A WEB-BASED APPROACH TO IMAGE-BASED LIGHTING USING HIGH DYNAMIC RANGE IMAGES AND QUICKTIME ™ OBJECT VIRTUAL REALITY

A Thesis

by

TAMARA CUELLAR

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

May 2008

Major Subject: Visualization Sciences

A WEB-BASED APPROACH TO IMAGE-BASED LIGHTING USING HIGH DYNAMIC RANGE IMAGES AND QUICKTIME ™ OBJECT VIRTUAL REALITY

A Thesis

by

TAMARA CUELLAR

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Chair of Committee,	Karen Hillier
Committee Members,	Frederic Parke
	John Keyser
Head of Department,	Tim McLaughlin

May 2008

Major Subject: Visualization Sciences

ABSTRACT

A Web-Based Approach to Image-Based Lighting Using High Dynamic Range
Images and QuickTime ™ Object Virtual Reality.
(May 2008)

Tamara Cuellar, B.E.D., Texas A&M University
Chair of Advisory Committee: Professor Karen Hillier

This thesis presents a web-based approach to lighting three-dimensional geometry in a virtual scene. The use of High Dynamic Range (HDR) images for the lighting model makes it possible to convey a greater sense of photorealism than can be provided with a conventional computer generated three-point lighting setup. The use of QuickTime ™ Object Virtual Reality to display the three-dimensional geometry offers a sophisticated user experience and a convenient method for viewing virtual objects over the web. With this work, I generate original High Dynamic Range images for the purpose of image-based lighting and use the QuickTime ™ Object Virtual Reality framework to creatively alter the paradigm of object VR for use in object lighting. The result is two scenarios: one that allows for the virtual manipulation of an object within a lit scene, and another with the virtual manipulation of light around a static object. Future work might include the animation of High Dynamic Range image-based lighting, with emphasis on such features as depth of field and glare generation.

ACKNOWLEDGEMENTS

I would like to thank several folks who have helped me and *provided hope* for me on this arduous journey!

My committee members John Keyser and Fred Parke lent their time and expertise to my project. I appreciate their picking up my project so quickly after my long absence. By their asking some tough questions about my process, I believe my work was really strengthened. My committee chair, Karen Hillier, was immensely supportive and the best kind of mentor for me through my time in the Viz Lab. In addition to being as excellent educator, critic and cheerleader, she understood the way I worked and was extremely accommodating when my schedule (or life plan) got thrown out of whack. I truly value her artistic vision for our degree program, her commitment to professionalism in all aspects of her life, and most of all – her wit and sense of humor. I look forward to "lunching" with her long after I have a thesis to discuss.

I'd like to also thank the many friends I've made during my time in the Viz Lab. The "vizzer" culture is unlike anything I have ever known, and I really loved being a part of it for as long as I did. From the support staff to the faculty, and every student that passes through its doors, the people that I've interfaced with over the past five years are wholeheartedly committed to innovation, pushing the limits of technology, and keeping the "Viz Lab" name not only relevant, but still impressive in the industry. I am proud to say I worked and played alongside all of them.

There are several women in my life that I would like to thank for their inspiration. I am so fortunate to have been affected by a network of such strong and intelligent women! These women include my mother Mary Cuellar, my sisters Monica Cuellar and Cristina Pena, my aunt Linda Cuellar, my best friends Suzannah Creech, Amanda Williams, and Jill Meyer, as well as my friends/colleagues throughout my time at A&M University – Elizabeth Heilman, Jackie Besinaiz-Thomas, Brooke Beane, and Holly Murphy-Brackin.

Lastly, I would like to give a special thanks to Alex Garza for keeping me motivated and happy through the hardest parts of the process. The journey was long and there were many bumps on the road. His quiet confidence in me kept me going and I can only hope I can do the same for him when his time comes.

NOMENCLATURE

3d	Three dimensional
CG	Computer Generated
LDR	Low Dynamic Range
HDR	High Dynamic Range
PTM	Polynomial Texture Maps
QTVR	QuickTime ™ Virtual Reality

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

FIGURE Page

FIGURE Page

CHAPTER I

INTRODUCTION

This thesis presents a web-based approach to interactively lighting three-dimensional (3D) geometry in a virtual scene. With this work, I will generate original High Dynamic Range images for the purpose of image-based lighting, and use the QuickTime ™ Object Virtual Reality framework to creatively alter the paradigm of object VR for use in object lighting. The result will be two scenarios: one that allows for the virtual manipulation of an object within a lit scene, and another with the virtual manipulation of light around a static object (Figure1).

Figure 1. Three frames from the "Bunny Scenario 2" QuickTime ™ movie sequence.

1. Motivation for Research Direction

Motivation for this research stemmed from my desire to further my knowledge of image-based lighting for computer graphics. Recognizing the breadth of the subject, I later narrowed my interest to include its effect on surface details of 3D objects.

This thesis follows the style and format of *IEEE Transactions on Visualization and Computer Graphics.*

This interest, accompanied by previous web development experience, influenced the my decision to ask a series of questions that in turn provided the foundation for this research. Web interactivity played a part in developing the premise, by providing a tangible and accessible product for the viewer. With these factors in place, a marriage of the two interests was born.

One of the related areas of research is a method called "Polynomial Texture Mapping" (PTM). Developed by Tom Malzbender at the Research Labs at Hewlett Packard, PTMs are photographic representations of stationary 3D objects that allow enhanced image quality and improved perception of the object's surface details [1]. Like bump mapping, this allows the *perception* of surface deformations. While considered an extension of texture maps, the PTM method is, however, image-based – that is, several photographs of a surface under varying lighting conditions are actually used to construct these maps (Figure 2). Unlike texture and bump maps, Polynomial Texture Maps (PTMs) also capture variations due to surface self-shadowing and interreflections, which enhance realism [1]. The Java applet written by Clifford Lyon, allows users to view the PTM format over the web. It provides a simple and user-friendly interface that allows the user to explore the details of the intended surface features [2].

Figure 2. PTM photographs showing influence of light and shadow on a bin of almonds (courtesy of HP labs).

The researcher soon recognized an opportunity for exploration, specifically in the case of the PTMs, where a physical light moves around a static physical object. The thought to extend that logic into the virtual world seems only natural. The questions are then raised, “What if a *virtual* light moves around a static *virtual* object” and the subsequent variation “What if a virtual *light* is to stay static and a virtual *object* moved?”

The web would prove to be a powerful vehicle in either scenario for enabling the audience to experience an enhanced level of involvement with a virtual 3D object. Surface details and subtle variations in specularity could be observed, with each incremental translation of light over an object. The same would also be true with each incremental rotation of the object within the statically lit environment. The result would be for the user to have the ability to manipulate the object in a virtual lighting setup (or vice versa) offering the viewer an opportunity to observe the virtual object in a unique way not previously possible.

‘Remember the phrase “Lights, Camera, Action”? Lights come first for good reason. Without them the camera and action are useless.’ [3].

Visual observation, in the real and computer generated world is often enhanced with the assistance of lights. In the following figure, the combination of lights and shadow are used to not only define the shape and form of objects, but also serve the basic visual function of defining a spatial relationship between objects. The source of the light is indicated by the highlight and the effect on surface details, while the shadows reveal the objects’ relationship to one another. Without any indication of the source of light, it is unlikely that one will be able to identify where exactly the objects are located in this space. With the addition of shadows however, the image on the left reveals more information about the objects’ relative position and depth in space than the image on the right. This observation might be a moot point in the real world, but in a computer-generated world, the opportunity for variance from physical reality is great. In other words, there are many ways in which CG lights can go wrong and interrupt visual consistency. We do not recognize the importance of this observation until we witness it performed *incorrectly,* or find it missing altogether or (Figure 3).

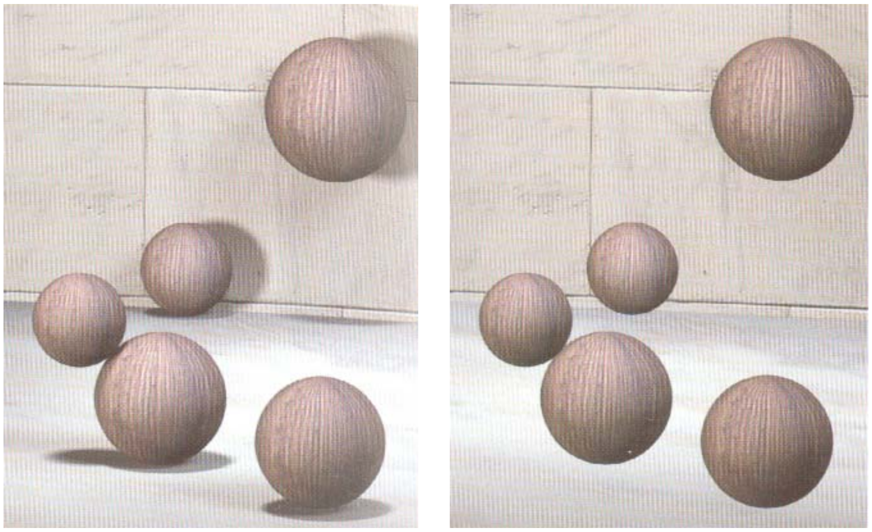

Figure 3. Images of spheres rendered with (right) and without (left) lights casting shadows [4].

The inverse square law is one rule that is obeyed in the real world and CG lights. This law explains how light's luminosity or the light's energy emission fades over distance [5]. Another example of CG lights simulating real world lights would be the way lights affect reflection and refraction. This simulation takes place using the rendering process known as raytracing.

There are also several fundamental differences between real world lighting and CG lighting. The approach to lighting a given scene differs a great deal in that both kinds of lighting artists (real world and CG) must take into account the limitations of their medium. The world of CG offers great flexibility in that there are many variable parameters that offer very tight control over a light's effect. 'Negative lights' for example, can subtly darken areas like corners of rooms, as well as be used to 'fake' shadows, which would otherwise be costly in computation time. CG artists also have the ability to exclude and include objects from a lights' influence, and the lighting design may reflect this. On the flipside of this, real world lighting artists rely on the exposure controls of the film camera to compensate for varying lighting conditions, while the CG artists adjust the lights intensity to simulate similar conditions and effects.

Because of its visual impact, the manipulation of lights in the real world has long since been an integral part of the cinematic experience. As advances in computer-generated imagery are made, mediums for a creative use of applied lighting (real world and CG) go beyond the cinematic, including educational, technological, artistic and scientific.

The importance of CG lighting can first be understood by recognizing that much of what is done in CG is done to simulate and emulate the look of the real world. The true advantage of CG lighting though, is that it can afford its artists the ability to not only simulate the real world, but do so with flexibility, control and creativity – ultimately *enhancing* the visual experience by convincingly obeying and deceivingly breaking what we know as physically-based lighting rules.

