HDR Lighting Capture of the Sky and Sun

Thesis by
Jessi Stumpfel

In Partial Fulfillment of the Requirements
for the Degree of
Master of Science

California Institute of Technology
Pasadena, California

2004

Contents

1 Introduction **1**

1.1 Real-world Capture of Natural Illumination 1

1.2 Related Work 2

2 Indirect Capture of Sun Intensity **4**

2.1 Lighting Board 4

2.2 Camera Calibration 7

2.3 Extracting the Lighting Environments 8

2.4 High Dynamic Range Recovery 8

2.5 Fresnel Reflection 10

2.6 Correcting for Sky Polarization 12

2.7 Proberware interface 15

2.7.1 Input Image Formats 15

2.7.2 Instructions 16

3 Direct HDR Capture of the Sky and Sun **19**

3.1 Capturing the Full HDR Range of the Sun and Sky 20

3.2 Camera and ND Filter Calibration 22

3.3 Adaptive HDR Imaging 24

3.4 Acquired Data 25

3.5 Rendering 25

4 Discussion **29**

4.1 Comparison 29

4.2 Error Source Analysis 31

4.3 Performance 32

4.4 Conclusion 33

4.5 Future Work 34

4.6 Acknowledgements 34

Bibliography **35**

Chapter 1

Introduction

In this thesis, two techniques for capturing the sky and sun are presented. The first technique was developed for use in the paper *Estimating Surface Reflectance Properties of a Complex Scene under Captured Natural Illumination* which has been conditionally accepted to ACM Transactions on Graphics [8]. The second technique has been accepted as a paper to Afrigraph 2004 as *Direct HDR Capture of the Sun and Sky* [27]. Sections of this thesis have been paraphrased or directly re-formatted from these publications. Chapter 2 describes the first technique which employs indirect calculations of sun intensity. The second technique directly observes sun intensity through a high dynamic range (HDR) sequence of images and is detailed in Chapter 3. The techniques have different performance advantages for use in inverse global illumination and rendering; these are analyzed in Chapter 4.

1.1 Real-world Capture of Natural Illumination

Advancements in digitizing the geometry, reflectance properties, and illumination of the real world have helped create rendering and virtual worlds that are more realistic than ever before. Datasets of captured real-world illumination such as [7]

have provided a new element of realism in offline rendering [6], as well as real-time [24, 26] rendering. These datasets make use of omnidirectional high dynamic range photography to record indoor and outdoor scenes. Because of the challenges in recording an outdoor scene that includes a directly visible sun, current techniques have not been able capture to this important type of lighting environment.

A traditional lighting environment describes, for every incoming direction, a color and intensity which can then be used for rendering or inverse global illumination. With a lighting environment, captured or synthetic, novel virtual scenes using that lighting can be rendered, such as in [27]. Given a photograph of a scene and a lighting environment taken at the same time, it is also possible to *un-light* the scene obtaining lighting independent material properties of objects in the scene as in [8].

For accurate capture outdoor lighting environments, the upper hemisphere including the sky and sun is of more importance than the ground plane. Capturing the upper 180° of the sky is easily done using a fisheye lens or by photographing a convex mirrored surface.

However, there are two additional challenges in creating an accurate record of outdoor natural illumination: the breadth of the dynamic range, and the absolute brightness of the sun. The sun can be on the order of 100,000 times (or seventeen stops) brighter than the sky and clouds. This is a greater range than can be covered with the practical shutter speeds of most commercially available digital cameras. Furthermore, the absolute intensity of the sun is much brighter than what cameras are designed to capture.

1.2 Related Work

To date, most computer graphics renderings illuminated by sun and sky light have been accomplished using analytic sky models as in [20, 28] and the "gensky" function in [29]. A model specifically designed for realistic rendering of large-scale environ-

ments was presented in [21]; this model used an atmospheric scattering model to compute not only the colors in the sky but also physically-based aerial perspective of surfaces near the ground. A night sky model was developed by [14] where they model the effect of the sun, moon, atmosphere, orbiting dust, stars, and the Milky Way. [33] adapted an analytic sky model so that it could be fit to a small set of sky intensity measurements for an inverse rendering application. Work has also been done to efficiently compute the lighting of outdoor natural illumination on synthetic scenes such as [5]. These models can produce synthetic skies that match ideal conditions well, but they do not model the more visually interesting aspects of cloud formation and motion. Several techniques have been proposed to simulate cloud formation, illumination, and evolution [13, 10, 19]. These techniques produce synthetic 3D clouds that are realistic in appearance, but do not necessarily represent the full range of atmospheric effects encountered in nature, or the specific types of cloud formation found in particular geographic locations.

In the meteorological community, instruments and cameras monitor weather year round at locations worldwide. The goal in this case is to compile weather statistics, such as average cloud cover and precipitation; one such study is presented by [18]. A limitation of their approach is they do not attempt to capture sun intensity, instead using a moving arm to block the view of the sun from the sensor.

Image-based techniques in Chapter 3 to directly image the sky, clouds, and sun are based on the illumination capture process of [6]. Also, in Chapter 3, capture of time-lapse natural illumination measurements throughout the day are accomplished in a manner similar to the high dynamic range video project of [15].

Chapter 2

Indirect Capture of Sun Intensity

The fundamental concepts and techniques described in this chapter were developed and used in [8]. A typical capture image and equipment setup can be seen in Figure 2.1. Additional refinements were made to this setup based on experience obtained processing the lighting captures for [8]. These refinements resulted in a new apparatus, shown in Figure 2.2. Since the original apparatus (Figure 2.1) is described in [8], this chapter focuses on the second generation setup (Figure 2.2). A technique for removing the polarization sensitivity of reflective surfaces used in capturing the sky is also described.

2.1 Lighting Board

For the second generation technique, a EOS-1Ds Canon digital camera with a 100mm macro lens was used to take images of a custom built lighting board as seen in Figure 2.2. A hemisphere shaped mirror placed curved side up and photographed from above reflects the entire sky. An optical grade hemisphere mirror

Figure 2.1: The original lighting capture board (aka. the *Prober*) included a steel ball bearing, a black plastic sphere, and a spray-painted grey diffuse sphere.

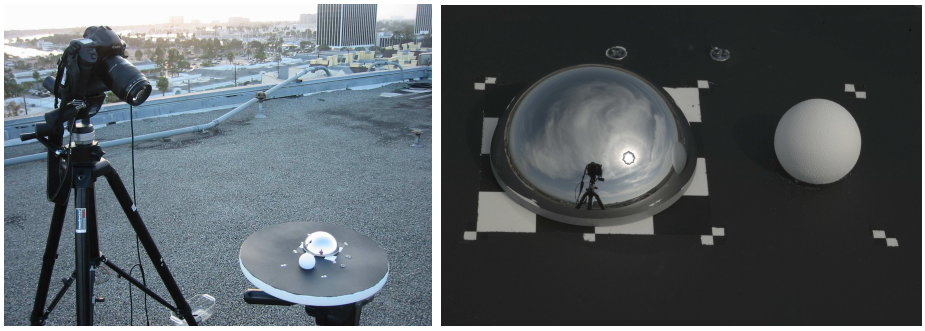

Figure 2.2: The second generation lighting capture board (aka. *Prober Jr.*) uses a hemispherical aluminum mirror and spherical door knob coated with a special diffuse paint.

was constructed using a 75mm diameter, 38.76mm radius plano-convex (PCX) lens from Edmund Industrial Optics (cat no. H45-367). This lens was then coated in protected aluminum by Reynard Corporation to provide greater than 88% reflective mirror surface. Protected aluminum mirrors are first coated with aluminum then protected with silicon monoxide to enhance durability.

The diffuse ball in the setup is a wooden round cabinet knob painted in Munsell diffuse reflectance coating (Edmund Industrial Optics cat no. NT53-696), a paint used for the interior of integrating spheres. This paint was chosen because it is calibrated to be 98% reflective across the visible spectrum, inherently diffuse, and

Figure 2.3: The *Proberware* interface processes lighting capture images to produce lighting environments including the sun intensity. This involves solving for camera position, raytracing the spherical mirror and diffuse ball to unwarp the lighting environment, and finding the sun intensity based on the diffuse ball.

atmospherically stable allowing for repeatable results. However, a cheaper and easier solution of grey or white matte spray paint, carefully applied, can also be used as in [8]. During the capture process a single low dynamic range (LDR) image is taken of the capture board and processed with the C++ application *Proberware* (Figure 2.3). Single images can be acquired in under 3 seconds for the EOS-1Ds camera when recording to a compact flash card and under 6 seconds when downloading the images to a computer. The board can also be filmed with a video camera for real-time lighting capture, though the resolution of most video cameras will produce significantly lower resolution lighting environments.

2.2 Camera Calibration

At the end of each capture session a series of images of a checkerboard calibration object were recorded according to the calibration procedures described in [34]. Using a Matlab toolbox released by Jean-Yves Bouguet [4], these images were used to solve for the camera and lens intrinsic parameters: focal length, principle point, and distortion coefficients. Because these parameters can change with the camera's focus distance, one intrinsic calibration is required each time the setup of the camera and lighting board is changed. Using the calculated intrinsic parameters and the six fiducial markers on the lighting board, the camera position and rotation were recovered [9]. The lens falloff for the 100mm lens was calibrated by Andrew Jones as in [8].

