Framework for Designing Catadioptric Projection and Imaging Systems*

Rahul Swaminathan, Shree K. Nayar and Michael D. Grossberg
Department of Computer Science, Columbia University
New York, New York 10027
Email: {srahul, nayar, mdog}@cs.columbia.edu

Abstract

Current projector systems assume a frontal planar display surface. In order to project onto arbitrary surfaces, one must warp the images digitally before projection. As a result, the image quality degrades due to image resampling. In this paper, we propose catadioptric projector systems that use lenses and mirrors, to optically warp images. The warp is specified as a map from points in the image to points on the display surface, called the *image-to-scene* map. The key problem then is to determine the mirror shape that implements this map. Previous methods for mirror design were case-specific and required considerable designer interaction and skill. In contrast, we present a *fully automatic* algorithm to determine the mirror shape for any image-to-scene map. We use splines to model the shape of the mirror and show that the parameters of the spline can be efficiently computed by solving a set of linear equations. Although we focus on the design of projection systems, our framework is directly applicable to the design of imaging systems. We demonstrate the effectiveness of our approach by computing the mirror shapes for both catadioptric imaging and projector systems.

1 Optical Image Warping

The past few years have seen a dramatic rise in applications that use both projectors and cameras, in the realm of ubiquitous computing and collaborative work environments. The central idea in these areas has been to seamlessly integrate computational interfaces into the space around a user. To achieve this, it is necessary to be able to use arbitrary surfaces around the user for display and control purposes [21, 19]. Towards this end, projectors must be able to project high quality images on any surface in the scene. Also, cameras used as input or calibration devices must be able to acquire high quality images of the display surfaces. This calls for greater flexibility in the design of both projector and camera systems.

Consider the scenario illustrated in Fig. 1 of a "smart room" with multiple projectors and cameras. Panoramic displays can be constructed using multiple projectors (**P3**, **P4**) with overlapping fields of view [7, 16, 20], or with a single wide field of view projector, such as **P1**. Similarly, occlusion effects can be removed by using multiple projectors and cameras [24]. Instead of compensating for occlusions, it is easier to avoid them altogether by moving the projector (see **P2**) close to display surface. In all these cases the images need to be digitally warped to minimize distortions in their projections. However, digital warping of images reduces their quality by introducing unwanted artifacts due to image re-sampling. More importantly, in such configurations,

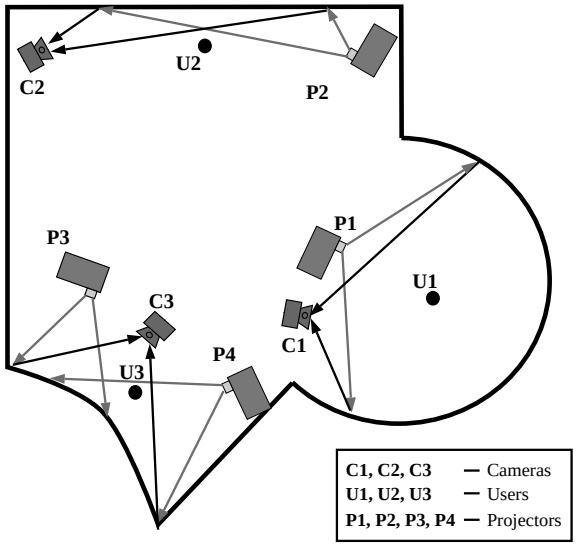

Figure 1: Schematic of a "smart room" using multiple projector-camera systems. Wide angle or panoramic projections are possible using multiple projectors (**P3, P4**) with overlapping fields of view (FOV), or with a single large FOV projector (**P1**). In both cases, the image needs to be warped prior to projection. Another example which requires image warping is when the projector **P2** is moved close to the display surface to avoid occlusions by the user. Typical digital warping techniques degrade image quality due to re-sampling issues. Furthermore, for certain applications, the limited FOV of the projectors makes it impossible to project onto the entire display surface as required. In contrast, catadioptric systems can be used to enhance the FOV as well as provide optical warping thus maintaining image quality without the use of computational resources.

*This research was supported by the DARPA HID Program under Contract No. N00014-00-1-0916.

the limited FOV of the projectors makes it impossible to project onto the entire display surface.

In this work, we propose to perform the image warping optically within the projector or imaging system. Optical warping preserves image quality without any computational overhead. In addition, we can also enhance the FOV of most projectors in unconventional ways. The designer of such an optical system need only specify a map between pixels in the image space and their desired locations in the scene. We call this the *image-to-scene* map. The optical components of the projector or imaging system must then implement this map to project (or acquire) images as desired.

Lenses could be crafted to perform the desired optical warping. However, it is not always feasible to fabricate such lenses due to their complex shapes. Furthermore, such lenses introduce strong optical aberrations. In contrast, mirrors do not suffer from such aberrations and hence are widely used for designing catadioptric (combination of lenses and mirrors) systems. Most of the research in catadioptrics within the last decade focussed on developing imaging systems for photography and computer vision (see [22, 28, 29, 6, 3, 18, 1, 25, 8, 13, 2, 9, 4, 17, 10, 15] for examples). In all cases, the catadioptric system is assumed to consist of some known *primary optics* (usually a perspective or telecentric lens) and a mirror. Constraints imposed by the image-to-scene map are used to derive partial differential equations (PDEs) that the mirror must satisfy. Then, analytical solutions for the PDEs are sought that determine the mirror shape. This approach of deriving and solving PDEs for each image-to-scene map is cumbersome and difficult to scale.