2. High Dynamic Range Images

High Dynamic Range images (HDR images) differ from traditional digital images in that they contain more illumination information than can be reproduced with a standard display device (such as an LCD or CRT monitor) or captured using a conventional 35 mm film or digital camera with a single exposure. From the brightest light source to the darkest shadow, the HDR image pixel values are directly proportional to the light energy in the scene [5]. This differs from most digital images in that, for example, the RGB color spaces’ pixel display values are not proportional to the light levels in the scene. Usually, light levels are encoded nonlinearly, or are “clipped” so they can effectively appear on nonlinear display devices such as cathode ray tubes or liquid crystal displays. Furthermore, standard digital images typically represent only a small fraction of the dynamic range—the ratio between the dimmest and brightest regions accurately represented—present in most real world lighting environments. When part of a scene is too bright, the pixels saturate to their maximum value (usually 255) no matter how bright they really should be. When used as a source for image-based lighting, the high dynamic range images yield higher contrast, more detail, with more realistic lighting effects than conventional digital images.

Figure 4a. HDR image of Old St. Paul’s Cathedral in New Zealand, by Dean Pemberton.

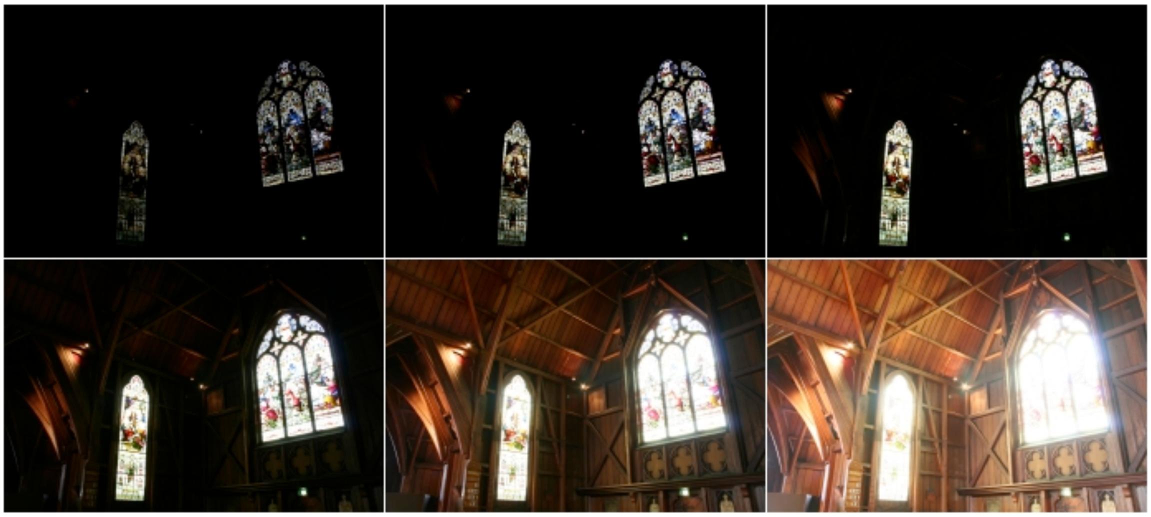

Figure 4b. The six individual exposures used to create the previous HDRI. Keeping f-stop constant, the exposure times from top left are: 1/40sec, 1/10sec, 1/2sec, 1sec, 6sec, 25sec. http://commons.wikimedia.org/wiki/Image:Old_saint_pauls_1.jpg. [6]

In the scope of this research, the creation of HDR images begins with capturing the broad range of light in a single scene. Due to the limitations inherent in most digital camera light sensors and film emulsions, it is not possible to capture this broad range of light in a single exposure. There are ways to get around this problem. By bringing together multiple exposures from a standard film camera, with the right software tools, it is in fact possible to create a single high dynamic range image.

By taking a range of multiple exposures as in Figure 4b, each image in the bracketed sequence will have different areas properly exposed, while other areas are under or overexposed. For example in the first of the series, one can observe great detail in the stained glass, while the rest of the image appears too dark to distinguish. On the opposite end of the spectrum, the last image possesses considerable detail in the woodwork, while the window is ‘blown out.’ One can see that none of photographs in this series individually exhibit the “happy medium” exposure settings that can be achieved with a high dynamic range image. But by acquiring the series of photographs, each with the same f-stop but with their varied exposure times, they can come together forming the final composited high dynamic range image (Figure 4a).

There is another method for creating an HDR file from photographs taken with multiple exposures. In this thesis project, we will limit the discussion only to manipulation within Photoshop, primarily because it is accessible and ease of use.

Part of the way Photoshop creates an HDR file is by reading the EXIF (exchangeable image format) information that is encrypted in each file. This information reveals specific details about how the image was captured, including shutter speed, aperture and ISO settings. Photoshop takes this information, in combination with the 16 or 24 bit depth color information to composite the images together and output luminance values in a 32-bit high dynamic range image format. A unique characteristic of the luminance values in an HDR image is that the color information is stored in floating point values that directly relate to the amount of light in a scene. This relationship of luminance values to the light in a scene then creates a relatively open-ended brightness scale [7]. This method is different from those used in 16 and 8-bit image files, in that low dynamic range images store their information (gamma encoded) within in a fixed range of luminance values – essentially “clipping” the lightest lights or the darkest shadows. To combine multiple low dynamic range images together with the flexibility of a 32 bit format then provides an interesting opportunity to see the widest range of values within one photograph – a high dynamic range image.

3. Image-Based Lighting

The method of placing individual lights around objects in a 3D scene for the purpose of achieving desired lighting conditions is one that has been well documented in the computer graphics world for years. Effective use of computer generated (CG) lights in a 3D scene however, requires the careful and specific placement of a variety of lights (Figure 5). In addition to noting the great *number* of lights that might be used in any scene, it is important to also note the many *types* of CG lights: point lights, spotlights, area lights, directional lights and ambient lights, just to name a few. Each of these types of lights has a set of common parameters, as well as other parameters that are specific to the individual type. The cumulative effect of these lights on a scene, in conjunction with the fine-tuning of their many parameters, can produce very high quality images.

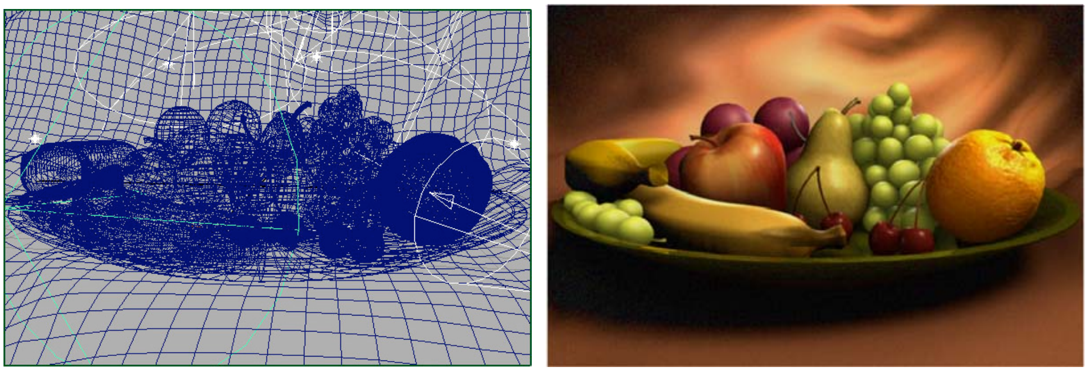

Figure 5. Positions of CG lights in a scene (left). Scene rendered with many CG lights (right).

There are other methods for achieving similar if not more photorealistic results. Image-based lighting (IBL) is the process of illuminating scenes and objects with *images of light* from the real world rather than using CG lights [7]. Image-based lighting works just like real world illumination in that the environment provides not just reflections on the object, but also captures the ambient light information from all natural or artificial light sources present. It has evolved from a reflection-mapping technique in which panoramic images are used as reflection maps on CG models to show how shiny objects might reflect real environments. When used effectively, IBL can produce realistic rendered appearances of objects and can be an effective tool for integrating CG objects into real scene.

Because HDR images record the full range of light in an environment, these images contain valuable information about the shape, color, and intensity of direct light sources, as well as the color and distribution of indirect light from surfaces in the rest of the scene. In addition to capturing multiple exposures to collect the most light information, HDR images also need to be *omnidirectional.* This process of capturing images that 'see' in all directions is used to mimic the light information in any traditional lighting setup, and to create the most effective environment [8].

The basic steps in Image-Based Lighting are:

1. Capturing real-world illumination as an omni-directional high dynamic range image.

This can be done in one of several ways including: capturing a HDR image with an expensive, specialty panoramic camera, "stitching" many images together with a photo-editing software like Photoshop, or creating a HDR image using the light probe method.

2. Mapping the HDR image onto a geometric representation of the environment.

While this 'environment' is not necessarily visible to the viewer in the final render, it is needed in set-up for achieving the desired lighting effect. This 'environment' is often a simple primitive such as a sphere or cube. A cubic environment map is created by displaying six coordinated images on the inside faces of a cube. The spherical environment map requires nonlinear image warping to compensate for distortion (Figure 6).

Figure 6. Cubic (left) and spherical (right) environment map projections

3. Placing the 3D model inside the virtual environment, aligned with the origin.

4. Simulating the light from the environment illuminating the 3D model during rendering.

This can be done in the render globals/settings dialog box of a standard 3D application.

5. Render, and post pocess the renderings as needed.

An important application of IBL is in the area of motion picture visual effects. As the ultimate aesthetic goal is for the computer-generated object to appear as though it was filmed within the live action scene, the challenge can often be met with image-based lighting techniques, utilizing HDR technology [7]. These IBL techniques allow for real illumination information to be captured at the location the CG object needs to be placed, and then used to light the CG element so that its shading, shadows and highlights are all consistent. Using IBL as a starting point, CG lighting artists can modify and adjust image-based lighting to achieve effects that are as dramatic as they are realistic.

4. QuickTime ™ Object Virtual Reality

QuickTime ™ Object Virtual Reality (QTOVR) is an extension of the QuickTime ™ digital multimedia framework that uses a series of sequential digital images to provide additional information about objects that could not be made apparent through the use of traditional 2D photography [9]. With Object VR, photographs are taken of a stationary object from incremental positions around it, and then combined in a way that mimics an animated sequence of images, to produce an interactive representation of the object. The result is an enhanced web experience for the viewer, where s/he can further investigate the object in a way not possible from a single static image.

A camera has three rotational degrees of freedom: pitch (pivoting about a horizontal axis), yaw (pivoting about a vertical axis) and roll (rotating about an axis normal to the view plane). The simplest QuickTime ™ object model possesses a sequence of images taken incrementally around an object in a full 360 degrees around one axis. When these frames are viewed along side each other, or put into a loop, the object appears to be rotating about its vertical axis (Figure 7).

Figure 7. Sequence of images rendered around a vertical axis. Model by Chris Wheeler.

