1

Introduction

"Make it so."

—Captain Jean-Luc Picard, 2305–

EVERYDAY we are exposed to an immense amount of visual information. The natural world is filled with beautiful and complex phenomena which have sparked the curiosity of humankind for millennia. Throughout history we have struggled to find new ways to understand, manipulate, and, in fact, create the world around us for purposes ranging from communication, to medicine, to art and entertainment.

The visual arts and sciences have been following parallel paths in our quest for understanding the nature of the world around us. These seemingly distinct disciplines have been inextricably linked in the development of this understanding. Artists have readily utilized existing scientific and mathematical knowledge to create more convincing works of art. For instance, the development of contrapposto by classical Greek sculptors such as Polykleitos in the 4th century BCE relied heavily on a system of ideal mathematical proportions initially investigated by Pythagoras and Euclid [Tanner, 2006]. The understanding of the behavior of light perhaps bridges these two disciplines most since it is light which determines the visual appearance of the natural world. Often discoveries were made independently in both fields, like when Renoir and Monet rediscovered indirect illumination by observing that shadows are not brown or black, but the color of the surrounding environment. At other times the physical sciences have only been able to explain phenomena years after artists had developed an intuitive understanding. Leonardo Da Vinci, for instance, introduced *aerial perspective* in his *Treatise on Painting* [1651] to describe the effects of light scattering in the atmosphere centuries before a scientific understanding of volumetric scattering was established. In fact, these disciplines are so intertwined that historically, practitioners of one discipline were often heavily ingrained in or influenced by the other. This quality is perhaps embodied most by Da Vinci, but other notable examples include Archimedes, Su Song, Michelangelo, Piero della Francesca, Albrecht Dürer, Vermeer, and M.C. Escher. Recently we have seen a convergence between many aspects of these disciplines with the introduction of photography, and now the computer.

Today, the understanding fueled by these two disciplines, combined with the advent of the computer, has provided us with unprecedented potential for simulating our world. The introduction of computers has spawned a new discipline at the crossroads of these two fields: computer graphics. The goal of computer graphics is to develop algorithms that allow a user to digitally synthesize and manipulate visual content. As in art, the visual content we are typically interested in synthesizing relates to the natural world and therefore depends on the physical sciences.

Physically based rendering is concerned with generating, or rendering, images by modeling the way the physical world behaves. A related goal is photorealistic rendering, which strives to synthesize images that are indistinguishable from photographs. Both of these disciplines are primarily concerned with simulating the physical behavior of light as it interacts in virtual environments in order to generate realistic images. Applications of realistic image synthesis include design, architecture, and education. In particular, we have recently witnessed an explosion in the widespread use of computer graphics and realistic image synthesis in journalism, games, television, movies, as well as virtually all forms of media. This is made evident by the entertainment industry's extensive use of computer generated imagery and special effects.

As the field has matured, computer generated images in many of the aforementioned areas are becoming increasingly complex, and the desire for unparalleled realism is always increasing. Theoretical understanding combined with increased computational power has allowed us to simulate virtually every observable or imaginable phenomenon. Unfortunately, many of the most interesting phenomena are also some of the most computationally intensive, fueling the need for more efficient, general purpose algorithms. One example of particular interest to the computer graphics community is simulating participating media.

Participating, or scattering, media influences virtually everything we see around us: from the color of the sky at sunset, to the appearance of milk, to the reduced visibility we encounter in fog or rain. Participating media gives rise to pyrotechnic effects such as fire, smoke, clouds, and explosions, which are of particular interest to the special effects industry. Furthermore, paper, marble, and virtually all organic materials exhibit subsurface scattering, which is a special case of the volumetric scattering inherent to participating media. See Figure 1.1 for a few examples of the striking and visually complex effects caused by participating media. Unfortunately, rendering images with participating media is an especially difficult problem. The behavior of light transport becomes significantly more complex when participating media is considered, since photons may interact at any point within the volume instead of just at surface boundaries. For this reason, many rendering algorithms assume that all objects are in a vacuum in order to simplify the lighting simulation needed. At times this can be a reasonable approximation, but there are many situations when the effects from participating media have a significant visual impact (see Figure 1.1). These are not pathological cases by any means, and are in fact of considerable interest to many disciplines outside of computer graphics, including heat transfer and neutron transport.