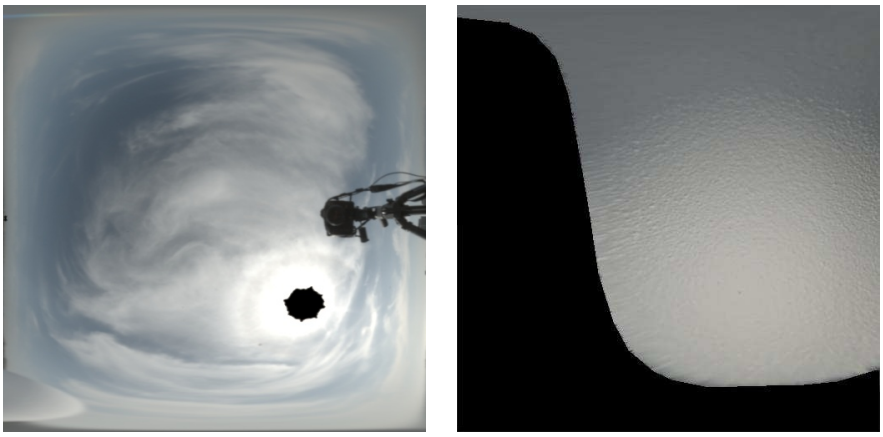

(a) (b)

Figure 2.4: (a) This is a latitude-longitude projection of the upper hemisphere showing the light which has been extracted from the image of the spherical mirror. (b) The same projection as seen from the diffuse ball. Because the camera can only see one side of the diffuse ball, assuming lambertian reflection, not all of the directions in upper hemisphere are visible in the diffuse environment.

2.3 Extracting the Lighting Environments

Given the intrinsic parameters of the camera, recovered camera position, and measured physical positions of the aluminum mirror and diffuse ball on the board, in *Proberware* casts rays from the camera to the spheres and reflects them into the sky. These rays are used to recover a color and intensity for every incident direction to the spheres on the lighting board. Using mapping along the reflected angle for the spherical mirror and mapping along the normal angle for the diffuse sphere, this was used to map the images of the lens and diffuse ball into two latitude-longitude images of the upper hemisphere as seen in Figure 2.4. The latitude-longitude projection simply maps the horizontal and vertical dimensions of the image onto latitude and longitude as on a globe. Notice that camera viewing angle and position do not effect the directions in the lighting environment; the environments are aligned with the orientation of the lighting board in the world.

2.4 High Dynamic Range Recovery

When the sun or bright clouds around the sun are present they usually produce saturated values in a small region of pixels in the mirrored sphere image. The goal of this section is to solve for the missing energy which is not recorded in the LDR image of the spherical mirror.

Given a curved, convex homogenous lambertian surface, such as the diffuse ball, and a distant source of illumination, such as the sky, the color of the surface can be described by a convolution of the lambertian BRDF and the lighting environment [23]. An implementation in latitude-longitude representation is shown in 2.1. $P_{x,y}$ is a pixel the color of the surface in the diffuse image being calculated. $\overrightarrow{S}_{i,j}$ is the unit direction of the source lighting environment represented by pixel i, j and $\overrightarrow{D}_{x,y}$ is the unit normal for the surface being computed. $I_{S_{i,j}}$ is the three channel intensity of pixel (i, j) in the source environment.

$$P_{x,y} = \sum_{i=0}^{sizex} \sum_{j=0}^{sizey} \overrightarrow{S}_{i,j} \cdot \overrightarrow{D}_{x,y} * I_{S_{i,j}} \tag{2.1}$$

An approximation of the intensity of the missing portions of the sky, α, can be reconstructed based on the image of the diffuse ball. To start, the position of the sun is estimated at the centroid of the clipped pixels in the sky reflection. A latitude-longitude image of a sun disk of 0.5323 degrees (the size of the sun viewed from earth) is then created at that position with an intensity of 1 in the three color channels, as seen in Figure 2.5 image S*.

The reflection of the sky in the aluminium lens then is corrected for fresnel reflection, as described in the next section. The diffuse ball environment is divided by .98 to compensate for the paint's reflectivity. We can break down the environmental light (image (S) in Figure 2.5) into light attributable to the non-clipped portions of the sky seen in the mirror image S$'$, plus an unknown intensity α times a unit sun. To solve for α, the equation is converted into the diffuse reflection realm through the use of the diffuse convolution 2.1. Then, since the image D of the diffuse ball can be directly observed, we can solve for α.

Since there are many pixels observed on the diffuse ball, this system is an overdetermined system and solving for sun intensity is accomplished by a least squares solution in each color channel to match the photographed ball.

Typical RMS solutions match the diffuse ball with under 5% percent error. However, this process will not recover values of the saturated pixels near the sun. The missing energy from these pixels will be added to the sun's intensity. If the sun is partially occluded by a cloud, the clipped region may not be always centered around the sun, and inconsistent sun positions could be recovered. This could be ameliorated by fitting a curve to estimated sun positions captured over a day of light or by referring to meteorological predictions of where the sun is at a given time of day.

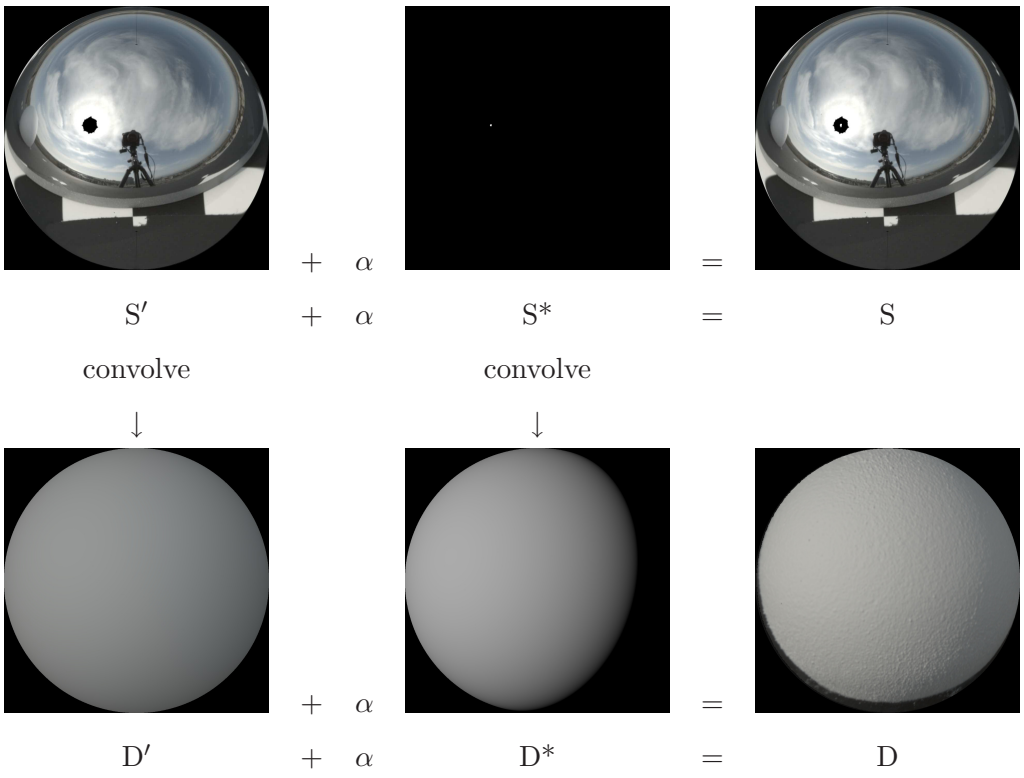

Figure 2.5: This equation describes a method of solving the missing energy caused by the clipped sun pixels. A diffuse-convolved (Equation 2.1) image of the environment as seen in the aluminum lens (D′) plus a convolved unit sun image (D*) multiplied by unknown sun intensity (α) must equal the observed image of the diffuse ball (D).

2.5 Fresnel Reflection

It is possible to accurately describe the reflectance behavior of mirrored spheres using the Fresnel equations (Equation 2.2 and Equation 2.3) [11]. The reflectance is defined as the fraction of the incident light energy that is reflected. n is the index of refraction which is a property of the specific dialectic or metal material. θ_i is the angle of incidence. $R_\perp$ is the reflectance of light polarized perpendicular to the plane of incidence. $R_\parallel$ is the reflectance of light polarized parallel to the plane of incidence.