Interactivity is added when the user is given opportunity to click and drag the cursor frame by frame, simulating a smooth transition. When frames are added incrementally around the object as in Figure 8, with the camera rotating about the object's *horizontal axis* (pitch), a true QTOVR is represented. This is called a global rotatable model as opposed to the more primitive single axis rotating display.

Figure 8. Sequence of images rendered from cameras positioned around both horizontal (pitch) and vertical axes (yaw). Model by Chris Wheeler.

a. Capabilities / Benefits of Using QT Object VR

The use of QuickTime ™ Object VR for the purpose of this research offers several benefits. First, the inherent interactivity of this popular viewer application enables the viewer to experience a more involved investigation of the object. By allowing the viewer to rotate and turn the object the object manipulation will appear more realistic and seamless. The interactive controls are simple, and intuitive to the average web user. Previous web experience and familiarity with the viewer application enables viewers to just click and drag to turn or rotate an object; to zoom in by holding down the option key (the shift key in Windows), and zoom out by holding down the control key. The manageable file sizes that result from available compression techniques make for easy access on the web. Files load quickly and are significantly smaller than other digitized video alternatives.

CHAPTER II

HISTORY

Methods for computer generated lighting simulation have continued to improve in recent years. This section discusses several of Paul Debevec's SIGGRAPH papers pertaining to image-based lighting research, as well as the progress feature film and gaming industries have made in the use of HDR technology. Also discussed are two methods that have been developed specifically for the web that allow the users to manipulate a lighted virtual scene via the web.

1. Image-Based Lighting in Motion Pictures

The past 15 years has seen great advances in image-based lighting. Some of the earliest practical applications for image-based lighting have been in feature film. Visual effects professionals first introduced audiences to image-based, reflection-mapping techniques in the Walt Disney's *Flight of the Navigator* in 1986 (Figure 9). In 1991, *Terminator 2* was released, and Industrial Light and Magic showcased a stunning metallic CG villain with detailed shiny reflections never before seen by audiences (Figure 10).

a. Early Reflection Mapping in Film

In both of these examples, the reflection mapping technique not only produced realistic reflections on the computer-generated object, but in the case of the *Terminator*, made the character appear to have actually been in the filmed environment. This was an important advance for realism in visual effects. Environment mapping consistently produced convincing results for shiny objects, but innovations were necessary to extend the technique to other varieties of computer-graphics models, most specifically models not possessing highly reflective surface characteristics.

Figure 9. Still from Randal Kleiser's 1986 film, *Flight of the Navigator*, demonstrating early reflection mapping in a feature film [10].

Figure 10. Industrial Light + Magic used the technique to create the look of the T1000 robot in the 1991 film *Terminator 2* [11].

b. Paul Debevec and HDRI Radiance Maps

At the 1997 SIGGRAPH conference, Paul Debevec presented his paper entitled "Recovering High Dynamic Range Radiance Maps from Photographs". Starting from where previous reflection and environment mapping techniques left off, this method took image-based lighting a step further. In a process that involves photographing the same scene several times with a wide range of exposure settings and combining these separate exposures into one image, the high dynamic range (HDR) image was born. A true HDR image captures a much greater *dynamic range* of the light information than other common image formats, making it possible to simultaneously record details of a scene from the brightest intensity values, to the darkest shadowed areas, as well as all the mid-range intensities in between. The pixel values of HDR images extend outside the common RGB color space, and take full advantage of the visible spectrum of light observed by the human eye [7]. HDR images are large files, stored in a 32 bit floating point format, and are not compressed or inherently ‘clipped’ to 0 to 255 integer values required by peripheral displays as LCD or CRT monitors.

c. Putting the Light Probe Method to Work

In his 1998 paper, Paul Debevec presented "Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography.” This paper discusses a method of capturing light information involving, like in traditional HDRI capture, photographing a subject under a variety of exposure settings (Figure 11). In the case of Debevec’s light probe, the subject of the photo is a twelve inch mirrored ball in a uniquely lit environment. The HDR capture of the mirrored ball serves as an omni-directional, panoramic image of the environment and is used to record the environmental appearance and illumination information. This approach added an unprecedented level of realism, supplying real-world lighting data to a whole lighting model. In 1999, Paul Debevec used High Dynamic Range images to take some impressive steps forward in the area of digital imagery and lighting. In his three-minute animation, entitled *Fiat Lux* (Figure 12), he introduces and utilizes the HDR *light probe*, to illuminate and measure the light affecting his virtual subjects. Paul Debevec’s choice in depicting a granite monotlith, glass pendulum, and hundreds of reflective chrome spheres each with choreographed movement through a virtual St. Peter’s Basilica, is a testament to the accuracy of his light probe [7]. The result is a hyper realistic and cinematic animation. The convincing detail with which the light and shadow were captured and adjusted in post-production further enhanced the sense of realism.

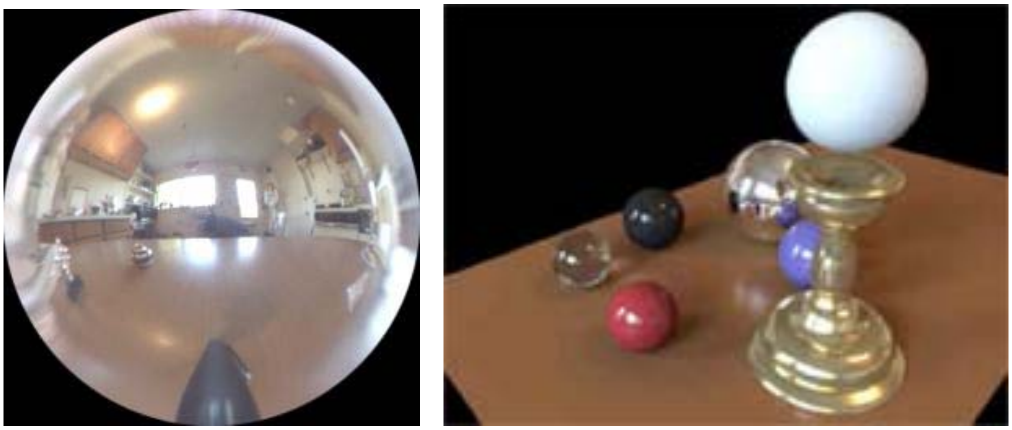

Figure 11. A collection of objects is illuminated by the light information obtained in an indoor kitchen environment. (Light Probe image by Paul Debevec) [7].

Figure 12. Panoramic rendering of St. Peter’s Basilica in Paul Debevec’s animated short film, *Fiat Lux.*

The techniques of High Dynamic Range Imaging and Image-Based lighting (HDRI and IBL), and the methods and systems derived from them, are now widely used in the visual effects industry. These methodologies have provided visual effects artists with new lighting and compositing tools that give digital actors, buildings, machinery and creatures a previously unattainable realism especially when composited into a live action environment. Examples of elements illuminated in this way include the transforming mutants in *X-Men* and *X2: X-Men United* (Figure 13), virtual cars and stunt actors in *The Matrix Reloaded*, and whole cityscapes in *Spiderman 2*.

Figure 13. Frank Vitz’s reference photos from the production of *X-men 2* (SIGGRAPH 2004 course). Images on the left are HDRI to be converted into light probes. Images on right are composites with the character, *Mystique* - and lit using the acquired light probe information [12].

2. Real-Time HDRI Developments in Video Games

These advances have lent themselves to the video gaming industry as well. With “next generation” consoles dominating the market today, consumers (or “gamers”) demand a high degree of realism from the evolving medium. Development of in-game HDRI has for a number of years seemed out of reach, being very taxing for PC hardware and console game engines. For years, the dramatic, movie-like quality that HDRI brought to video game imagery was only seen in pre-rendered cinematics, trailers and in-game movies, not real-time play.

At the 2004 E3 convention, gamers saw the first glimpse of real-time HDRI.

Valve Software announced its Unreal Engine 3, part of the “next generation” of game technology. Making the calculations required to enable real-time HDRI less taxing, this advancement breathed new life into the visual quality of games, blurring the lines between film and game entertainment (Figure 14).

Though in-game HDRI is still a rather new concept, recent games with this capability have met rave reviews from gamers, and this feature will likely become a standard as the technology matures.

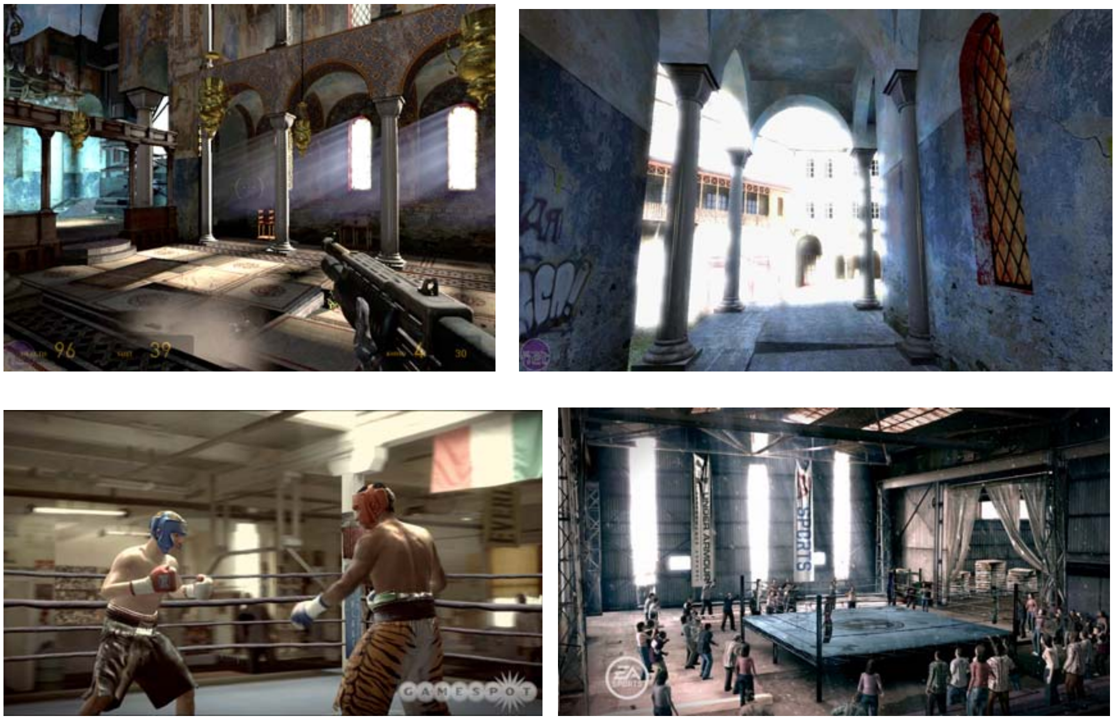

Figure 14. The first example of High Dynamic Range technology in games: the playable technology demo, *Half Life 2:Lost Coast* (top images) 2005 [13].
“Next Generation” console - Xbox 360 game, *FightNight3* screenshots (bottom images) [14].