It is highly desirable to have one method for deriving mirror shapes for arbitrary image-to-scene mappings. Hicks was the first to pose this general problem [14] using "geometric distributions". He also presented a novel method to test if a mirror that implemented the desired map existed in theory. However, his framework still requires the user to guide the process by imposing constraints (generally PDEs) and selecting case-specific tools to solve for the mirror.

The goal of this paper is to develop a single algorithm that can compute a mirror shape for a catadioptric system that implements any image-to-scene mapping with no human intervention or effort. Towards this goal our paper makes the following key contributions:

- **Linear Computation of Mirror**: At the core of our algorithm is a linear and yet highly flexible and efficient representation of the shape of the mirror that is based on tensor product splines. This allows us to compute the mirror shape for any given image-to-scene mapping using an iterative linear approach.
- **Handles all possible primary optics**: The framework allows the use of primary optical components with arbitrary known projection geometries including orthographic (telecentric lens), perspective, and other generalized models [11, 26].
- **Finds mirror with least image distortion**: As opposed to previous methods, the mirror we compute is guaranteed to minimize image or scene projection errors. This is achieved by formulating the error metric correctly in the image (or scene) space.
- **Simple method to compute the caustic**: Our method can design systems with a single point of projection or a locus of points, called a *caustic* [26]. The caustic completely describes the geometry of the catadioptric system. Our framework also includes a simple technique to compute the caustic of any designed catadioptric system.

We demonstrate the power of our algorithm by computing the mirror shapes and caustics for previously developed catadioptric systems as well as new ones. In the case of previously derived mirrors, we show that our design matches them perfectly. We believe that the algorithm we propose is a powerful tool that is extremely general and throws wide open the space of projector and imaging systems' that can be designed.

2 Image-to-scene Map using Catadioptrics

The flexibility of our framework comes from the fact that the designer can define a map from pixels in the image to points in the scene. We call this the *image-to-scene* map. The catadioptric system we wish to design is assumed to consist of some *known primary optics* and an *unknown mirror*. Given the image-to-scene map and a model for the primary optics, we determine the mirror shape that best implements the map. Throughout this paper we will address the problem as one of designing catadioptric imaging systems for convenience of notation. However, the same analysis applies directly to projector design as well.

2.1 Modeling the Primary Optics

The catadioptric system is assumed to consist of some known *primary optics* and a mirror. For example, the para-catadioptric camera [18] uses a telecentric lens (orthographic projection) along with a parabolic reflector. Our framework accommodates the use of primary optics possessing any projection model: *perspective*, (Fig. 2(a)), *orthographic* (Fig. 2(b)), or the *generalized model* [11, 26] (Fig. 2(c)). Thus, the primary optics could be a simple perspective lens or a complex catadioptric system itself.

As shown in Fig. 2(c)), in the general case a pixel (u, v) in the image possesses a viewpoint $\mathbf{S_1}(u, v)$ and a view-

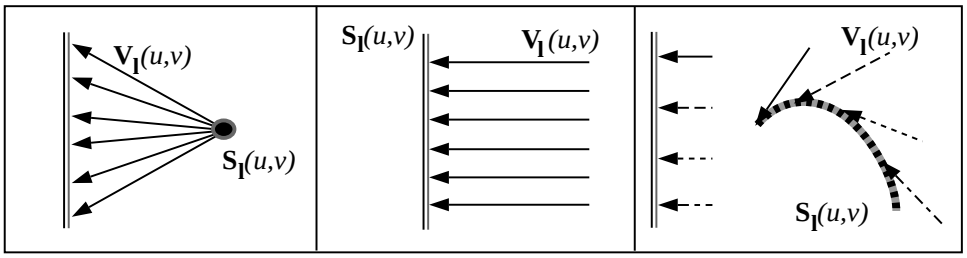

(a) Perspective (b) Orthographic (c) Generalized

Figure 2: The different models of the primary optics for which our method can compute a mirror shape. (a) A perspective imaging model found in most projectors and imaging systems. (b) Orthographic projection, typically obtained with telecentric lenses. (c) The generalized imaging model (caustic model) [11], wherein each pixel can have its own associated viewpoint and viewing direction. This last model is most general and allows the most flexibility in designing catadioptric systems.

ing direction $\mathbf{V_l}(u,v)$. For a perspective lens, the viewpoints $\mathbf{S_l}(u,v)$ collapse to a single point (see Fig. 2(a)).

2.2 Modeling the Mirror Shape

We now describe our model for the mirror shape. For convenience, we express the mirror shape $\mathbf{S_r}$ in terms of the model for the primary optics as:

$$\mathbf{S_r}(u,v) = \mathbf{S_l}(u,v) - D(u,v)\mathbf{V_l}(u,v), \qquad (1)$$

where $D(u,v)$ is the distance of the mirror from the viewpoint surface. We model D using *tensor product splines* in order to facilitate simple and efficient estimation of the mirror shape. We define $D(u,v)$ to be:

$$D(u,v) = \sum_{i=1}^{K_f}\sum_{j=1}^{K_g} c_{i,j} f_i(u) g_j(v), \qquad (2)$$

where $f_i(u)$ and $g_j(v)$ are 1-D spline basis functions, $c_{i,j}$ are the coefficients of the spline model, and $K_f \cdot K_g$ are the number of spline coefficients.