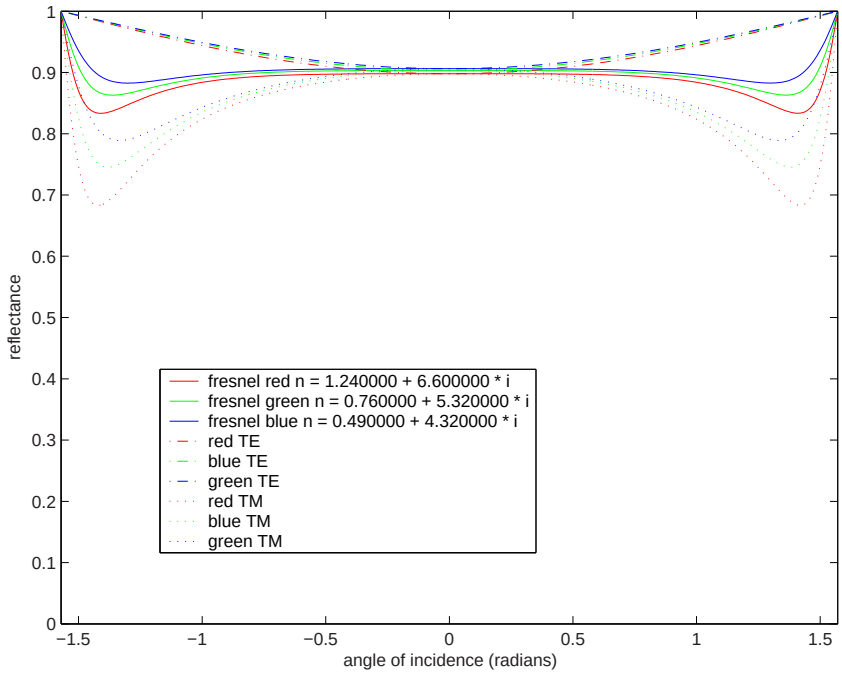

Figure 2.6: This is a plot of the fresnel reflection for a protected aluminum mirror. TE and TM represent light that is polarized perpendicular and parallel, respectively, to the incident plane.

$$R_{\perp} = \left(\frac{cos\theta_i - \sqrt{n^2 - sin^2\theta_i}}{cos\theta_i + \sqrt{n^2 - sin^2\theta_i}} \right)^2 \tag{2.2}$$

$$R_{\parallel} = \left(\frac{-n^2 cos\theta_i + \sqrt{n^2 - sin^2\theta_i}}{n^2 cos\theta_i + \sqrt{n^2 - sin^2\theta_i}} \right)^2 \tag{2.3}$$

According to Fresnel's reflectance equations and the complex index of refraction, n, of protected aluminum [25], the reflection off the aluminum mirror could be darker or brighter depending on the direction and degree of polarization of the incoming light.

0° 20° 40° 60° 80° 100° 120° 140° 160°

Figure 2.7: A canon D30 camera and 50mm lens were used to image the sky with a polarizer. Nine photographs were taken at even increments while the polarizer was rotated 180 degrees. The first photograph was taken with the polarizer perpendicular to the direction of the sun.

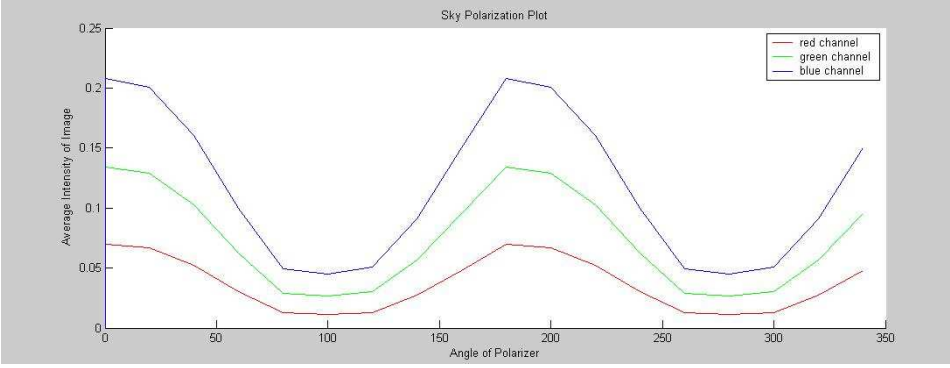

Figure 2.8: The average intensity the three channels of the images in 2.7 were plotted against the angle of the linear polarizer. The plot displays the expected sinusoidal curve appearance.

2.6 Correcting for Sky Polarization

Whenever light reflects, refracts, or scatters, polarization may alter how much light is transmitted. While the error introduced by ignoring polarization effects is often small and visually unnoticeable, it can be an important consideration when capturing the sky accurately, for inverse global illumination or matching lighting environments, especially when rendering objects with specular surfaces.

In these cases, polarization effects are important to consider when capturing natural illumination conditions; the blue color of the sky is the result of Rayleigh scattering off gas molecules and the resulting light can be more than 70% linearly polarized in

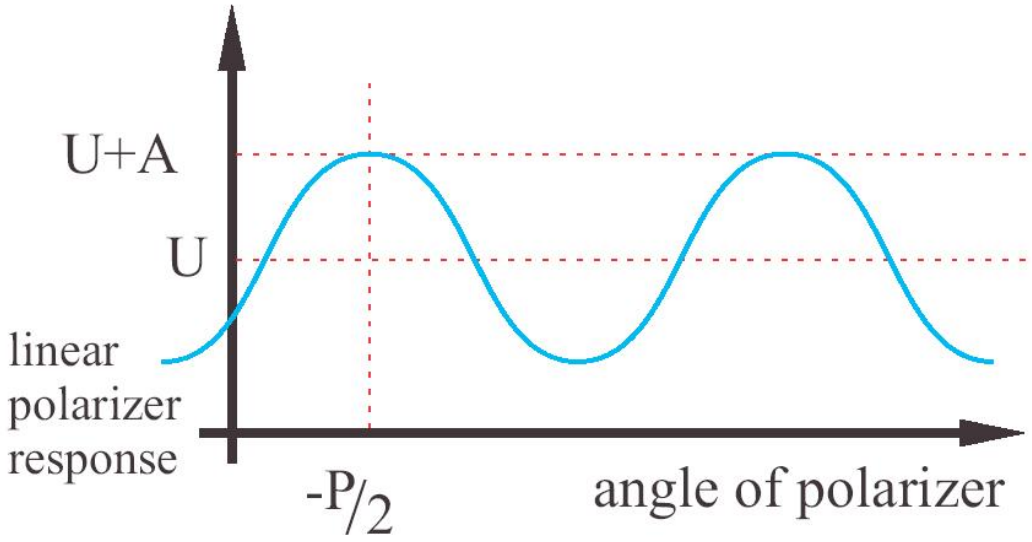

Figure 2.9: The sinusoidal curve of the response of sunlight to a linear polarizer. The direction of polarization is -P/2, the amount of unpolarized light is U-A, and A is amount of linearly polarized light

the region perpendicular to the direction of the sun. This effect is easily observed with a polarizer on a clear day, as in Figure 2.7. The polarization of the sky is further complicated by the presence of clouds and dust in the atmosphere. These larger particles exhibit a different type of scattering known as Mie scattering, which is largely unpolarized.

In order to observe and correct for the effect of the varying polarization across the sky, the incoming polarization is captured at each pixel. Since the sky is partially polarized, the light at each point in the sky can be characterized as a mixture of linearly polarized light and unpolarized light [17] [11].

The response of sky light (I) to a linear polarizer at angle θ, as seen in Figure 2.8, can be represented as a sinusoidal curve with three variables (Equation 2.4), as in Figure 2.9: polarization orientation (P), average response (U), and the amplitude

of response variation (A). In all cases, the frequency is fixed at two.

$$I(\theta) = U + Acos(2\theta + P) \tag{2.4}$$

It is possible to solve for U, P, and A given at least three samples [32], $I_{0...n}$, of a polarizer at three known orientations, $\theta_{0...n}$, by finding the least squares solution to the linear system in Equation 2.5, and substituting into Equations 2.6 and 2.7.

$$\begin{bmatrix} 1 & sin(2\theta_0) & cos(2\theta_0) \\ 1 & sin(2\theta_1) & cos(2\theta_1) \\ \vdots & \vdots & \vdots \\ 1 & sin(2\theta_n) & cos(2\theta_n) \end{bmatrix} \begin{bmatrix} U \\ A_1 \\ A_2 \end{bmatrix} = \begin{bmatrix} I_0 \\ I_1 \\ \vdots \\ I_n \end{bmatrix} \tag{2.5}$$

$$A = \sqrt{A_1^2 + A_2^2} \tag{2.6}$$

$$P = \arctan\left(\frac{-A_2}{A_1}\right) \tag{2.7}$$

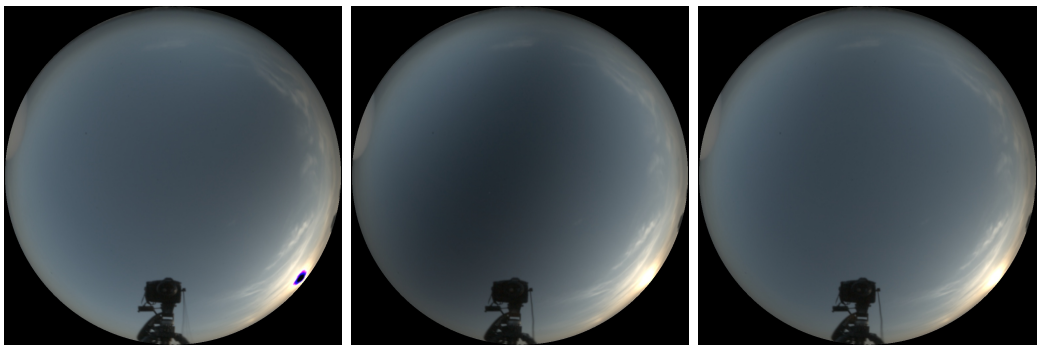

Figure 2.10: Three images have been warped from the image of the aluminum mirror to imitate a 180° fisheye lens looking straight up. They are taken with a polarizer over the camera lens at 3 orientations ($\theta = 0°$, 90°, and 135°) and reveal the polarization of the light entering the lens.