3. Integrating HDRI Technology on the Web

Further progress has been made outside of film and games. The web is a vast medium in which developments are being made constantly. With a wide variety of coding languages made available and rich content/media becoming ever more prevalent, the opportunities are vast for HDRI exploration. In 2003, Masaki Kawase developed an online demonstration of “Real-time High Dynamic Range Image-based Lighting.” A stunning extension to Paul Debevec’s work in that he references his library of light probes, Kawase engages the user in an experience that showcases HDRI, and lets the user “play” in real-time with such settings as glare generation, automatic exposure adjustment and depth of field blur (through Microsoft’s DirectX 9.0 API) [15]. The demonstration though, while rich with visually stunning features, is just that, - a demonstration. The user has a choice of only two light probe environments and a handful of objects (Figure 15). Functionality is limited to what the author packaged as the demo product. Another weakness is that it relies on a Windows-only executable file, again limiting who can operate the demonstration.

It will be interesting to see where this can go, and it would be most helpful if more documentation were made available. Already five years old, it is difficult to determine if there will be more development. After many failed attempts at contacting Kawase, I know that both time and language can still be valid obstacles in achieving effective collaboration.

Figure 15. Screenshots of the “Real Time High Dynamic Range Image-Based Lighting” demo by Masaki Kawase.

Lastly, it is worth discussing the work of Tom Malzebender at the Hewlett Packard Research Labs. His web-friendly method for increasing photorealism in texure maps is called “Polynomial Texture Mapping.” This method relies on coefficients of a biquadratic polynomial that are stored per texel, and used to reconstruct the surface color under varying lighting conditions. Like bump mapping, this allows the perception of surface deformations. However, because this method is image-based, a large array of photographs of the surface under varying lighting conditions are needed to construct these maps (Figure 16).

This method was found particularly useful in paleontology as can be seen below. The Java applet written by Clifford Lyon, allows users to view the PTM format over the web. It provides a simple and user-friendly interface that allows the user to explore the details of the intended surface features. It is limited though in how the user spatially interacts with the object and by the fact that it can only interpret real objects.

Figure 16. PTM images of cuneiform tablet (courtesy of Hewlet Packard Labs). Original photograph of a 4000 year old cuneiform tablet. (top left), PTM reconstruction (top right), Synthetic specular highlights created from extracted surface normals (lower left), Synthetic highlights added to PTM reconstruction (lower right).

CHAPTER III

METHODOLOGY

This project consists of several parts: obtaining a light probe, preparing photographs for HDR use, rendering the 360 degree views of an object in two different scenarios (moving object with static light, and static object with moving light) and finally, sequencing the images for the QuickTime ™ Object VR manipulation. For flexibility and variety, the environments will be comprised of one outdoor scene, one interior scene, and one indoor studio lighting setup. The exterior environment is located in a creek bed in my hometown neighborhood in Helotes, Texas on a bright, sunny day. The Jack C. Williams Administration building on the Texas A&M Campus provides the indoor environment, with a unique colored glass effect and ornate interior architectural details. Finally, a three-point studio setup was created in the fourth floor studio of the Visualization Lab at Texas A&M University.

1. Image Capture - Obtaining a Light Probe

To capture the lighting of a scene in all directions, the most effective HDR images are often 360º panoramic images. These images can be captured in a variety of ways. Documented methods include stitching a series of tiled images together, capturing images taken with a fish-eye lens, and using the 360º scanning feature of expensive and specialized panoramic cameras. In this thesis, I will utilize another effective method of omnidirectional photography. It is a process of capturing images that, like the other methods, ‘see’ in all directions from a single point in space, however offers a convenience more specifically applicable to the multiple exposure method of creating HDR photography.

Paul Debevec’s “light probe” method includes photographing a mirrored ball at several different exposures and combining these low dynamic ranges images to create one high dynamic range image. The idea behind using a mirrored ball as the single reference point in space, stems from the fact that the mirrored ball can very nearly 'see' 360 degrees of its environment and that unique view can cleverly be captured with a typical camera lens. That is, anything visible from the viewpoint of the mirrored ball will be visible to the camera as a reflection in the ball. In this thesis, I use the light probe method, particularly for its transportability, its low cost and quick results.

2. Image Capture and HDR Image Generation

By taking a series of digital photographs of the mirrored ball at varying exposures, one can begin to understand the true meaning of dynamic range. As in most interior photography for example, it is apparent that a traditional film or digital camera can only expose for one of the two distinct contrast areas at any given time - the interior space or the exterior space (Figure 17). The contrast between the two is usually too great for the camera to accurately respond, which explains why, in most photography we see today, interior shots have blown out exteriors, and shots properly exposing for exterior detail have very dark interior spaces. The beauty of HDR images then, is that by bringing the full range of exposures together, the viewer can finally experience a visual representation of a scene that is more true to life, that is previously possible with traditional single exposure photographs.

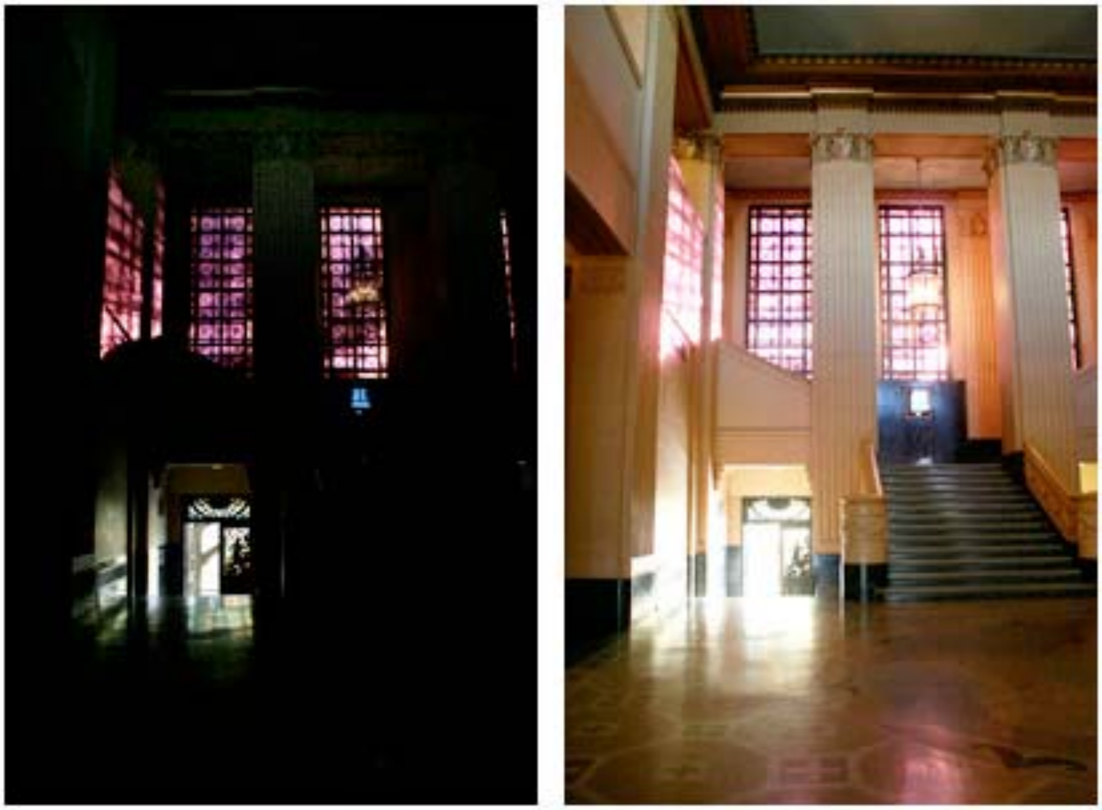

Figure 17. Photograph of Administration Building interior, exposed properly for viewing the tree outside (left). Photograph exposed properly for viewing the interior details.

To begin the process, the camera's aperture is narrowed to f/5.6 to allow sufficient depth of field within the environment. Then by systematically varying the shutter speed to capture different levels of light – from 1/5 seconds to 1/200 of a second (depending on environment), a set of images is created (Figure 18). In each exposure, the time the shutter is open is doubled. This method is repeated until the image is very obviously over exposed, with the darkest object in view finally becoming clear and detailed. In reviewing the six bracketed photographs, one can notice that the best levels of detail in the stained glass window are obtained by taking photographs with a very fast shutter speed, and alternately, the best levels of detail in the interior space are obtained by photographs with very long shutter speeds.

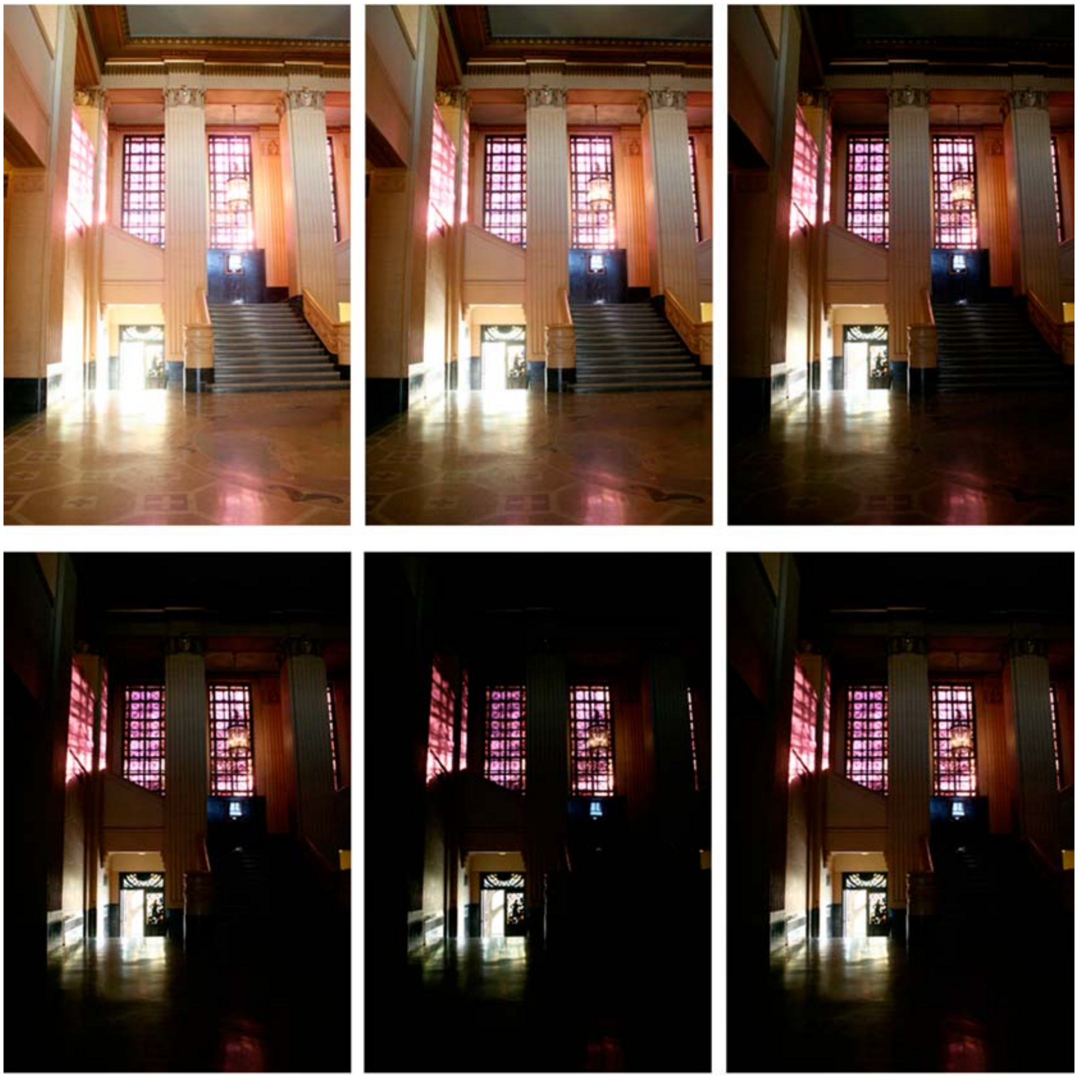

Figure 18. Keeping f-stop constant, the exposure times from top left are: 1/5 sec, 1/10 sec, 1/25 sec, 1/50 sec, 1/100 sec, 1/200 sec.

