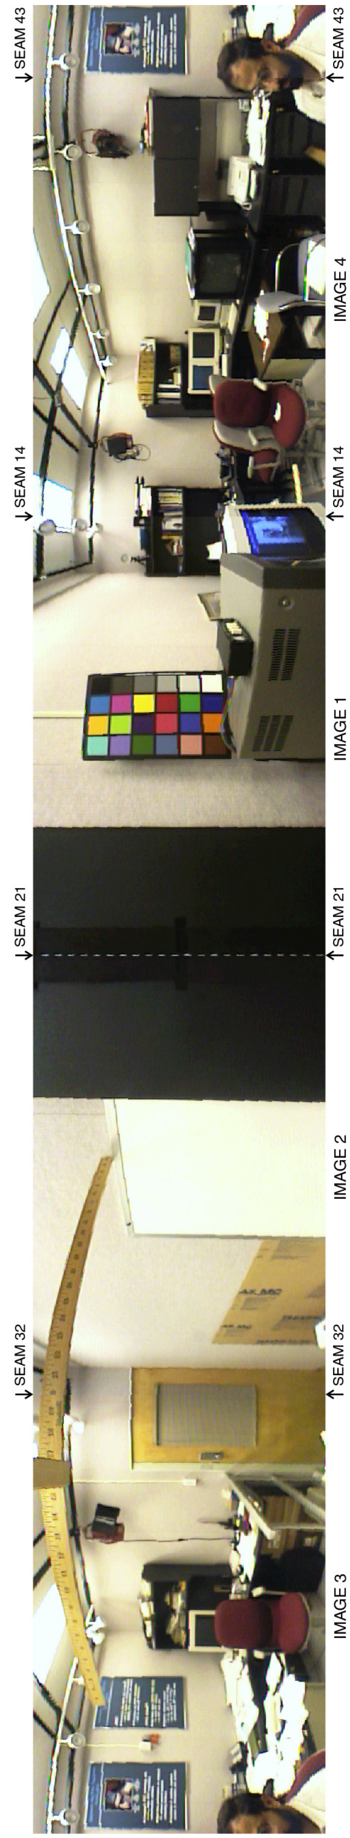

Figure 20(b) Omni-directional 360° cylindrical image unwrapped.