Using the method described in Equations 2.5, 2.6 and 2.7, the polarization of the

light coming into the camera was solved for. The intensity and polarization of the light coming from the sky was calculated by inverting the Fresnel reflection. For each ray reflecting off the aluminum mirror to the camera, the solution for the polarized and unpolarized light was projected onto directions that are parallel and perpendicular to the plane of incidence. In [Equation 2.8] ϕ is the angle between the direction of maximum polarization and the direction perpendicular to the reflection's incident plane. $I_{\parallel}$ and $I_{\perp}$ represent the intensity of light in the parallel and perpendicular directions.

$$I_{\parallel} = \frac{U-A}{2} + Acos^2\phi \qquad I_{\perp} = \frac{U-A}{2} + Acos^2\phi \tag{2.8}$$

These can then be scaled by the inverse of the Fresnel reflectance for the corresponding directions, producing polarized intensity values. Figure 2.11 shows a corrected sky image taking into account the polarizing effect of the mirror, the corrected image is as much as 5% brighter in the region perpendicular to the sun direction than the lighting environment assuming unpolarized light. It should be noted that for less reflective mirrors, such as the steel ball bearing used in [8], this effect is increased because of the increased difference between the parallel and perpendicular Fresnel cases.

2.7 Proberware interface

This is user documentation of software seen in Figure 2.3.

2.7.1 Input Image Formats

Proberware can read .hdr, .pfm, and .fim high dynamic range images. Proberware expects saturated pixel as non-finite values (#INF). It assumes that images are radiometrically linear. Also, if processing an entire directory, Proberware requires

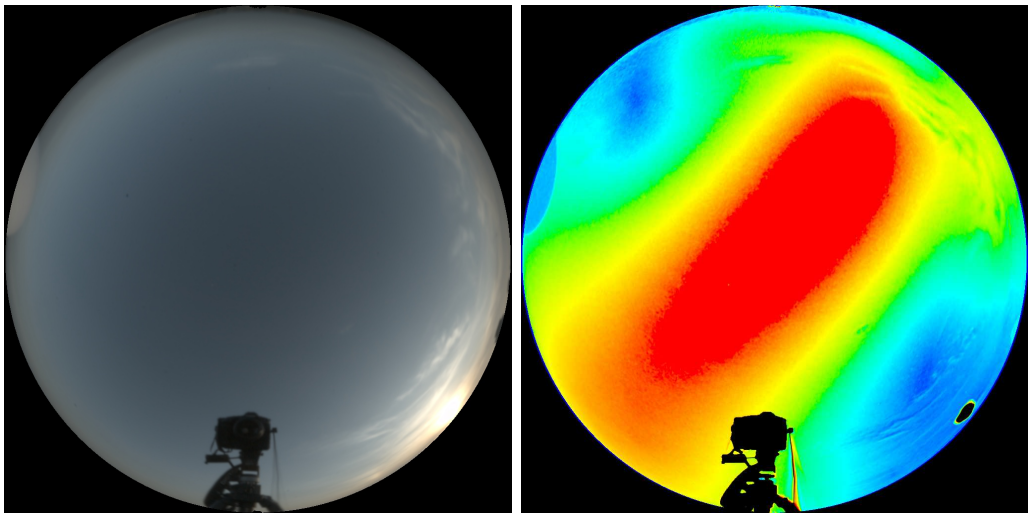

Figure 2.11: (left) This sky image has been created taking into account Fresnel effects of the spherical mirror and the captured polarization of the sky. (right) This is a false color mapping of captured sky polarization. The red color represents skylight that is 40% polarized and the blue color represents unpolarized light. The red band of max polarization is 90 ° from the sun; atmospheric conditions such as dust, pollution, and clouds cause the light to be less polarized toward the horizon.

the last four digits before the file extension to be a sequential number (for example probe_0123.hdr, probe_0124.hdr, probe_0125.hdr, etc.).

2.7.2 Instructions

Steps to processing an individual lighting environment:

1. Press **Load Camera Calibration** and select the proper camera calibration .m file.

2. Press **Load Capture Image** button to load the capture image. Hold and drag right mouse button to pan the image. Hold right mouse button and use middle mouse wheel to zoom

3. Mark fiducial markers in the image with points by selecting the point number

with the numeric keyboard and then clicking with the left mouse button. It should be noted that capture images should be seen with the aluminum mirror on the right side, diffuse on the left. The fiducial markers are numbered as follows 0 for top right, 1 for top middle, 2 for top left, 3 for bottom right, 4 for bottom middle, 5 for bottom left.

4. Press **Save Points**

5. Press **Find Camera**, once to initially estimate camera position and at least once more to optimize camera position. The average error (projected fiducial markers to marked fiducial markers) will be shown on the top of the window. Ideally this error would be below 1.0. If the number remains higher than this it could be due to error in the camera calibration, fiducial markers, board measurements, and camera position/orientation optimization.

6. Press **Find Spheres**

7. Press **Show Sphere Pixels** so Proberware displays the pixels that it considers sphere pixels. If the pixels shown are different from the sphere pixels the physical positions of the sphere can be changed to help minimize the error.

8. Press **Show Tweak Dialog** to change the positions of the spheres and re-run **Find Spheres**.

9. Press **Output Lighting Environments** to output the lighting environments. The final output of Proberware is the radiance scene file (probename.rad) and the Arnold scene file (probename.ass). Both of these scene files refer to the lighting environment anglemap (probename_anglemap.hdr) also output after pressing this button.

Alternatively if an entire directory of lighting envs is to be processed:

1. Press **Load HDR Image** to load the first image of the sequence.

2. Mark and save the points on the first capture image as described above.

3. Press **Find Camera** to perform an initial estimate of the camera position.

4. Press **Process Entire Directory** and select that same first image. It will process and output lighting envs for this image and the rest of the sequence, finding the fiducial markers based on where it marked them in the previous capture. It also initially estimates the camera to be where it was in the previous capture.

To minimize repeated calculations Proberware saves the marked points (probename.pts), camera (probename._camera.txt), rotations to axis align the environment to the Parthenon (currentRotations.txt), and values entered into the Tweak Dialog (currentOffsets.txt) . If these files are present Proberware will read and use this information, if you want Proberware to recalculate the camera, Parthenon rotation, or tweaked values, delete these files and run Proberware again. The currentRotations.txt and currentOffsets.txt are used for an entire sequence if **Process Entire Directory** has been pressed.

Chapter 3

Direct HDR Capture of the Sky and Sun

This chapter presents a different technique for capturing the full dynamic range of natural illumination environments including the sun and sky, which has presented a challenge for traditional high dynamic range photography processes. Through careful selection of exposure times, aperture, and neutral density filters, it was found that this full range can be covered in seven exposures with a standard digital camera. There are particular calibration issues such as lens vignetting, infrared sensitivity, and spectral transmission of neutral density filters which must be addressed. The results are demonstrated by showing time-lapse renderings of a complex scene illuminated by a high-resolution, high dynamic range natural illumination environments.

In this paper, these challenges are overcome by using a combination of varying shutter speed, aperture, and a properly chosen neutral density filter. With optimized settings, the full dynamic range of the sky is captured using at most seven exposures using a camera with 12-bit linear-response sensor.

To demonstrate these techniques, several full days of natural outdoor illumination

Figure 3.1: The camera with 180° fisheye lens and laptop on the roof of the USC's Institute for Creative Technologies building (left). This is a typical assembled light probe (right) taken with the fisheye lens and neutral density filter spanning the 17 stops of the sky and sun.

were captured and used to render an outdoor scene under time-lapse illumination conditions. Finally, an adaptive capture technique is discussed that requires fewer exposures when the sun is not visible.