With the newest release of Adobe Photoshop, one can merge and assemble conventional 8 and 16 bit images into 32 bit High Dynamic Range Images, by using the “HDR merge” command. This automated method of layering multiple exposures, and then compressing the information into a single file will proves an efficient method for creating HDR images.

The first trial of producing these HDRI light probe images, showed a few problems that needed to be worked out. First I found that the photographs of the mirrored ball inevitably include myself, the photographer. Upon further inspection, it also becomes apparent that a *single* photograph of a mirrored ball ultimately obscures a small portion of the environment that the camera cannot ‘see.’ This obscured area is illustrated by the yellow field in Figure 19. In other words, taking a photograph from a single location will result in a far from perfect panoramic image. The law of reflectance states that “the angle of incidence equals the angle of reflection" [Weinhard]. Consequently, we have to determine just what the mirrored ball reflects, by looking at the angle that the camera's view makes - to a point on the mirrored ball. This is also known as a tangent as illustrated in the line of Figure 19.

What is found when determining this tangent angle then, is that the farther away the camera is located from the ball, (or the smaller the ball is) the more of the environment is reflected (left), or the less the environment is “missing”. The location of the photo shoot can inevitably determine how easily photographs are taken, as space may be limited.

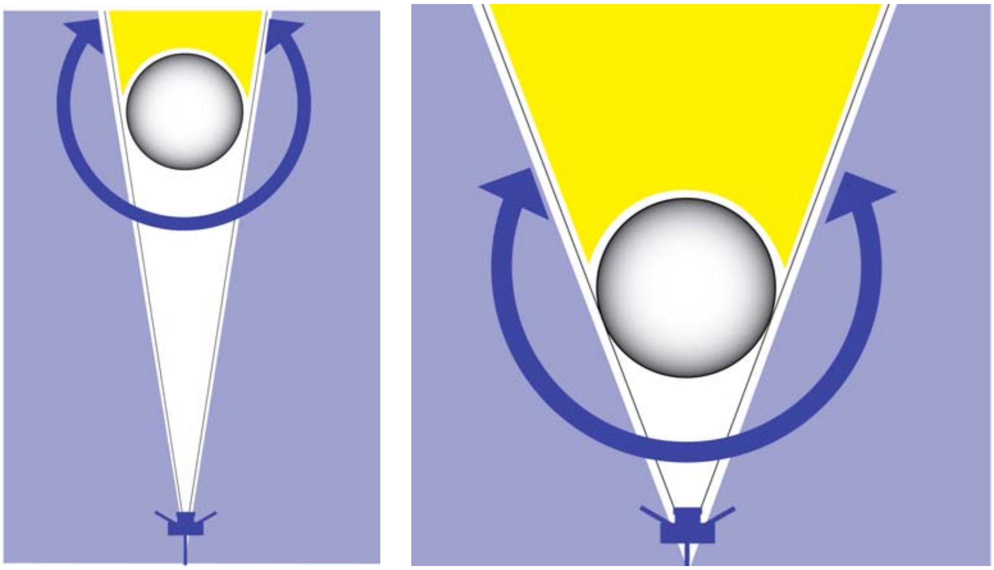

Figure 19. The camera’s tangent line, and thus angle of reflectance, is dictated by the camera’s distance from the mirrored ball and/or the size of the mirrored ball.

Secondly, things that are reflected near the edge of the ball become extremely stretched and distorted, resulting in an unsatisfactory result when it is unwarped in a software program like HDRShop. Since the mirrored sphere reflects a *nearly* 180 view of the environment in both directions, a full 360 degree view can be achieved only by filling in for the missing degrees of reflectance from the first camera position.

To achieve this compensation and alleviate these problems, the process of photographing the mirrored ball at the specific shutter speed intervals is repeated from another location. Since the two 'bad' spots in the mirrored ball are directly towards the camera (and the photographer), and directly away from the camera (the yellow area obscured by the ball), the two pictures should be taken from positions 90° apart from each other (Figure 20). This way the regions of bad sampling and camera interference will be in different locations in the two images. These affected regions will be masked out later, after converting the images to HDRI.

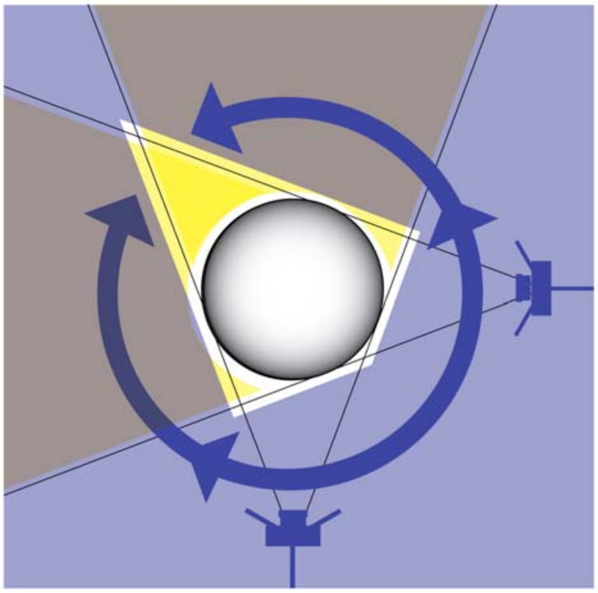

Figure 20. By combining two images, each taken at 90 degrees from each other, I can compensate for distortion and blind spots.

The product of the methods so far have yielded two HDR images, each with the photographer and several distortion errors present. By comparing the similarities of images against one another, it can be seen that the remaining issues can be remedied by compositing them in an application like Photoshop and cleverly masking the unwanted areas out. The result then would be a single, uninterrupted, high dynamic range image of a 360 degree space.

3. Rendering Multiple Object Views for Two Scenarios

a. Scenario One: Moving Object, Static Lights

The next step will be to use the HDR images to simulate the lighting environment for the 3d model. This is done in Autodesk Maya, by mapping the HDR image onto a geometric representation of the environment, such as a large encompassing sphere. The 3d model will reside at the center of the sphere, and a camera will be placed a designated distance away from it, with its focus on the model. I will then write a script to dictate the rotation of an animated camera in increments of 10 degrees in both equatorial and azimuth angles around the model.

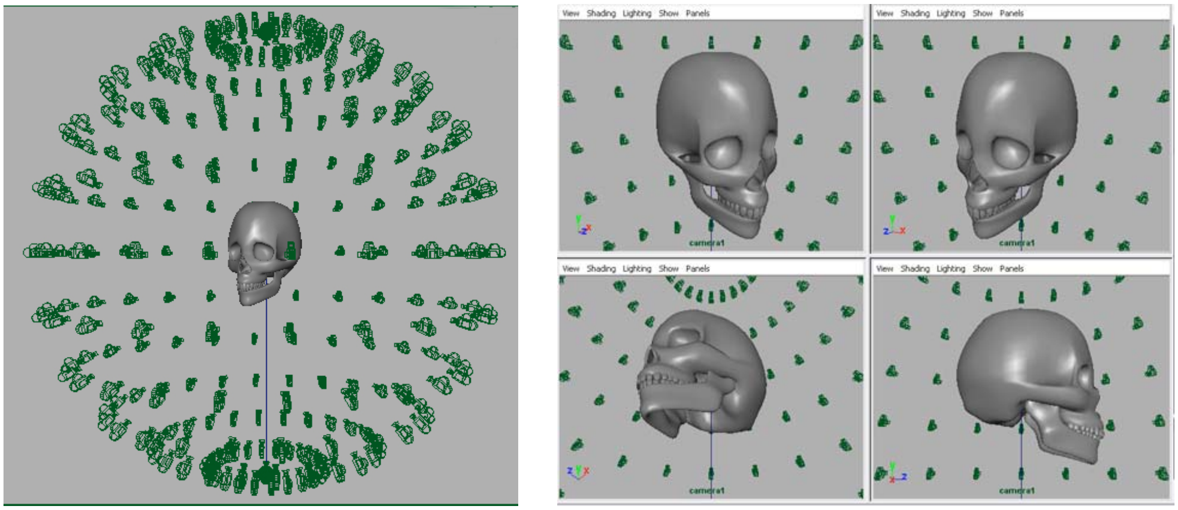

Figure 21. The orientation and location of this animated camera will, in essence, create a simulated “dome” of cameras. Dome of cameras (left). View from random cameras (right).

This dome-shaped array of cameras is needed to render images of the object from the various views (Figure 21). These rendered images will provide the information needed later in the QuickTime ™ Object VR software to simulate an interactive manipulation of the model.

b. Scenario Two: Static Object, Moving Lights

To execute the second scenario, it is important to recognize how the clever manipulation of lights and cameras can provide the viewer with *the illusion of a specific perception.* The frames rendered in the first scenario reflect exactly what is in front of a moving camera. In this second scenario however, the camera remains still and the environment moves - giving the viewer an interesting perspective – and the perception that lights are moving around the object.

Like the first method, the 3d model will reside at the center of the HDR environmental sphere. In this case however, there is only one designated position for the camera, as opposed to the previous method's virtual dome of cameras positioned around the object. The distance of the camera from the object is fixed and the model will sit on a level plane. This planar piece of geometry is necessary to see how the model's contact shadows are affected by varying light conditions. There will be as many rendered images as there are environment rotations. An expression similar to the first scenario's is used to describe how the environment will rotate (from top to bottom and 360 degrees around). The result is 684 images that show varying lighting conditions, each with subtle differences that will become apparent when the images are viewed with the QuickTime ™ viewer.