We now have a simple linear model for the mirror surface, which is locally smooth and yet flexible enough to model arbitrary mirror shapes.

2.3 The Image-To-Scene Map

Fig. 3 shows a catadioptric system used to image some known scene. The primary optics can be perspective, orthographic or the generalized projection model.

The user provides a map $\mathcal{M}$ from points (u,v) in the image I to points $\mathcal{M}(u,v)$ in the scene. The mirror surface $\mathbf{S_r}(u,v)$ implements the mapping $\mathcal{M}$ by reflecting each scene point $\mathcal{M}(u,v)$ along the scene ray $\mathbf{V_r}(u,v)$ into the primary optics, where:

$$\mathbf{V_r}(u,v) = \frac{\mathbf{S_r}(u,v) - \mathcal{M}(u,v)}{|\mathbf{S_r}(u,v) - \mathcal{M}(u,v)|}. \qquad (3)$$

This constrains the surface normal of the mirror $\mathbf{N_r}$ as:

$$\mathbf{N_r}(u,v) = \frac{\mathbf{V_l}(u,v) - \mathbf{V_r}(u,v)}{|\mathbf{V_l}(u,v) - \mathbf{V_r}(u,v)|}. \qquad (4)$$

We have now presented all the basic ingredients of our framework for designing general catadioptric systems. These constitute the model for primary optics, the spline-based model for the mirror, and the image-to-scene map.

3 Computing the Mirror Shape

We now present our algorithm to compute the mirror shape for a general catadioptric imaging or projection system. We begin by assuming that the surface normals of the mirror are known. Later, we relax this constraint to present an iterative linear solution for the mirror shape in the general case, where the normals depend on the relative location of the mirror and the scene.

3.1 Normals Known: A Linear Solution

Many catadioptric systems are designed assuming the scene to be very distant (theoretically at infinity) from the mirror. In such cases, the image-to-scene map essentially maps a point in the image (u,v) to the reflected ray direc-

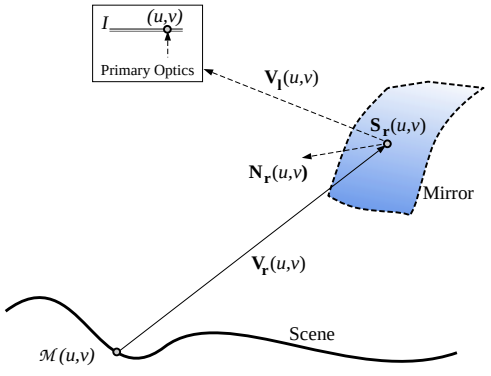

Figure 3: A catadioptric imaging system consisting of some known primary optics and a mirror. In general, a pixel (u,v) in the image maps to the scene point $\mathcal{M}(u,v)$ after reflecting at $\mathbf{S_r}(u,v)$ on the mirror. This forces constraints on the surface normals $\mathbf{N_r}(u,v)$ of the mirror.

tion $\mathbf{V_r}(u,v)$. Since the primary optics are also known, we can derive the required mirror surface normals using Eq.(4) (see Fig.3). This is often used to design imaging systems (see [6, 13, 23], for examples).

The tangent vectors on the mirror surface must be orthogonal to the normal in Eq.(4). This provides two constraints on the mirror shape:

$$\begin{aligned} \frac{\partial \mathbf{S_r}(u,v)}{\partial u} \cdot \mathbf{N_r}(u,v) &= 0, \\ \frac{\partial \mathbf{S_r}(u,v)}{\partial v} \cdot \mathbf{N_r}(u,v) &= 0. \end{aligned} \tag{5}$$

Rearranging the terms and substituting Eq.(1) into Eq.(5), we get:

$$\begin{aligned} (\frac{\partial D}{\partial u}\mathbf{V_l} + D\frac{\partial \mathbf{V_l}}{\partial u}) \cdot \mathbf{N_r} &= \frac{\partial \mathbf{S_l}}{\partial u} \cdot \mathbf{N_r}\,, \\ (\frac{\partial D}{\partial v}\mathbf{V_l} + D\frac{\partial \mathbf{V_l}}{\partial v}) \cdot \mathbf{N_r} &= \frac{\partial \mathbf{S_l}}{\partial v} \cdot \mathbf{N_r}. \end{aligned} \tag{6}$$