3.1 Capturing the Full HDR Range of the Sun and Sky

A roof with an unobstructed view of the horizon in all directions was accessible for this research. The entire sky could be imaged from this vantage point by setting up a Canon EOS-1Ds camera pointing up, equipped with a 8mm Sigma fisheye lens [Figure 3.1]. While the Canon camera provides automatic exposure bracketing, using this to adjust the exposure between 1 second and 1/8000 of a second yields only 14 stops (a factor of 16384). Shutter speed times longer than 2 seconds are impractical because of the motion of the clouds. As shutter speeds increase past 2 seconds noise levels in the image increase as well. This is insufficient to capture the

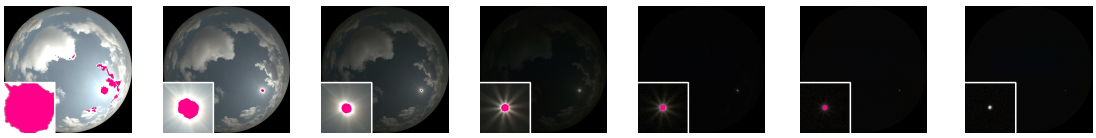

f/4, 1s f/4, $\frac{1}{4}$s f/4, $\frac{1}{32}$s f/16, $\frac{1}{16}$s f/16, $\frac{1}{125}$s f/16, $\frac{1}{1000}$s f/16, $\frac{1}{8000}$s

Figure 3.2: This is a typical HDR sequence spanning the 17 stops of the sky and sun in 7 exposures.

17 stop range (131,072x) of a typical sunny sky.

A library provided by Canon [1] allowed the camera to be controlled from a laptop, while downloading the images to the laptop's disk. Andreas Wenger and Andrew Jones wrote a program to adjust the aperture and exposure to the required settings, acquiring an image sequence detailed in Table 3.1 and 3.2. This sequence spanned the needed 17 stops in a manner sufficient for reconstructing the HDR image, and required approximately 50 seconds to acquire and download. These images are at most 3 stops apart which is within the range covered by the camera's sensor, thus there are no gaps in the dynamic range between images.

Aperture	Exposure Time	Effective Exposure
f/4	1 s	0 stops
f/4	1/4 s	-2 stops
f/4	1/32 s	-5 stops
f/16	1/16 s	-8 stops
f/16	1/125 s	-11 stops
f/16	1/1000 s	-14 stops
f/16	1/8000 s	-17 stops

Table 3.1: This table shows the camera settings used to capture the full dynamic range of the sky.

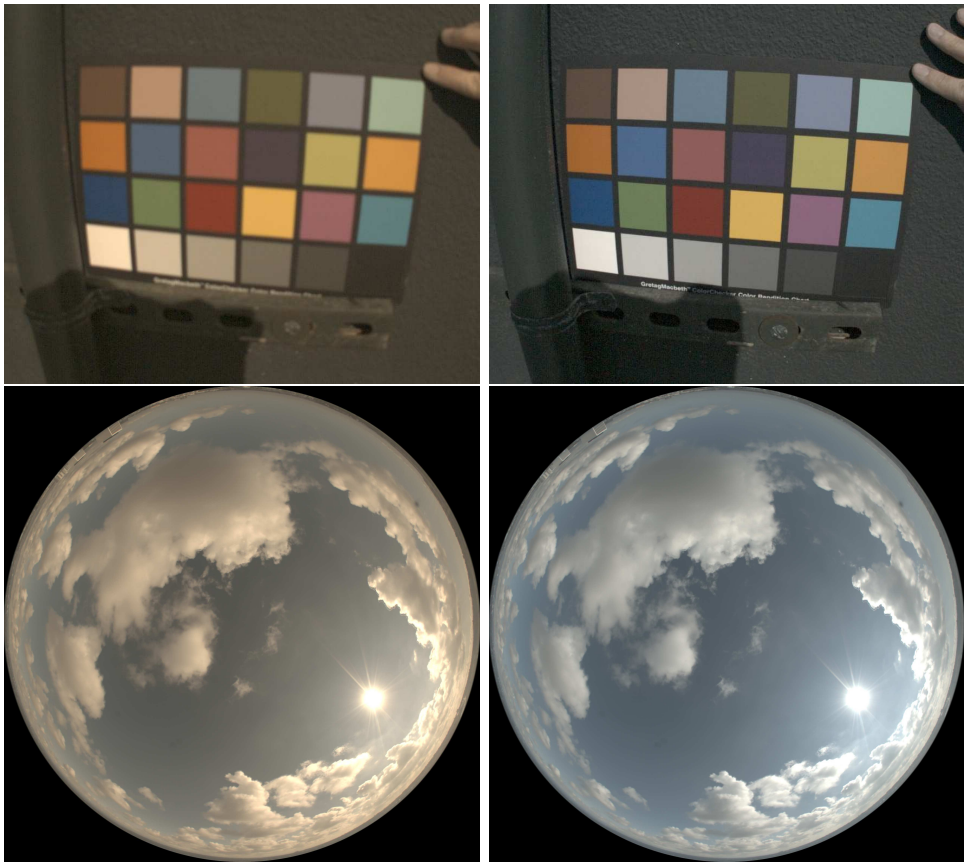

Figure 3.3: The neutral density filter, while reducing the amount of light, introduces a chromatic change as well. To correct for this, a color chart was photographed with and without an ND filter. These images were then used to solve for a color correction matrix in a program written by Andrew Jones.

3.2 Camera and ND Filter Calibration

To achieve accurate results, it was necessary to calibrate the fisheye lens geometrically and photometrically. Geometric fisheye calibration was accomplished (by Andrew Jones) by marking scene correspondences while rotating the camera around the average nodal point. For the Sigma 8mm fisheye lens a perfect fisheye model was assumed and the lens center and radius were optimized. A more sophisticated geometric model could be based on a more accurate measurement of the lens construction [16], or by fitting an additional function to the lens distortion. Also, there

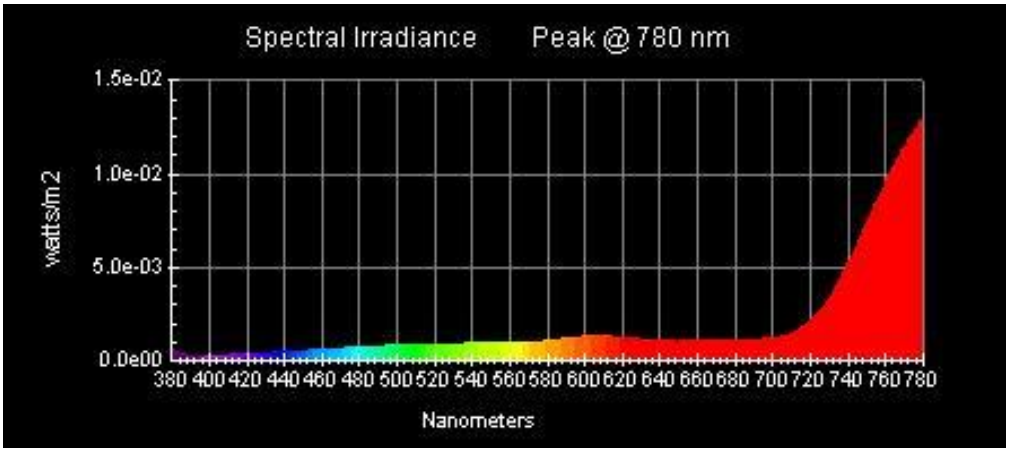

Figure 3.4: The spectral calibration of the 3.0 ND Kodak WRATTEN Filter, showing a gentle increase in the visible range from blue to red. A significant amount of IR light is transmitted.

is significant radial intensity falloff and vignetting associated with the fisheye lens. This was accounted for by Andrew Jones as in [27].

Given the EOS-1Ds camera's image sensor sensitivity, at least a 3.0 Neutral Density (ND) filter was necessary to bring the sun within range. This filter in theory allows 0.1% of the light to pass through. ND filters are usually placed in front of the lens, but none of the standard filters fit the curved front surface of the Sigma 8mm fisheye lens. This was solved by cutting a piece of a Kodak WRATTEN 3.0 Neutral Density gelatin filter and placing it internally in between the lens and the camera.

Figure 3.3 shows an uncorrected image of the sky with reddish-brown clouds taken with the ND filter. Contrary to their name, neutral density filters are often not chromatically neutral. A correction for the chromatic shift in the visible portion of the ND filter spectrum was estimated in a program written by Andrew Jones in the following manner: a Macbeth color chart was photographed in sunlight with and without the ND filter [Figure 3.3], then solved for a 3x3 linear color transform was calculated that most closely matched the color samples between the images.

To further investigate the effect of the ND filters, using a Photo Research PR-650

spectroradiometer the ND filter's spectral transmissivity was measured [Figure 3.4]. Not only is the filter more transmissive in the red region than the blue, but it also displays a large amount of light transmitted in the near IR region above 720nm which ND filters are not designed to block.

This can be potentially problematic, as many digital cameras have a significant sensitivity to near IR light. For example, a Sony VX1000 clearly responds to light above 720nm. Fortunately, experiments performed by Andreas Wenger showed that the Canon EOS-1Ds does not have a significant near IR response. However, as sun light has a very strong near IR component, any camera that exhibits IR response should use IR-cutoff filters.

It should be noted that the fisheye 8mm Sigma lens was tested for the effects of polarized light by placing a polarizing gel over a diffuse light and photographing it at different points in the camera image. It was found that at the edge of the image the light varied by 25% as the direction of the polarizer was rotated. This effect would introduce some error to the lighting environments. The fisheye lens is polarization independent in the center of the lens where the light with the highest degree of polarization is imaged.