4. Sequencing Images in QuickTime ™ VR Authoring Software

Once the images have been rendered, the next step is to take these images and organize them for use in the QuickTime ™ VR authoring software. By rendering images from all of the surrounding positions, rather than just from around a single axis, a *global rotatable* model is created. This 360 degree view of the object can better exhibit the form of the object, as opposed to the more primitive *single axis rotatable model*, which simulates only a turntable view.

For smoother transitions, the Apple's standard recommendation is to have increments of 10 degrees in both equatorial and azimuth angles, which results in a hemisphere of 19 rows of 36 images each, or 684 total images [16]. Increments of 20 degrees can also be used for similar results, at a lower cost, due to less render time. On the QuickTime ™ movie timeline, the rows are laid out linearly, adjacent to one another. Horizontal mouse movements move the timeline frame-to-frame, while vertical mouse movements jump to the same relative frame in the appropriate adjacent group.

After the images have been sequenced, they are ready to be combined to form the QuickTime ™ movie. The cumulative image data collected is large, exceeding 200 MB. By moderately compressing the data, the web storage requirements can be reduced to a more manageable size, with little visible degradation of image quality.

CHAPTER IV

IMPLEMENTATION

This chapter documents the implementation of the methods developed in the previous chapter to form a pipeline that involves creating a light probe, preparing original photographs for HDR use, and rendering the 360 degree views of an object for two distinctly different scenarios. This chapter also describes my use of this pipeline to create results that can be displayed on the web in the form of a Quicktime interactive movie. A more detailed account of the challenges that I encountered throughout the working process is fully noted.

1. Getting Started : Selection of Environment Subject Matter

In selecting appropriate environments for this experiment, I looked at the careful choices made by Paul Debevec in his extensive HDRI research. In his many experiments, he often returned to a handful of light probes, each having very specific light qualities. St. Peter's Basilica, for example, had an almost angelic quality to it, as the skylight illuminates the subject from above. The Kitchen light probe offered a casually bright, but equally diffuse light quality. The eucalyptus grove light probe offered another rich light quality, with bright ambient light and deep, lacey shadows throughout. To look at the environments collectively, one might notice a few differences between the environments including color, contrast and "ceiling' height. It was important to me to capture a similar variety of environments, ultimately to be afforded the opportunity to create different and interesting results with each scenario.

For the first location, I chose an interior environment that I found to possess a light quality comparable to that found in Debevec's St. Peter's Basilica light probe. The interior lobby of the Jack Williams Administration building on the Texas A&M Campus radiates with an amber glow and in the afternoon is bathed in a saturated magenta from the stained glass windows. Though no overhead sunlight is present, an architecturally detailed ceiling and intricate staircase give the space a regal if not dramatic presence.

The second environment was a creek bed in my hometown. Similar to the eucalyptus grove, it is a scenic and wooded area, void of any kind of urban elements. On the day of shooting, the sky was bright, clear and blue through the tree branches and vines. The uneven ground was to be expected, however there were times when I felt I was testing the limits of the tripod being used.

The last scene proved challenging, despite the control of the space. I chose to emulate a professional three point studio setup, and in this, I was posed with the task of not only producing an *effective* HDR image, but I needed to create a very *clean* image as well. The creek bed and administration building environments lent themselves to intricate reflections. Unlike the other scenes though, the studio setup would need to reflect very sharp, clear images of softboxes. Because of this, it was important to digitally 'clean up' distracting areas.

At the end of the several days of shooting, I found that my experiences with each location were as different as the results that came from them.

2. Capturing Light Probes

The process of capturing a light probe began with gathering my tools and being ready to document the choices I made. I approached my locations armed with a Canon Digital Rebel camera and two sturdy tripods: one for the camera and one for the mirrored ball. Because I did not want to obscure any of the light information and wanted to keep the reflection of distracting elements to a minimum, I chose to use the ball on the tripod only when necessary, as the legs protruded a bit out from underneath the mirrored ball. At the creek bed location, the mirrored ball was affixed to the tripod with duct tape (state of the art, eh?). In the case of the studio setup and Administration Building location, I also used a cardboard mailing tube to support the mirrored ball. The flat floor surface was ideal for balancing the ball on the tube. The mailing tube worked well in both the Administration building and the studio because it hardly appeared in the images.

The choice in mirrored ball evolved with some trial and error. I ended up being very pleased with the images gathered with the ten-inch garden gazing ball as opposed to smaller (yet still very effective) Christmas tree ornaments or chrome balls. My preference stemmed from the fact that I could more easily focus on my reflection, guaranteeing sharp detail and aim, with the gazing ball at a distance of four feet than I could with a 2.5” inch chrome ball or ornament at an equivalent distance.

Each location required slightly different preparation (finding ‘just the right spot’) as well as different aperture settings, but the process remained fairly consistent. As it was necessary to take photographs of the mirrored ball from two angles for compositing later, it was important to try to eliminate any variables in method. For this reason, I also brought along a tape measure to assure that there was no discrepancy in distance from one shooting position to another.

Taking the photographs was straightforward. After centering the mirrored ball within frame, I zoomed in on the ball so that it took up almost the full height of the frame. Capturing the full range of light was accomplished by the repetitious bracketing of shutter speed. In documenting the process, it was again important to eliminate variables, so I brought a notepad for field notes. Images reflecting the replicated shutter speed settings for the two angles are shown in (Figure 22).

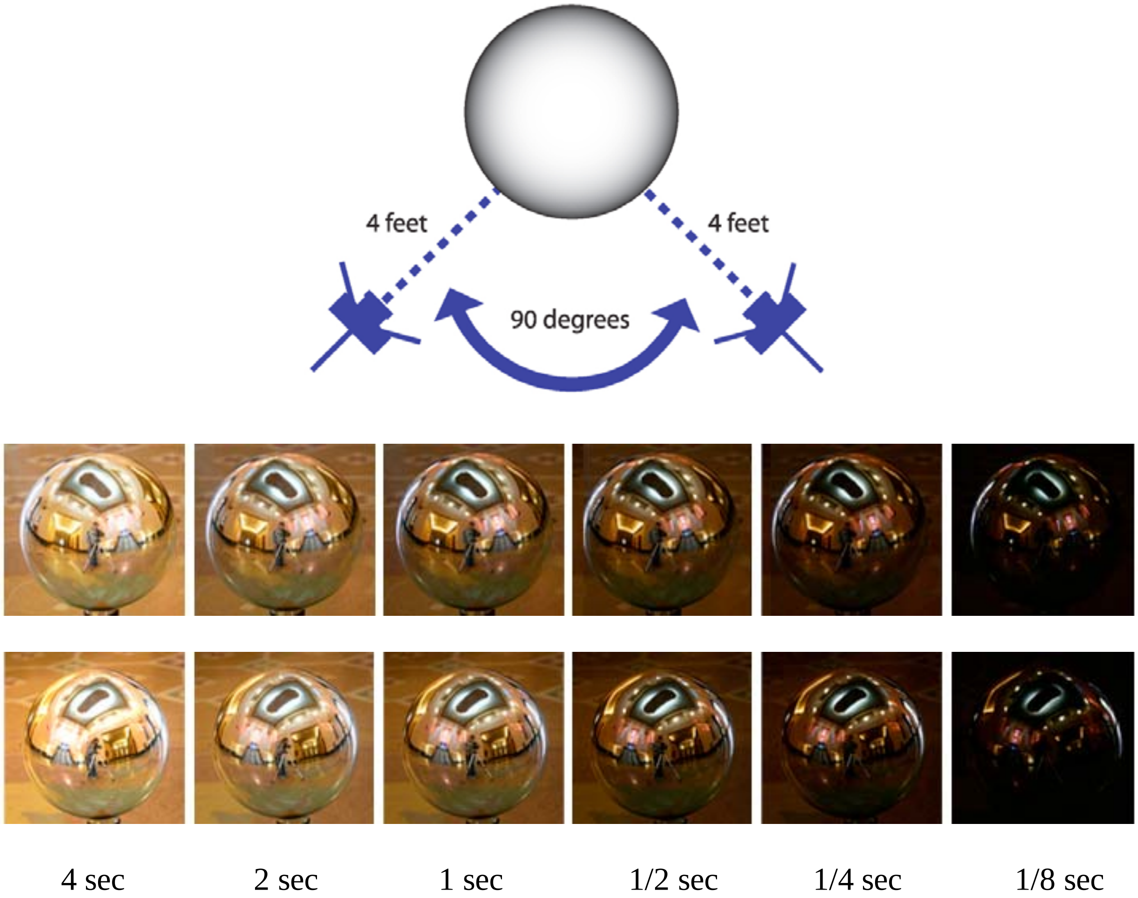

Figure 22. “Lefty and Righty” series of photographs. Keeping aperature constant, the shutter speeds for the Administration Building shoot ranged from 4 seconds to an 1/8 of a second.

The next step involved organizing and processing the images, and preparing them for conversion to HDR. I used Adobe Bridge as a virtual thumbnail “lightbox” after downloading the files to my computer. The strength in a program like Bridge is that the EXIF information embedded into each file (containing relevant information like aperture and shutter speed) is visible and offers easy organization and selection of the images to be used. When I selected a series of final images for each environment, I isolated and placed them in folders that I then referred to as “Righties” and “Lefties,” as indicated by where the noticeable feature (in this case the ceiling ornamentation) was located in the photograph. It did not take long into the first test photo shoot to find out that organization was key in terms of managing files and maintaining a naming convention. With dozens of photos to closely examine, Adobe Bridge's instant thumbnail browsing made this process much less of a headache than it could have been without it.

Next, I automated a set of crop and masking actions in Photoshop that I performed on each "lefty" and "righty" directory of images. Cropping the excess out and masking the areas surrounding the mirrored ball made for clean images and Photoshop's Merge to HDR command simplified the composite (Figure 23). This process of selecting, organizing and processing was consistent for each location.