Now, substituting Eq.(2) into Eq.(6), we get two new constraints:

$$\left(\mathbf{V_l}\sum_{i,j} c_{i,j} f_i'(u) g_j(v) + \frac{\partial \mathbf{V_l}}{\partial u}\sum_{i,j} c_{i,j} f_i(u) g_j(v)\right)\mathbf{N_r}\,, = \frac{\partial \mathbf{S_l}}{\partial u}\mathbf{N_r}$$

$$\left(\mathbf{V_l}\sum_{i,j} c_{i,j} f_i(u) g_j'(v) + \frac{\partial \mathbf{V_l}}{\partial v}\sum_{i,j} c_{i,j} f_i(u) g_j(v)\right)\mathbf{N_r} = \frac{\partial \mathbf{S_l}}{\partial v}\mathbf{N_r}.$$

where, $f_i'(u)$ and $g_j'(v)$ denote the partial derivatives of $f_i(u)$ and $g_j(v)$, respectively. The above constraints are linear in the spline coefficients $c_{i,j}$ and therefore can be re-written in the form:

$$\mathbf{A} \cdot \mathbf{c} = \mathbf{b}, \tag{7}$$

where, $\mathbf{c}$ represents the set of unknown coefficients $c_{i,j}$ of the spline. Every point (u,v) in the image provides two constraints. Thus, we need at least $\frac{K_f K_g}{2}$ mirror points at which the normals are known. Typically, we solve for the mirror surface at the resolution of the image itself.

Eq.(7) measures the algebraic residue of the dot product between the tangent and surface normal. When a mirror shape exists for the prescribed map, the shape obtained by minimizing the above constraint also minimizes image projection errors. However, if no mirror shape exists, then the computed mirror shape need not correspond to minimum image (or scene) projection error.

Ideally, we should compute the mirror shape that minimizes image (or scene) projection error. Generally speaking, minimizing this metric makes the problem non-linear and unwieldy. However, we circumvent this problem by weighting (with w) the linear constraints (see [27] for details). We finally compute the mirror shape by solving the set of linear weighted constraints for the mirror shape parameters c:

$$(w\mathbf{A})^T(w\mathbf{A}) \cdot \mathbf{c} = (w\mathbf{A})^T w\mathbf{b}, \tag{8}$$

where the weights w are related to the image (or scene) projection errors.

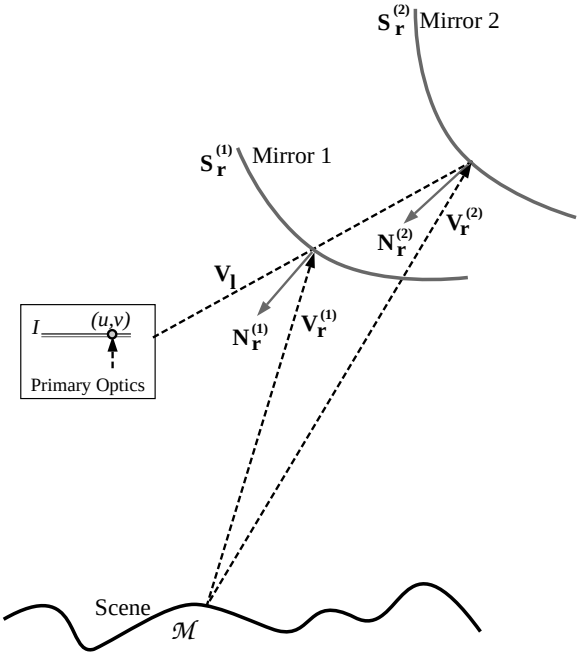

Figure 4: The direction along which a prescribed scene point $\mathcal{M}(u,v)$ is viewed by the sensor depends on the position of the point of reflection on the mirror. As shown here, changing the location from $\mathbf{S}^{(0)}$ to $\mathbf{S}^{(1)}$ alters the reflection direction, thus changing the surface normal.

3.2 The General Iterative Algorithm

In the previous section, we derived linear constraints to compute the mirror shape given a set of surface normals. When the scene is at infinity, this is all we need to compute the mirror shape. We now extend this method for the general case of arbitrary scenes and image-to-scene maps.

In general, the mirror can lie anywhere within the field of view of the primary optics. Fig. 4 shows the mirror at two possible locations, $\mathbf{S_r}^{(1)}$ and $\mathbf{S_r}^{(2)}$, reflecting the scene point $\mathcal{M}(u,v)$ into the primary optics. The direction $\mathbf{V_r}(u,v)$ along which the scene point $\mathcal{M}(u,v)$ is viewed clearly depends on the mirror location. This in turn influences the surface normals on the mirror and hence its shape.

We resolve this cyclical dependency using an iterative

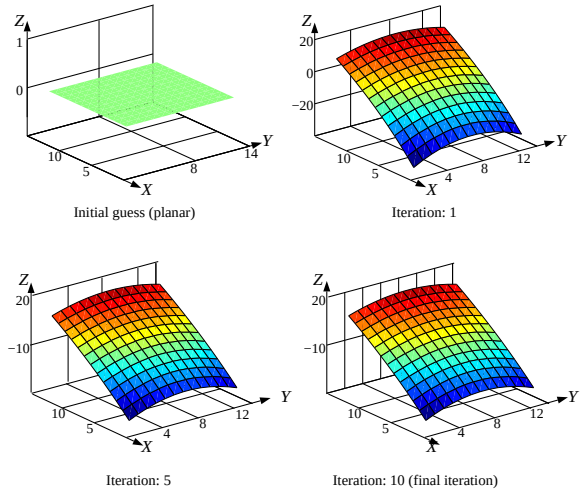

Figure 5: Convergence of the mirror shape designed for the skewed plane projection in Fig. 7(row 3). (a) Initial planar guess for the mirror with facets having different surface normals. (b) After the first iteration, the mirror shape is already close to the final solution. (c) After five iterations, the mirrors changes a negligible amount showing fast convergence. (d) The final mirror shape obtained after 10 iterations.