3.3 Adaptive HDR Imaging

On the first pre-dawn morning on the roof, it was discovered that the f/4 1 second exposure was not long enough to capture the appearance of dark clouds before and during dawn. In general, over the course of the day the brightness of the sky can change drastically with sun position and weather.

To address this problem, as well as to increase the speed of the capture and reduce wear on our camera, Andreas Wenger wrote a program to adaptively change the exposure and aperture when capturing an HDR sequence [27]. This program analyzes images on the fly to determine when pixels in the sky are underexposed or

saturated.

Using the adaptive capture program an HDR image could be captured every 40 seconds with sequences consisting of three to seven pictures. The adaptive capture also handled the highly increasing and decreasing light levels during the entire day from pre-dawn to post-dusk.

3.4 Acquired Data

Using these techniques several days of light have been captured with different weather conditions. Figure 3.5 shows example lighting environments. These datasets have allowed the visualization the full dynamic range of natural illumination over the course of the day, as seen in Figure 3.6. Exploring the practical issues of capturing the full high dynamic range of the sky using a fisheye lens with a neutral density filter allows the natural dynamics of the sky and sun to be represented faithfully in our renderings. Several captured HDR days of light can provide versatility and flexibility for choosing lighting for novel outdoor renderings. The datasets taken for this thesis will be made available at: www.ict.usc.edu/graphics/SkyProbes/.

3.5 Rendering

The captured lighting presented here could be used to render any outdoor scene, providing realistic illumination and a sky backdrop. However, due to the extreme contrast between the sun and sky, standard monte-carlo illumination algorithms will produce a large amount of noise, as they will have difficulty sampling the small disc of the sun reliably and sufficiently.

Algorithms have been developed for intelligent sampling of the environment with priority given to high intensity regions such as in [3]. In this case, a simpler solution was utilized for approximating the sun as a directional light, while using traditional

sampling techniques for the remainder of the sky. Energy from pixels in the lighting environment that are above a high threshold are removed from the lighting environment and approximated with a directional light, placed at the centroid of the removed pixels.

This method was applied to the whole-day datasets and used in conjunction with the Arnold global illumination system written by Marcos Fajardo. Several frames are shown in Figure 4.3. For these renderings, a program written by Chris Tchou, called ZIGI, was used to handle the Parthenon data and make calls to Arnold.

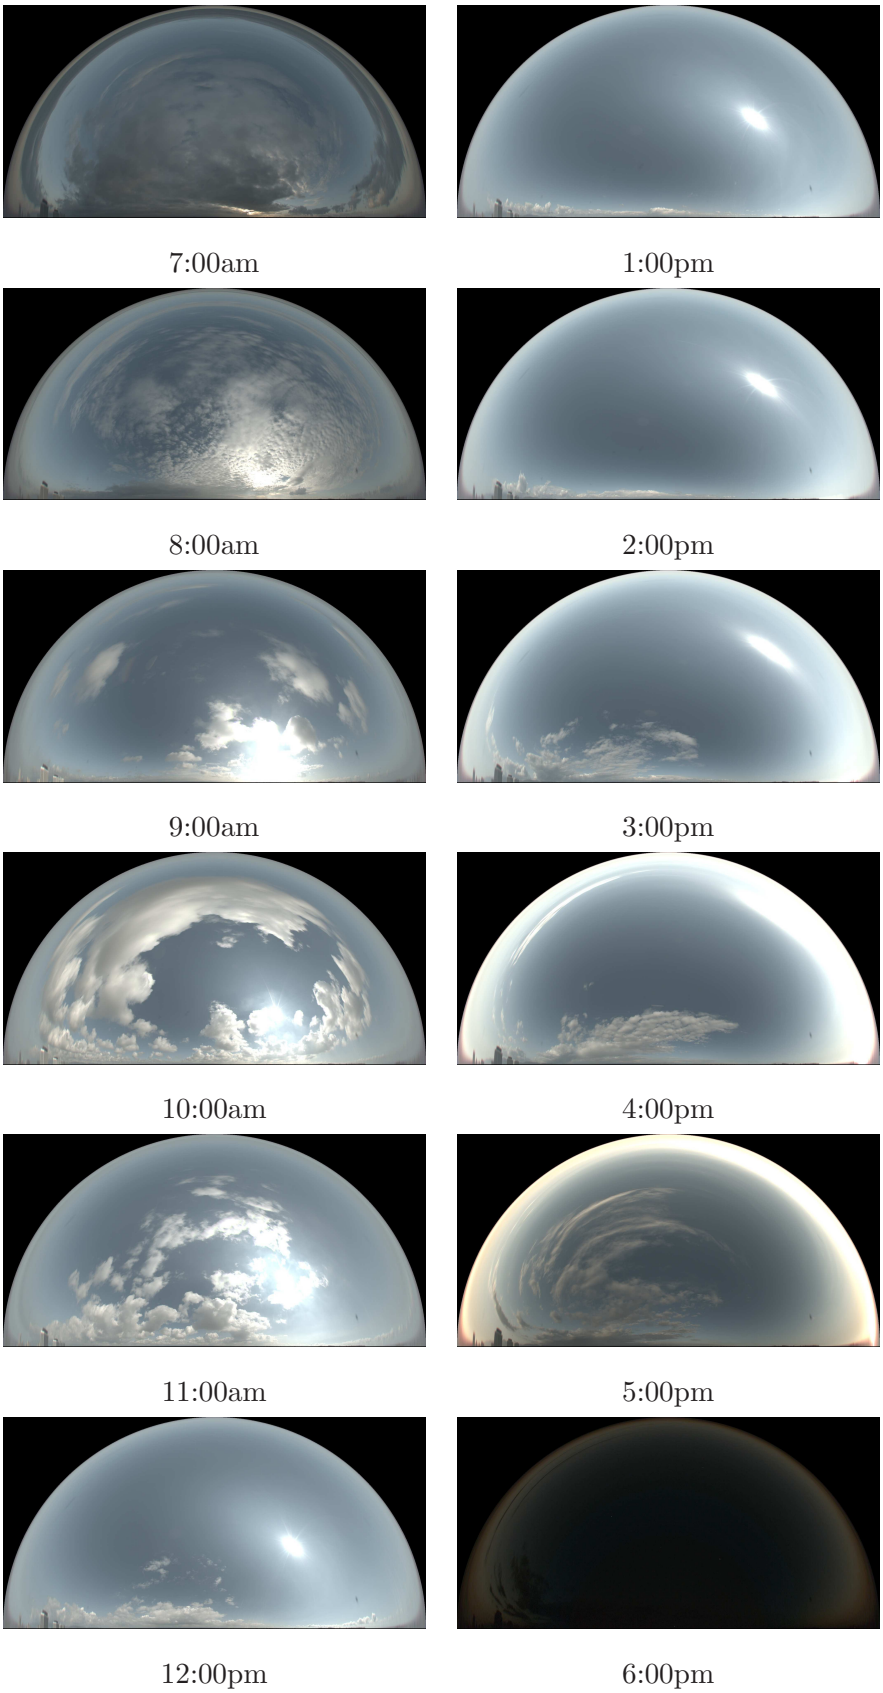

Figure 3.5: This is a sampling of the 700 captured lighting environments taken each minute of the sky in Marina Del Rey, CA on February 23, 2004.