Due to increased computational power and the ubiquity of participating media, there has been a steady increase in techniques capable of rendering volumetric effects. Due to the complexity of the problem, many of the most successful techniques have made simplifying assumptions about the medium being simulated. For example, homogeneous media with a high scattering albedo, such as milk, can be modeled accurately using a diffusion approximation. Diffusion was originally developed to explain the transport of neutrons in nuclear reactors [Glasstone and Sesonske, 1955], but also leads to very efficient rendering algorithms [Stam, 1995; Jensen et al., 2001b]. On the other extreme, the assumption of non-scattering but absorbing media makes the light transport problem significantly more tractable. Tenuous media with a low scattering albedo can also be effectively approximated by assuming only a single scattering event between the light source and the eye [Blinn, 1982; Nishita et al., 1987]. In isotropic media, which scatters light equally in all directions, the lighting simulation can be computed in a preprocess and rapidly visualized from any arbitrary viewpoint [Rushmeier and Torrance, 1987]. Many efficient

Figure 1.1: Volumetric scattering due to participating media is responsible for the appearance of a number of striking visual effects including aerial or atmospheric perspective (top), shafts of light (left), clouds and crepuscular rays (right), as well as interstellar dust and nebula (bottom). The bottom image is courtesy of T.A.Rector (NOAO/AURA/NSF) and the Hubble Heritage Team (STScI/AURA/NASA).

algorithms have been developed which work exceptionally well for specific classes of scattering materials. Unfortunately, not all participating media effects fall within these narrow confines, in which case we must resort to more general, but typically less efficient, approaches.

The goal of this dissertation is to develop efficient, general-purpose methods for solving light transport in scenes containing participating media. To meet our goal of generality, we concentrate on Monte Carlo methods. These methods can easily handle arbitrary medium properties and make few underlying assumptions about the scenes being simulated. Unfortunately, unbiased Monte Carlo techniques suffer from long computation times and significant noise. To satisfy our goal of efficiency we focus on caching-based Monte Carlo methods. These methods exploit the often smooth nature of illumination by computing and caching lighting at a small set of locations. These cache values are then reused to estimate lighting at all locations needed to render the image. Aside from this property, the methods developed in this dissertation make no assumptions about the properties of the medium being rendered.

1.1 Summary of Original Contributions

The work in this dissertation builds upon the irradiance caching and photon mapping methods. We outline our major contributions below.

Volumetric Radiance Caching. Our first contribution is a novel radiance caching method for efficiently rendering participating media using Monte Carlo ray tracing. The *volumetric radiance caching* method handles all types of light scattering, including anisotropic scattering, and it works in both homogeneous and heterogeneous media. Our method gains efficiency by sparsely sampling and interpolating radiance within the medium. The method provides several orders of magnitude speedup compared to path tracing and produces higher quality results than volumetric photon mapping. Furthermore, it is view-driven and well suited in large scenes where methods such as photon mapping become costly.

Accurate Illumination Gradients in Participating Media. Another key contribution in this dissertation is a technique for computing gradients of radiance evaluated in the presence of participating media. These gradients describe how the incident radiance field changes with respect to translation both on surfaces and within the medium. Our gradient derivations take the full path of the scattered light into account including the changing properties of the medium in the case of heterogeneous media. We compute gradients for single scattering from lights and surfaces and for multiple scattering. For surface irradiance, we additionally take into account the effects of visibility on the gradient. We apply our gradient derivations to provide higher quality interpolation within both the irradiance caching and volumetric radiance caching methods.

The Beam Radiance Estimate. We also introduce a new photon mapping method for more efficiently simulating the scattering of light within participating media. We present a theoretical reformulation of volumetric photon mapping, which allows us to develop a novel photon gathering technique for participating media. This approach explicitly estimates the accumulated radiance along the length of an entire ray instead of sampling the radiance at individual points. This optimization provides a significant savings making it substantially faster than regular volumetric photon mapping.

1.2 Organization of the Dissertation

The dissertation is divided into nine chapters. In the remaining chapters, we first focus on background material for the first two contributions above. In Chapter 2 we introduce the fundamentals of light transport in the absence of participating media. We describe irradiance caching and derived methods, which form the inspiration for our volumetric radiance caching approach, in Chapter 3. We also provide derivations which are missing from the literature in previous work. In Chapter 4 we extend our theoretical understanding by describing how light interacts with participating media. Based on this foundation, in Chapter 5 we develop the volumetric radiance caching method in detail and derive expressions for computing the gradient of in-scattered radiance within participating media. In Chapter 6 we extend our derivations for illumination gradients by considering the behavior of surface irradiance in the presence of participating media and occlusions. We then move on to the photon mapping method, which we describe in detail in Chapter 7. After highlighting the shortcomings of the volumetric photon mapping method, we present a theoretical reformulation in Chapter 8, which allows us to develop the beam radiance estimate. Chapter 9 summarizes our contributions and discusses possible avenues of future work.

In the body of this dissertation, we use a number of mathematical tools while providing only minimal overview and no derivations of associated properties. We therefore provide three appendices which cover, in more detail, these mathematical concepts. Appendix A provides a detailed introduction to Monte Carlo integration upon which all our methods are based. We cover spherical harmonics in Appendix B and review their benefitial properties and associated operations. Lastly, we review the theory behind density estimation methods such as photon mapping in Appendix C.