algorithm. We first approximate the mirror by a set of facets on a plane (say, one for each pixel), whose distance from the primary optics is chosen by the designer. We then estimate the initial set of surface normals using Eqs.(3,4). These normals are used to linearly solve for the mirror shape using Eq.(8). We now iterate as shown in Fig. 6, by using the computed mirror shape in each iteration to obtain a better estimate of the surface normals for the next iteration, until the mirror shape stops changing. Typically, the shape converges within 10 iterations. An example of how the mirror shape evolves from the initial planar guess to its final shape is shown in Fig. 5. In most cases, the mirror shape after the first iteration is already close to its final shape.

4 Caustics of Catadioptric Systems

The catadioptric systems our framework can design are not constrained to have a single point of projection. That is, all the light rays entering or leaving the system need not pass through a single point. In general, a locus of points called a caustic [11, 26] is formed.

Caustics are important as they completely describe the geometry of the catadioptric system. With respect to imaging systems they describe the effective viewpoint locus. For projector systems caustics describe the effective projection model of the catadioptric projector (see Fig.2). In the Appendix, we present a simple method to compute caustics of catadioptric systems using our framework.

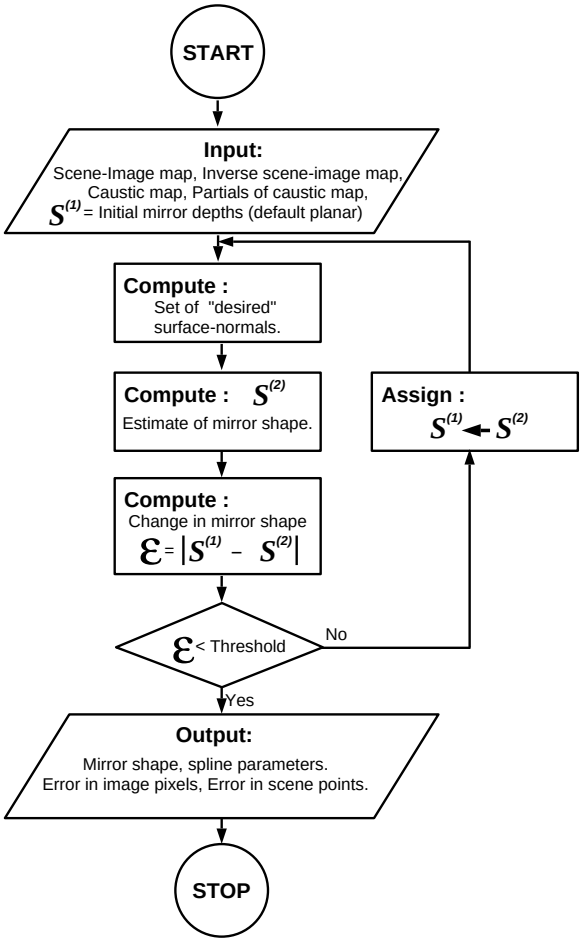

Figure 6: Flow-chart for the spline-based method to compute mirror shapes for general catadioptric imaging systems. The user specifies a map between pixels in the image and corresponding scene points as well as the geometry of the primary optics (shown in Figs. 2(a,b,c)). Using these as inputs, our algorithm computes the required mirror shape automatically.

5 Example Mirror Designs

We now use our algorithm to compute the mirror shapes for four catadioptric systems. These include the single viewpoint para-catadioptric camera [18], the cylindrical panoramic system [12, 23], and two novel projector and imaging system designs. We also present the caustic surfaces for each of the systems we design.

5.1 Single Viewpoint Sensors

We begin by designing a single viewpoint catadioptric imaging system. Such systems can acquire panoramic images from a single virtual viewpoint. In general, only a few mirror shapes and lens combinations yield single viewpoint systems [1]. Therefore, these systems hold a special place within the realm of catadioptric imaging.

We computed the mirror for a para-catadioptric [18]