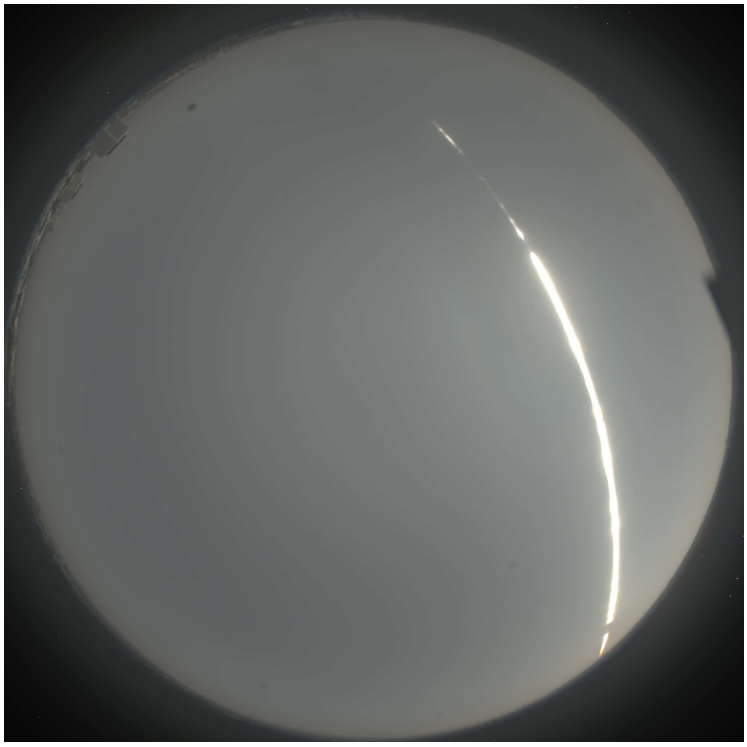

February 19, 2004

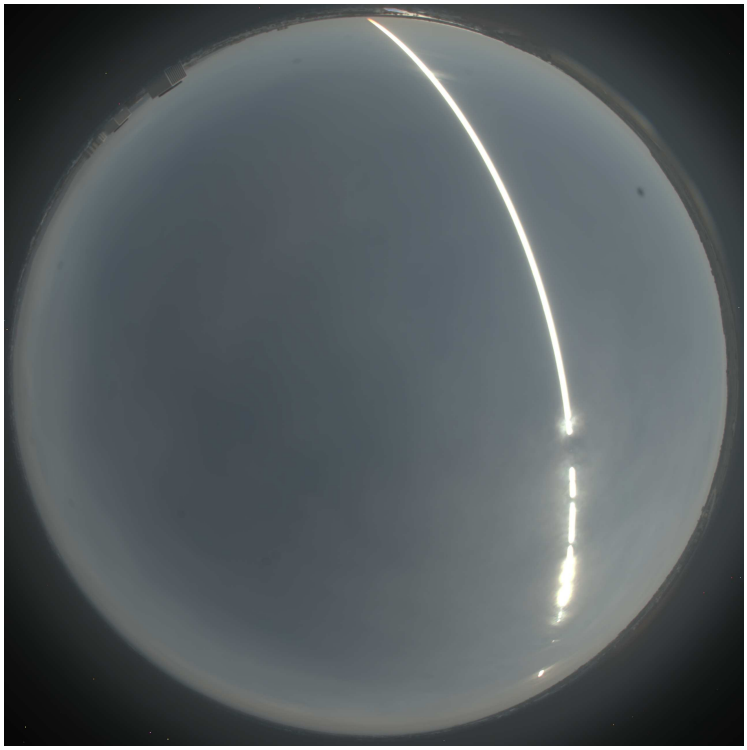

February 23, 2004

Figure 3.6: These images are the average of about 700 full dynamic range images taken of the sky at one minute intervals over two separate days of light. The sun streak peaks in and out of the clouds over the day. Individual clouds are averaged to a constant color.

Chapter 4

Discussion

4.1 Comparison

An experiment was performed to validate and compare the two capture techniques presented in this thesis. This was performed by capturing the sky using the two techniques at the same time at as close to the same position as possible. The resulting sky captures were then compared. Because individual camera, lens, and filter setups affect the recorded light in different ways, this required color and intensity calibration of each setup used in this experiment by photographing a Macbeth color chart. The two resulting images were used to solve for a color matrix that most closely matches images taken with the EOS-1Ds Canon camera, an ND filter, and a 8mm fisheye lens to images taken with the D60 Canon camera, a 50 mm lens, and a linear polarizer. This is the same technique that was used in Chapter 3 to correct for the chromatic shift caused by the ND filter.

The EOS-1Ds capture images were mapped into the D60 color space and then rotated to match the position of the building in the indirect capture. Both captures were then warped into angular maps as seen in Figure 4.1. The sun in the fisheye lighting environment was then approximated as a directional light, thresholding the

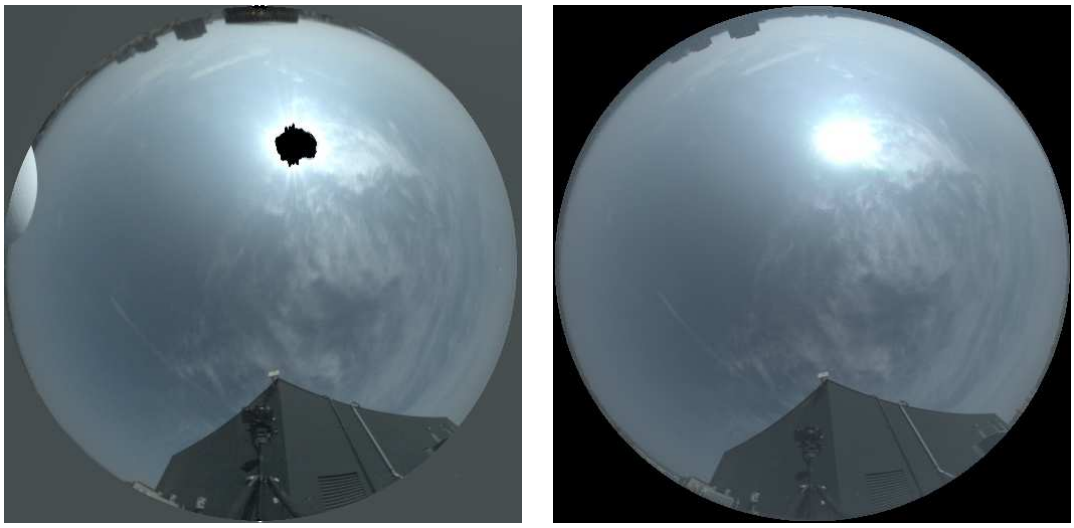

(a) (b)

Figure 4.1: These are (almost) simultaneous sky and sun captures taken with the technique outlined in (a) Chapter 2 (indirect sun capture with LDR image using spherical mirror and diffuse ball) and (b) Chapter 3 (direct HDR capture)). Because this required two cameras capturing at the same time (b) was imaged with the EOS-1Ds Canon camera and (a) the Prober Jr. was imaged with a Canon D60 camera, 50 mm lens, and linear polarizer.

pixels at approximately the same pixel values that saturated on the D60 camera sensor, as seen in Figure 4.2. The resulting sun intensities were then compared. The RGB sun intensity values solved for using the diffuse ball were 774454, 879767, 855087, and the values of the directional light approximating the sun in the HDR capture were 828340, 959394, 1008448,. This shows the difference between the two captures to be for red 3.36%, for green 4.33% and for blue 8.23%. The direct HDR sun intenisty was in general brighter than the indirect Prober capture. The difference in the blue channel is the highest; this could be caused by polarization effects on the fisheye lens, or non-ideal color calibration of the camera/lens/filters setups used in the captures.

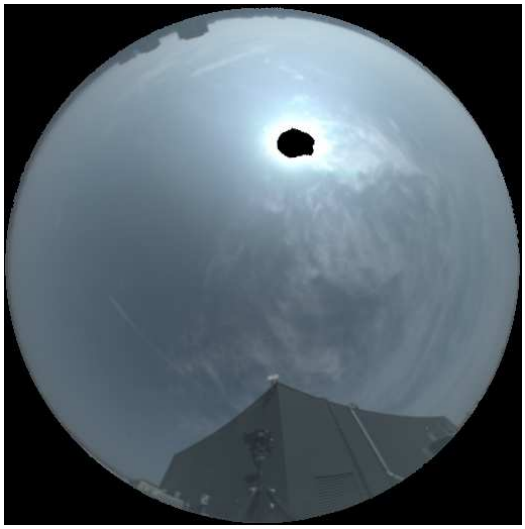

Figure 4.2: This shows the sun area approximated with a directional light for the direct HDR capture. The intensity of this directional light was used as a comparison to the intensity of the sun recovered in the indirect capture technique.

4.2 Error Source Analysis

For the indirect technique, possible sources of error include the recovered camera position and the approximated position of the sun. There are global illumination effects, such as indirect rays reflecting off the mirror, on the diffuse ball and texture of the diffuse paint which could cause some small error. There is also information lost by not directly observing the circumsolar region.

Possible sources of error for both the indirect and direct capture techniques includes abberations from a linear response curve for the camera sensor.

The accuracy of the direct HDR capture technique is affected if the reciprocity of the camera's sensor is non-perfect. It is also affected by effects caused by the ND filter, such as blurring. The fisheye lens is sensitive to skylight polarization, which have not been corrected. The correction for the chromatic shift caused by the ND filter, is only an approximate correction and we can assume the ND filter introduces some non-linear chromatic change due to spectrum of light passing through the

filter. Finally, for the direct HDR capture taking a sequence as long as 7 images long for each environment is relatively slow (40 seconds) and the position of clouds in the sky can change.

4.3 Performance

Each of the methods for capturing the sun intensity and upper hemisphere has its advantages and disadvantages. Capturing the sun with the fisheye, ND filter, and a high dynamic range sequence of images is convenient and simple. It can provide lighting environments that are high enough resolution to serve as background plates as seen in Figure 4.3. Additionally, the fisheye lens technique sees the full shape and character of the sun and the bright region around it, without making assumptions about the sun's position and size. The direct technique does not record the reflection of the camera and tripod as seen in the aluminum mirror technique.

The aluminum lens technique makes assumptions about the sun and sky and has limited resolution. Also getting the entire reflection of the sky as seen in the aluminum mirror in focus is difficult with some lenses due to depth of field issues.

However the aluminum mirror technique has advantages as well; this technique can acquire measurements as fast as the camera can take pictures, does not use an ND filter, and it is possible with two additional pictures to correct for polarization issues. The aluminum lens technique requires some materials and building but can be photographed with a cheaper camera and lens. Also taking environments with the aluminum lens makes no assumptions about the reciprocity of the camera's sensor since we are not changing exposure or aperture.