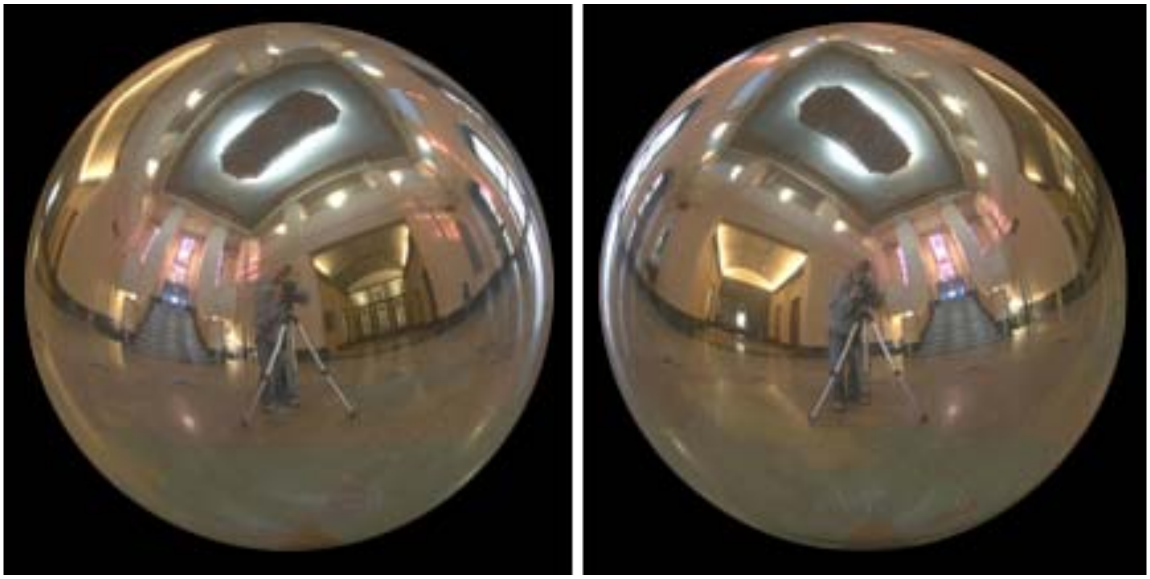

Figure 23. Composite HDR results from the Administration Building photo shoot, reflecting the environment from camera located 90 degrees apart and with the extraneous area surrounding the mirrored ball masked out.

Having completed the merge automation for each of the three environments, I am left with 6 HDR images: two from each environment, each taken from adjacent angles of the mirrored ball. The final step in producing an effective light probe then is compositing the two images from each environment together to create one light probe. HDRShop employs an image distortion operator that can effectively “unwrap” the ‘lefty” image to match corresponding points in the “righty” image, but cannot mask out the photographer.

I used Photoshop to create the mask. By layering “righty” on top of “lefty” and adjusting transparency to evaluate the match, one can see how the each individual image compensates for the other’s problem areas (Figure 24). The masking technique uses pixels whose RGB values are 0 where we wish to use image A, and 1 where we wish to use image B, and an intermediate value when we wish to blend between them. The following set of images describes the operation.

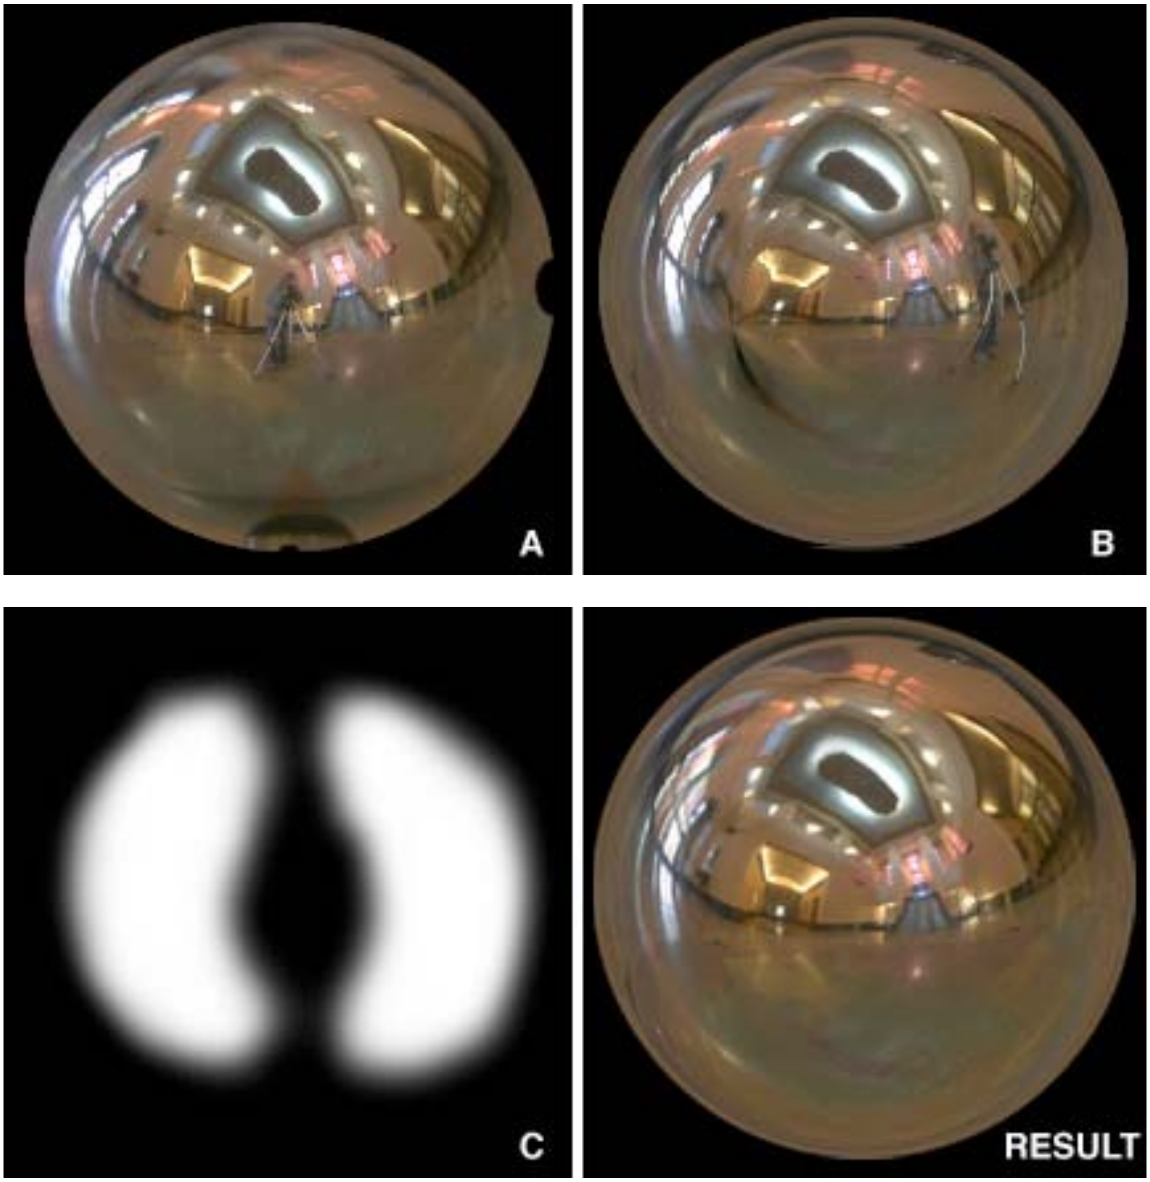

Figure 24. Images used in the composite process. The top images (A) and (B) are digitally transformed and unwrapped by matching coordinated points and layering on top of one another. The black and white image is used to mask out the errors. The composite image is a result of an operation that can be described as $A*C+B*(1-C)$.

Figure 25. The result of a final panoramic transformation. The photographer is no longer present and most warping issues have been resolved.

3. Rendering Multiple Object Views for Two Scenarios

The next step involves taking the light probes I created from each photo shoot and using their HDR information to simulate environmental lighting for the 3d model. This is done within the Autodesk Maya application and Mental Ray renderer by mapping the HDR image (Figure 25) onto a geometric representation of the environment, in this case, a large sphere. Image-based lighting is then calculated by extending rays from the object onto the surface of the image within the environment. The rendering of the object within the three HDR environments involved branching off into two methods, or scenarios - each with similar processes but distinctly different results. Within these scenarios, it is important to recognize a kind of paradigm that is shifted when evaluating the final images. Since the camera provides the “eye” into this virtual world, it is interesting to know how the perception *from* that ‘eye’ can be manipulated. By creatively altering this paradigm - that is, shifting the movement *from the camera*, to the environment *around the camera*, an effect is ultimately had on the "eye." It will be evident in the resulting renders from two scenarios: one that allows for the virtual manipulation of an object within a lit scene, and another with the virtual manipulation of light around a static object.

a. Scenario One: Moving Object, Static Lights

With the 3D model at the center of the sphere, a single camera is placed the desired distance away from the object, so that the rendered image captures as much of the object without any part of the object falling out of frame. This camera is then animated and rotated in increments of 10 degrees in both equatorial and azimuth angles around the model. The rotation of this camera around the object was controlled by a mel script that I wrote to automate the rendering and keep this distance constant. The expression simulates an animated "dome of cameras" in that it translates the camera 10 degrees in yaw for a 360 degree rotation (36 increments) before dropping 10 degrees in pitch (another 19 increments). The animation continues in this fashion, keeping an identical distance from the object, moving every frame for a total of 684 frames. The render specifications use the HDR environment as the source for IBL. The images are finally rendered at 640 X 480 resolution. These resulting rendered images provide the source information needed to import into a QuickTime ™ Object VR authoring software like VR Worx for the last step.

b. Scenario Two: Static Camera, Moving Lights

In this scenario, the 3d model again resides at the center of the environment, with a single camera at a fixed distance from it. As opposed to the first scenario, where *the camera* rotates, this scenario will call for *the HDR environment* to rotate in the exact same increments around the model. In this scenario, I modified the mel expression from the first scenario to apply the same translations of pitch and yaw, but this time to the environment, instead of the camera. The rotation of this environment around the object and camera will now be giving, from the perspective of the camera, the illusion that the environment (and ultimately the lighting provided by it) is actually movable around the object (Figure 26).

Figure 26. Rendered images of sequence from fixed camera showing varying color and lighting effects due to environment rotation. Exterior HDR image used for lighting.

This static camera then renders the exact number of frames as in scenario one and provides the information needed to import into the VRWorx software.

4. Sequencing Images in QuickTime ™ VR Authoring Software

Now that the images from all scenarios have been rendered, the last step would be to take these images and organize them for use in the QuickTime ™ VR authoring software. I used VRWorx to import images and publish Quicktime movies for both scenarios, for all three environments. For scenario one, the desired result is to gather the rendered images of the object to create a *global rotatable model* within a lit space, or a Quicktime object VR movie. For scenario two however, the cumulative effect of the rendered images offers a 360 degree manipulation *of the lights* - providing an opportunity to rotate the *effect of the lighting* around both horizontal and vertical axes, the same way one would rotate an object in a standard object VR movie.

VRWorx allows the user to indicate the rotation specifications for both pitch and yaw, so that the application can know just how many images to import. Because the sequence of rendered frames for each scenario is already specified and in the desired sequence (as dictated by the mel scripts) the import process and publishing is easy and straightforward.

Variable compression settings for the final movie range from no compression to jpeg (quality also be adjusted within each option), giving the user a range of options. Knowing that compression can ultimately affect image quality, I evaluated a number of movies at different compression settings. After I compared the results, I found the jpeg compression at high quality to yield the best image quality with the smallest file size.

CHAPTER V

RESULTS

The following grids of images (Figures 27-32) document the resulting rendered images of the object within the unique environments.