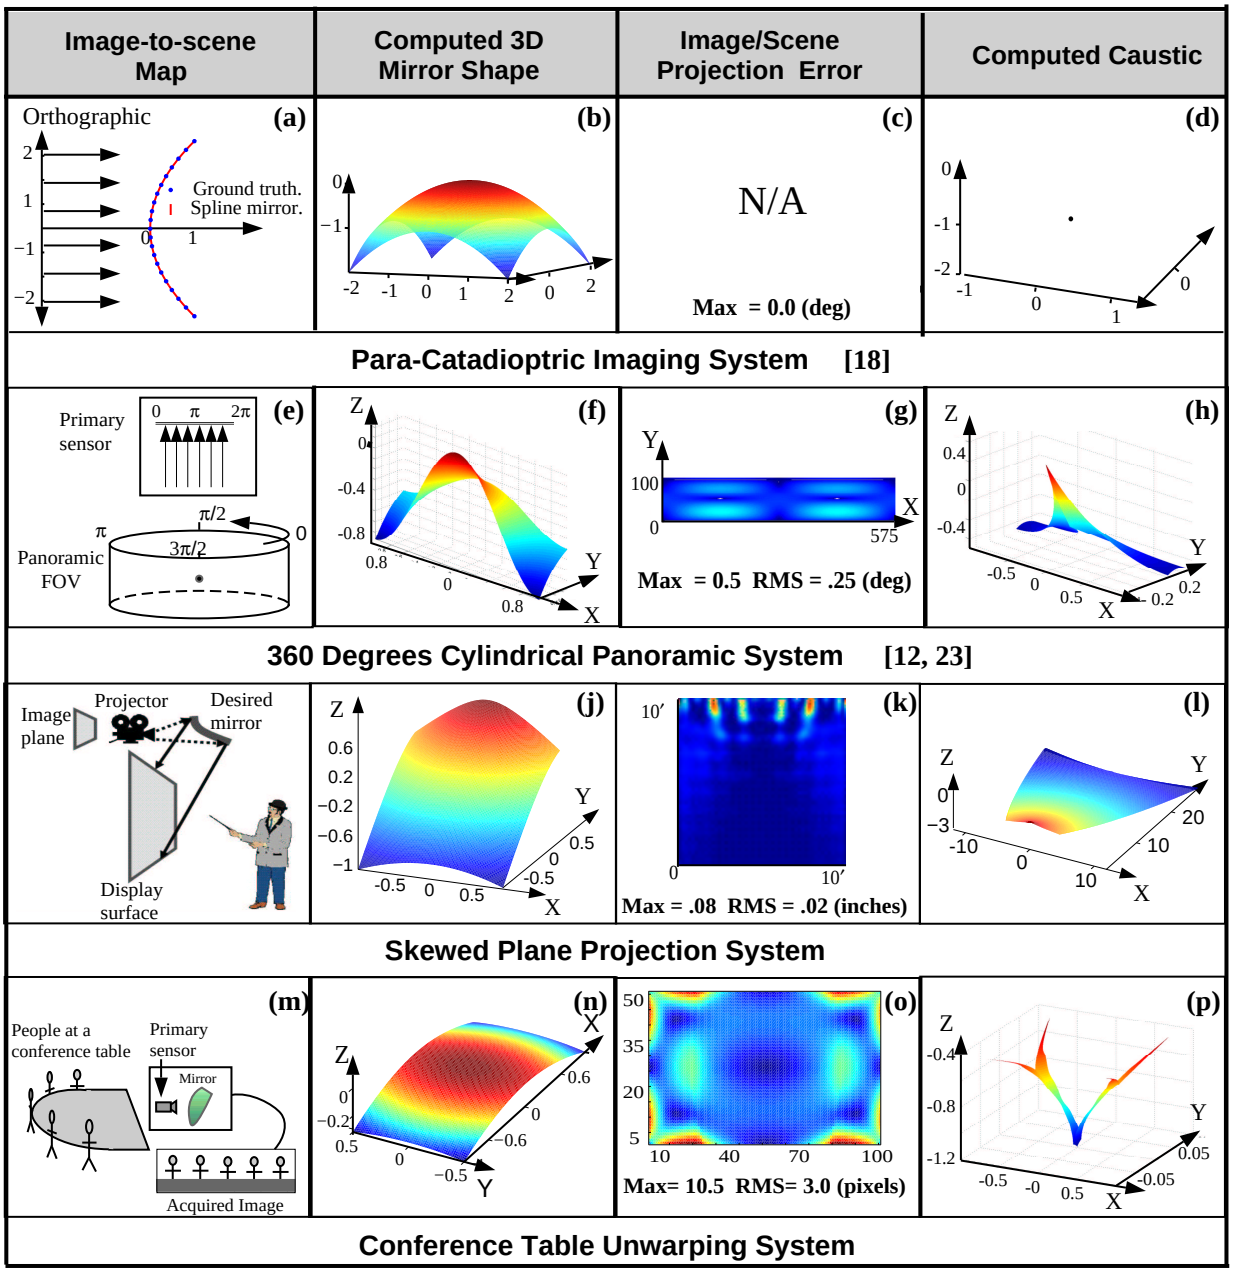

Figure 7: Results of applying our general mirror design method to a wide range of systems. From top to bottom, para-catadioptric imaging system [18], panoramic imaging/projection system [12, 23], skewed-plane projection system and the conference table unwarping system. For all the above systems we present the 3D shapes of the mirrors, the image/scene projection errors and the associated caustic surfaces. The caustic is important as it completely describes the geometry of the system.

imaging system using a telecentric lens for the primary optics (orthographic projection). The theoretical mirror shape for this system is known to be parabolic. The mirror was designed by specifying an image-to-scene map such that all the imaged rays pass through a virtual viewpoint located $1cm$ below the apex of the reflector. As seen in Fig. 7(a – d), the profile of the computed mirror using our method matches precisely with the analytic parabolic profile. The mirror shape is also shown in 3D for visualization. The error in viewing direction at each point in the image is zero. Finally, the caustic surface for this system was computed and found to be a very compact cluster of points (essentially a single viewpoint).