Figure 4.3: This is a rendering of the Parthenon with lighting from 7:04am (top left), 11:16am (top right), 4:11pm (bottom left), and 5:37pm (bottom right). Capturing outdoor lighting environments in over 17 stops of high dynamic range, high resolution, and time lapse allows for realistic lighting at different times of day and effects such as post-processed sun flare by Chris Tchou (top left).

4.4 Conclusion

For rendering scenes that include a visible sun the background, or timelapse HDR environments, the fisheye and ND technique is a logical choice to achieve suitable resolution and detail, especially with the easy capture process and reduced post-processing.

On the other hand, for applications such as rendering synthetic objects into real scenes and inverse global illumination, the aluminum mirror technique is a good option as it is more accurate at reproducing the appearance of non-mirror surfaces. Since we are specifically solving for the missing energy of the sun based on the appearance of a non-mirror surface, the diffuse ball, without the complexity introduced by the ND filter. The indirect method is faster it might be suitable when the

clouds and sky are rapidly changing and it is important to capture these changes.

4.5 Future Work

There are many avenues to explore in future research. It may be possible to reduce or remove lens flare caused by the sun to improve the appearance of the captured environments. Also further calibration of the 8mm Sigma fisheye lens could result in reduced chromatic abberation.

Methods of recovering spectral material properties in natural illumination, or reconstructing specular surfaces [22] in natural illumination could be possible. An apparatus for fast and passive automatic polarization capture for the fisheye and aluminum mirror could be designed ([32]). Also, it could be interesting to modify a renderer to take a ray's polarization into account, such as in [31], [12], and [30] but render with polarization from captured skies.

For the indirect method, it would be ideal to take into consideration global illumination effects on the diffuse ball, such as indirect lighting from the aluminum mirror to reduce error. Also further calibration of the cameras and lens in standard color spaces, such as ICC-Profiles [2], would be useful.

Since speed is a concern with the fisheye lens technique, methods to increase the speed of this method such as multiple cameras or HDR image sensors could be explored. Finally, visual flow fields could be implemented to interpolate between frames to synthesize sky sequences with increased temporal resolution.

4.6 Acknowledgements

I gratefully acknowledge Dr. Paul Debevec at University of Southern California who advised and provided technical expertise for every part of this thesis. I would also

like to salute Dr. Jim Arvo and Dr. Pietro Perona for their distinguished support and advice.

This work was made possible by the encouragement and contributions of Chris Tchou. I would also like to recognize my officemate, Andrew Jones, for his additions to this research. Many thanks to my co-workers at USC's institute for Creative Technologies, Andreas Wenger, Tim Hawkins, Andrew Gardner, Brian Emerson, Marc Brownlow for their constant creativity. I acknowledge Marcos Fajardo and his Arnold rendering system. This work was also made possible by the production efforts of Lora Chen and Diane Piepol, and the support of David Wertheimer, Richard Lindheim, and Neil Sullivan. Also thanks to Shane Chen for his expert spray painting skills, Richard DiNinni for disabling the roof alarm when necessary, and Visnu Pitiyanuvath for his instant messaging and late night advice. I gratefully acknowledge Charles Stumpfel for his optics consulting and Kathleen Stumpfel for her proofreading and general moral support. This work has been sponsored by the University of Southern California, U.S. Army contract number DAAD19-99-D-0046. and a graduate fellowship from the National Science Foundation.

Bibliography

[1] Canon digital camera software developers kit, 2004. http://www.powershot.com/powershot2/customer/developer.html.

[2] International color consortium, 2004. http://www.color.org/.

[3] Sameer Agarwal, Ravi Ramamoorthi, Serge Belongie, and Henrik Wann Jensen. Structured importance sampling of environment maps. *ACM Transactions on Graphics*, 22(3):605–612, July 2003.

[4] Jean-Yves Bouguet. Camera calibration toolbox for matlab, 2004. http://www.vision.caltech.edu/bouguetj/calib_doc/index.html.

[5] Katja Daubert, Hartmut Schirmacher, François X. Sillion, and George Drettakis. Hierarchical lighting simulation for outdoor scenes. In *Eurographics Rendering Workshop 1997*, pages 229–238, June 1997.

[6] Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, pages 189–198, July 1998.

[7] Paul Debevec. Light probe image gallery, 1999. http://www.debevec.org/Probes/.

[8] Paul Debevec, Chris Tchou, Andrew Gardner, Tim Hawkins, Andreas Wenger, Jessi Stumpfel, Andrew Jones, Charis Poullis, Nathaniel Yun, Per Einarsson,

Therese Lundgren, Philippe Martinez, and Marcos Fajardo. Estimating surface reflectance properties of a complex scene under captured natural illumination. *Conditionally Accepted to ACM Transactions on Graphics*, 2004.

[9] Paul E. Debevec, Camillo J. Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. *Proceedings of SIGGRAPH 96*, pages 11–20, August 1996. ISBN 0-201-94800-1. Held in New Orleans, Louisiana.

[10] Yoshinori Dobashi, Kazufumi Kaneda, Hideo Yamashita, Tsuyoshi Okita, and Tomoyuki Nishita. A simple, efficient method for realistic animation of clouds. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 19–28, July 2000.

[11] Grant R. Fowles. *Introduction To Modern Optics*. Dover Publications Inc, New York, second edition edition, 1975.

[12] Edward R. Freniere, G. Groot Gregory, and Richard A. Hassler. Polarization models for monte carlo ray tracing. In *Optical Design and Analysis Software, Proceedings of SPIE*, volume 3780.

[13] Geoffrey Y. Gardner. Visual simulation of clouds. In *Computer Graphics (Proceedings of SIGGRAPH 85)*, volume 19, pages 297–303, July 1985.

[14] Henrik Wann Jensen, Frédo Durand, Michael M. Stark, Simon Premoze, Julie Dorsey, and Peter Shirley. A physically-based night sky model. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 399–408, August 2001.

[15] Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High dynamic range video. *ACM Transactions on Graphics*, 22(3):319–325, July 2003.

[16] Craig Kolb, Don Mitchell, and Pat Hanrahan. A realistic camera model for

computer graphics. In *Proceedings of SIGGRAPH '95 ACM SIGGRAPH*, pages 317–324.

[17] Raymond L. Lee. Digital imaging of clear-sky polarization. *Applied Optics*, 37(9), March 1998.

[18] CN Long and JJ Deluishi. Development of an automated hemispheric sky imager for cloud fraction retrievals. In *In 10th Symposium on Meteorological Observations and Instrumentation*, pages 171–174, Boston, MA, 1998. American Meteorological Society.

[19] R. Miyazaki, S. Yoshida, Y. Dobashi, and T. Nishita. A method for modeling clouds based on atmospheric fluid dynamics. In *9th Pacific Conference on Computer Graphics and Applications*, pages 363–372, October 2001.

[20] Tomoyuki Nishita and Eihachiro Nakamae. Continuous tone representation of three-dimensional objects illuminated by sky light. In *Computer Graphics (Proceedings of SIGGRAPH 86)*, volume 20, pages 125–132, August 1986.

[21] A. J. Preetham, Peter S. Shirley, and Brian E. Smits. A practical analytic model for daylight. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, pages 91–100, August 1999.

[22] Stefan Rahmann and Nikos Canterakis. Reconstruction of specular surfaces using polarization imaging. In *2001 Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 1, pages 149–155, December 2001.

[23] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 117–128, August 2001.

[24] Ravi Ramamoorthi and Pat Hanrahan. Frequency space environment map rendering. *ACM Transactions on Graphics*, 21(3):517–526, July 2002.

[25] L. G. Shulz. The optical constants of silver, gold, copper, and aluminum. *J. Opt. Soc. Am.*, 44(5):357–368, 1960.

[26] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 21(3):527–536, July 2002.

[27] Jessi Stumpfel, Andrew Jones, Andreas Wenger, Chris Tchou, Tim Hawkins, and Paul Debevec. Direct hdr capture of the sun and sky. In *submitted to Afrigraph*, 2004.

[28] Katsumi Tadamura, Eihachiro Nakamae, Kazufumi Kaneda, Masshi Baba, Hideo Yamashita, and Tomoyuki Nishita. Modeling of skylight and rendering of outdoor scenes. 12(3):189–200, 1993.

[29] Gregory J. Ward. The RADIANCE lighting simulation and rendering system. In *SIGGRAPH 94*, pages 459–472, July 1994.

[30] Alexander Wilkie, Robert F. Tobler, and Werner Purgathofer. Combined rendering of polarization and fluorescence effects. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, pages 197–204, June 2001.

[31] Lawrence B. Wolff and David J. Kurlander. Ray tracing with polarization parameters. *IEEE Comput. Graph. Appl.*, 10(6):44–55, 1990.

[32] L.B. Wolff. Advances in polarization vision. In *ARPA94*, pages I:605–615, 1994.

[33] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *In SIGGRAPH '98 Conference Proceedings, Annual Conference Series, ACM Siggraph*, pages 207–218.

[34] Zhengyou Zhang. Flexible camera calibration by viewing a plane from unknown orientations. *IEEE*, 1999.