Figure 27. Images from Administration Building render sequence. Scenario 1 uses an animated camera to give the effect of an object moving in a realistically lit space.

Figure 28. Images from Administration Building render sequence. Scenario 2 uses a sequence of images from fixed camera showing varying color and lighting effects due to environment rotation.

Figure 29. Images from Exterior Landscape render sequence. Scenario 1 uses an animated camera to give the effect of an object moving in an realistically lit space.

Figure 30. Images from Exterior Landscape render sequence. Scenario 2 uses a sequence of images from fixed camera showing varying color and lighting effects due to environment rotation.

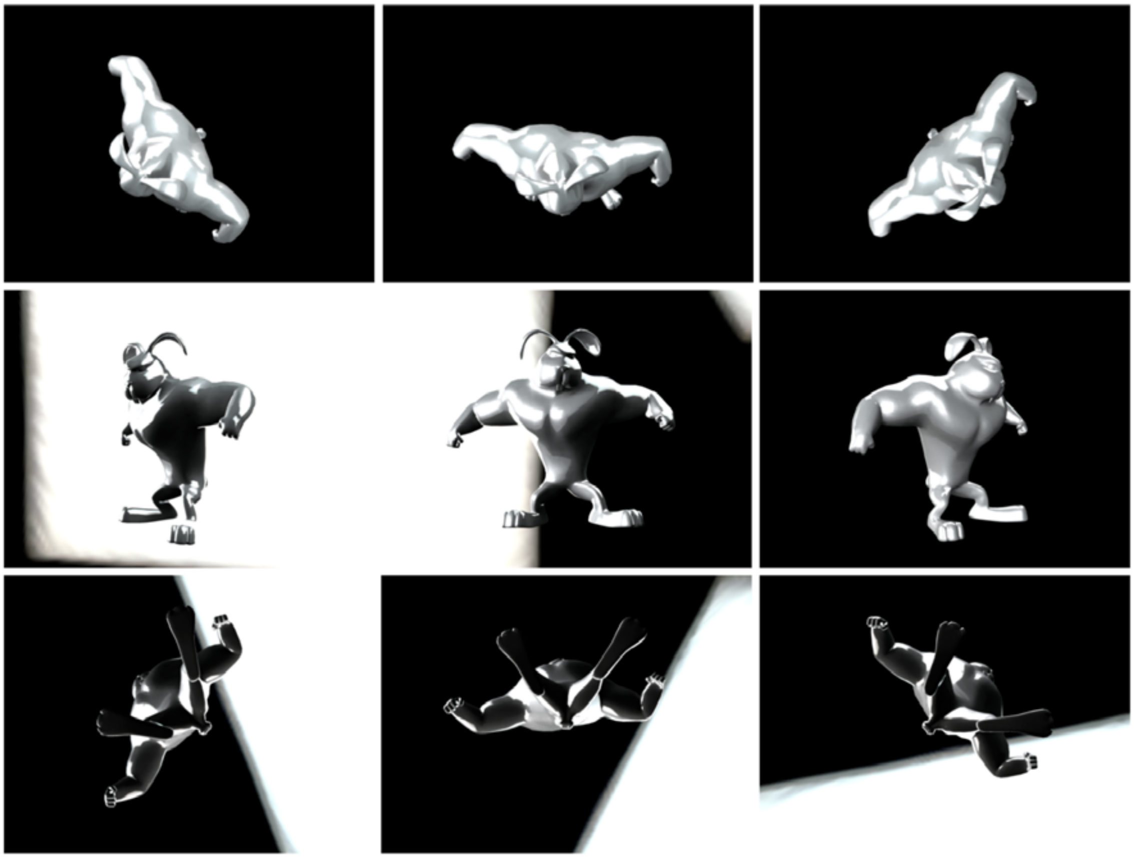

Figure 31. Images from Studio Lighting render sequence. Scenario 1 uses an animated camera to give the effect of an object moving in a realistically lit space.

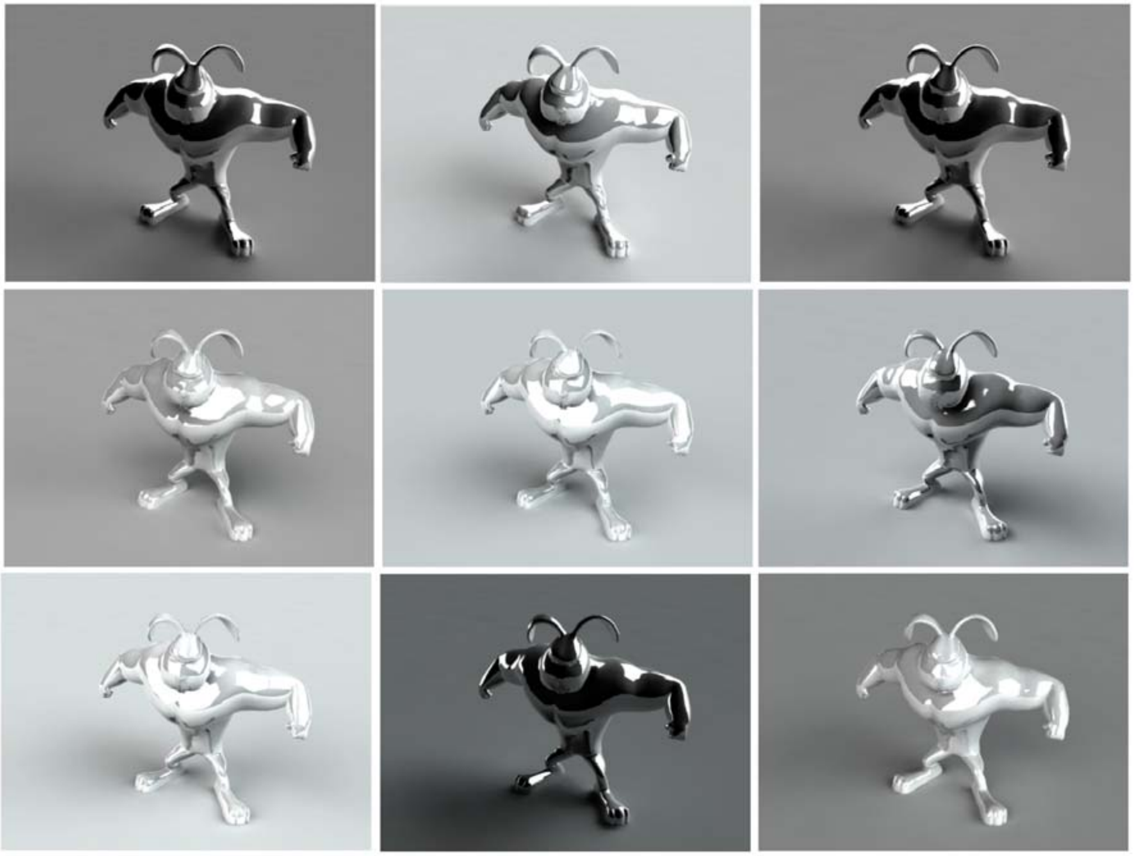

Figure 32. Images from Studio Lighting render sequence. Scenario 2 uses a sequence of images from fixed camera showing varying color and lighting effects due to environment rotation.

CHAPTER VI

CONCLUSION AND FUTURE WORK

The ability to interact with a 3d model on the web will be one that can enhance the user experience, and enable artists to illustrate, in greater depth, the complexity of their work. Even with the most realistic lighting and highest quality render, static images limit the viewer's ability to fully observe a 3d object in space. The methods discussed in this research have provided the artist a new and unique opportunity to offer an enhanced user experience to her audience. Using High Dynamic Range Images to interactively relight the object offers a photorealistic way of observing the object. With the web as the vehicle to communicate an artist's skill in modeling, the user benefits as well, in that they see more of what the artist intended, and more of what they want to see – the object.

In terms of extending this work, I believe there are several communities that could benefit from the methods described in scenario two, for example. In the field of architecture, for example, preliminary rendering of a space with realistic environmental lighting, along with the ability to manipulate that lighting could be very useful for designers and clients to communicate color and design direction. Online merchandise vendors, such as jewelry designers might also enjoy the ability to have customers fully inspect their work.

Other possibilities for future development might include the *animation* of High Dynamic Range Image-Based Lighting, with emphasis on such features as light exposure and depth of field. Offering interactivity in terms of animating these features can extend the user experience with an object on the web.

REFERENCES

[1] T. Malzbender, D. Gelb, and H. Wolters, “Polynomial Texture Maps.” *Proc. ACM SIGGRAPH 2001*, pp. 9-16, 2001.

[2] C. Lyon, “Viewing Polynomial Texture Maps Using Java.” Harvard Extension School. pp. 4-12, 2004. http://materialobjects.com/ptm/overview.pdf, Accessed January 2006.

[3] C. Gloman and T. Le Tourneau, *Placing Shadows*. Focal Press. 2000.

[4] J. Birn, *Digital Lighting and Rendering, 2nd Edition*. New Riders Publishing. 2006.

[5] D. Brooker. *Essential CG Lighting Techniques.* Focal Press. 2002.

[6] Image of St. Paul’s Cathedral. http:www.commons.wikimedia.org/Image:Old_saint_pauls_1.jpg Accessed January 2007.

[7] P. Debevec, “Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography.” *Proc. ACM SIGGRAPH 1998*, pp. 223-231, 1998.

[8] P. Debevec, “A Tutorial on Image-Based Lighting,” *IEEE Computer Graphics and Applications 2002,* Vol 21 Issue 3, pp.171-178, 2002.

[9] S.E. Chen, “QuickTime VR - An Image-Based Approach to Virtual Environment Navigation.” *Proc. ACM SIGGRAPH 1998*, pp. 129-138, 1998.

[10] R. Kleiser and M. Baker, *Flight of the Navigator*. Walt Disney Pictures, 1986.

[11] J. Cameron, *Terminator 2: Judgement Day*. Pacific Western Productions, 1991.

[12] B. Singer and Z. Penn, *X2:X-men United.* Marvel Enterprises, 2003.

[13] *Half-Life 2:Lost Coast*. Valve, Microsoft Windows platform, 2005.

[14] *Fight Night 3*. EA Sports, Xbox 360, 2006.

[15] M. Kawasa, Software demo “Real Time High Dynamic Range Image-Based Lighting,” http://www.daionet.gr.jp/~masa/rthdribl/, Accessed June 2006.

[16] S. Crawford, “Extending Quicktime’s Object VR Interactivity.” BioGraphix, LLC. Fort Collins, Colorado, 2001.

APPENDIX 1

Several interactive QTVR movies discussed in this thesis can be viewed in the “movies” file.

VITA

Name: Tamara Cuellar

Address: 108 North Brown Street,
Bryan, Texas 77802

Email Address: tamara.cuellar@gmail.com

Education: B.E.D., Texas A&M University 2002
M.S., Visualization Sciences, Texas A&M University, 2008