5.2 Panoramic Unwarping Sensor

The para-catadioptric camera designed above acquires an omnidirectional image but needs computational resources to unwarp it into a panorama. Hicks and Srini-

vasan [12, 23] independently presented a design of a catadioptric camera that acquires an optically unwarped panoramic image. Such catadioptric systems are also suitable for projecting directly onto panoramic display surfaces.

In Fig. 7(e – h) we show the setup used to design this system, the computed mirror shape [1], the projection errors, and the associated caustic surface. We again assumed orthographic projection for the primary optics. Every pixel (u, v) is mapped directly to the the required viewing direction $\mathbf{V_r}(u, v)$ of a panorama assuming the scene to be at infinity. Therefore, we can compute the mirror shape in one iteration of our algorithm.

5.3 Skewed Plane Mapping

We now present a catadioptric projector design geared towards avoiding occlusions. We accomplish this by moving the projector close to the display surface. To project an undistorted image on the display surface the desired system must optically warp the image.

Referring to Fig. 7 (i), we note that the display surface lies 1' below and away from the projector and spans a 10×10' square region. The projector's image plane was assumed to be parallel to the XY-plane. The computed mirror shape, its associated scene projection errors and caustic are shown in Fig. 7 (i – l). As expected from the setup, the mirror is symmetric about a single plane. Note that, in spite of the projector being only a foot away from a large screen, the errors in projection are negligible.

5.4 Conference Table Rectification

As our last example we present a new application for imaging systems. Consider the scenario of imaging people seated at an elliptic conference table. We would like to display the acquired image directly, such that the curved edge of the table appears straight. Thus, all the people would appear as if seated along a desk. In general, we might wish to acquire images or video with some predetermined warp, so as to project them directly without the use of any computational resources.

We call the sensor used in the conference table scenario, as the *conference table rectifying sensor*. The setup of the mirror and table for the "rectifying system" are shown in Fig. 7(m – p). The table consists of a semi-circle section of radius $30''$ with two extended straight sides, each $30''$ long. The camera is assumed to lie roughly $30''$ behind this table facing away from the scene into a mirror. The computed mirror shape, the image projection errors and the associate caustic surface are also shown in Fig. 7(row 4). As can be seen the mirror is symmetric about two orthogonal planes. It should also be pointed out that no mirror shape exists that provides the required image-to-scene map. The computed mirror approximates such a map by minimizing image projection errors. It was important to use a weighting scheme discussed in Section 3 to minimize image projection errors.

A point to note is that the same algorithm was used to estimate all the mirror shapes. We could compute single viewpoint as well as non-single viewpoint systems just as easily. Thus, making our method truly flexible and applicable to a large space of mirror design problems.

[1] Note that the mirror shape for this image-to-scene map does not always exist. In particular, the correct aspect ratio for the image needs to be selected in order to estimate the right mirror shape.

6 Conclusions

In this paper we presented a framework to design general catadioptric imaging and projector systems. Such systems consist of some known primary optics and a mirror. The advantage of such systems is that one has immense flexibility in designing novel projection and imaging systems. The user specifies the map between image points and scene points which we called the *image-to-scene* map. Our method then automatically computes the best mirror shape that implements the desired image-to-scene map. By best we mean that mirror shape we compute is guaranteed to produce the least distorted image or scene projection.

A major advantage of our method is that it is flexible enough to be used with all possible projection models for the primary optics, including perspective, orthographic and generalized [11]. Furthermore, the same method can be used to design a very large class of mirror shapes, including rotationally symmetric and asymmetric systems.

We finally presented results by designing many projection and imaging systems. For each designed system we also computed its caustic. Caustics are important as they completely define the geometry of the catadioptric system. For imaging systems they define the viewpoint locus. For projection systems, they represent the effective projection model for the system.

We believe that the algorithm we proposed is a powerful tool that is general and throws wide open the space of projector and imaging systems that can be designed.

References

[1] S. Baker and S. K. Nayar. A Theory of Catadioptric Image Formation. In *Proc. ICCV*, pages 35–42, 1998.

[2] R. Benosman, E. Deforas, and J. Devars. A New Catadioptric Sensor for the Panoramic Vision of Mobile Robots. In *Proc. OMNIVIS*, 2000.

[3] S. Bogner. Introduction to Panoramic Imaging. In *IEEE SMC Conference*, volume 54, pages 3100–3106, 1995.

[4] A.M. Bruckstein and T.J. Richardson. Omniview Cameras with Curved Surface Mirrors. In *Proc. OMNIVIS*, pages 79–84, 2000.

[5] D. G. Burkhard and D. L. Shealy. Flux Density for Ray Propoagation in Gemoetrical Optics. *Journal of the Optical Society of America*, 63(3):299–304, 1973.

[6] J. Chahl and M. Srinivasan. Reflective Surfaces for Panoramic Imaging. *Applied Optics*, 36(31):8275–8285, 1997.

[7] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the cave. In *Proc. SIGGRAPH*, pages 135–142, 1993.

[8] S. Derrien and K. Konolige. Aproximating a Single Viewpoint in Panoramic Imaging Devices. In *International Conference on Robotics and Automation*, pages 3932–3939, 2000.

[9] S. Gächter. Mirror Design for an Omnidirectinal Camera with a Uniform Cylindrical Projection when using the SVAVISCA Sensor. Technical Report CTU-CMP-2001-03, Czech Technical University, 2001.

[10] J Gaspar, C. Decco, J. Okamoto Jr, and J. Santos-Victor. Constant Resolution Omnidirectional Cameras. In *Proc. OMNIVIS*, page 27, 2002.

[11] M. Grossberg and S.K. Nayar. A General Imaging Model and a Method for Finding its Pa rameters. In *Proc. ICCV*, pages 108–115, 2001.

[12] A. Hicks. Differential Methods in Catadioptric Sensor Design with Applications to Panoramic Imaging. Technical report, Drexel University, Computer Science, 2002.

[13] R.A. Hicks and R. Bajcsy. Catadioptric Sensors that Approximate Wide-Angle Perspective Projections. In *Proc. CVPR*, pages I:545–551, 2000.

[14] R.A. Hicks and R.K. Perline. Geometric Distributions for Catadioptric Sensor Design. In *Proc. CVPR*, pages I:584–589, 2001.

[15] R.A Hicks and R.K Perline. Equi-areal Catadioptric Sensors. In *Proc. OMNIVIS*, page 13, 2002.

[16] A. Majumder. Intensity seamlessness in multiprojector multisurface displays. Technical report, University of North Carolina, Chapel Hill, 1999.

[17] F. M. Marchese and D. G. Sorrenti. Mirror Design of a Prescribed Accuracy Omni-directional Vision System. In *Proc. OMNIVIS*, page 136, 2002.

[18] S. K. Nayar. Catadioptric Omnidirectional Cameras. In *Proc. CVPR*, pages 482–488, 1997.

[19] C. Pinhanez, M. Podlaseck, R. Kjeldsen, A. Levas, G. Pingali, and N. Sukaviriya. Ubiquitous Interactive Displays in a Retail Environment. In *Proc. SIGGRAPH*, 2003.

[20] R. Raskar, M.S. Brown, R. Yang, W. Chen, H. Towles, B. Seales, and H. Fuchs. Multi-projector displays using camera-based registration. In *Proc. IEEE Visualization*, pages 161–168, 1999.

[21] R. Raskar and K. Low. Interacting with Spatially Augmented Reality. In *Proc. International Conf. on Virtual Reality, Computer Graphics, Interaction and Visualization in Africa*, 2001.

[22] D. Rees. Panoramic Television viewing System. *United States Patent No.3,505,465*, 1970.

[23] M.V. Srinivasan. New Class of Mirrors for Wide-Angle Imaging. In *Proc. OMNIVIS*, 2003.

[24] R. Sukthankar, T.J. Cham, and G. Sukthankar. Dynamic Shadow Elimination for Multi-Projector Displays. In *Proc. CVPR*, pages II:151–157, 2001.

[25] T. Svoboda, T. Pajdla, and V. Hlaváč. Central Panoramic Cameras: Design and Geometry. In *Proc. of Computer Vision Winter Workshop in Slovenia*, pages 120–133, 1998.

[26] R. Swaminathan, M. D. Grossberg, and S. K. Nayar. Caustics of Catadioptric Cameras. In *Proc. ICCV*, pages II:2–9, 2001.

[27] R. Swaminathan, S. K. Nayar, and M. D. Grossberg. A General Framework to Design Catadioptric Imaging and Projection Systems. Technical Report CUCS-017-03, Columbia Univeristy, 2003.

[28] Y. Yagi and M. Yachida. Real-Time Generation of Environmental Map and Obstacle Avoidance Using Omnidirectional Image Sensor with Conic Mirror. In *Proc. CVPR*, pages 160–165, 1991.

[29] K. Yamazawa, Y. Yagi, and M. Yachida. Omnidirectional Imaging with Hyperboloidal Projection. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1029–1034, 1993.

Appendix: Computing the Caustics

We now present a simple method to compute the caustics of catadioptric systems designed using the proposed spline based method. The use of splines to model the mirror shape also simplifies the estimation of its caustic surface. We begin by deriving the caustic surface for the general imaging system, and then describe ways to compute it numerically.

The caustic can be defined in terms of the reflector surface $\mathbf{S_r}(u, v)$ and the set of incoming light-rays $\mathbf{V_r}(u, v)$ as it lies along $\mathbf{V_r}(u, v)$ at some distance r_v from the point of reflection $\mathbf{S_r}(u, v)$ given by:

$$\mathcal{L}(u, v, r_v) = \mathbf{S_r}(u, v) + r_v \cdot \mathbf{V_r}(u, v). \qquad (9)$$

In order to determine r_v we employ the *Jacobian method* [5] by constraining the determinant of the Jacobian of Eq.(9) to vanish. This gives a quadratic equation in r_v, the solution to which gives us the caustic.

Note that, we do not know the analytic forms for $\mathbf{S_r}(u, v)$ and $\mathbf{V_r}(u, v)$. To compute their partial derivatives, we first fit splines to $\mathbf{S_r}(u, v)$ and $\mathbf{V_r}(u, v)$. Then, we compute the required partial derivatives numerically (see [27] for details). Thus, using the framwork of spline, we can easily determine the caustic of any general projector or imaging system.