4

Light Transport in Participating Media

"Thus, if one is to be five times as distant, make it five times bluer."

—Leonardo Da Vinci, 1452–1519

IN previous chapters we assumed that all lighting interactions occurred at surfaces. In particular, the rendering equation from Chapter 2 assumes that the radiance leaving a surface remains unchanged until it hits another surface. The property that radiance stays constant between these two points is expressed in Equation 2.10. However, this relationship is only true when the surfaces are embedded within a vacuum, since photons are able to travel unobstructed as they propagate from one surface to another. In reality, however, most of the scenes we observe from day to day do not satisfy this constraint. The space between objects is typically occupied by some host medium (e.g. air, water) which additionally contains impurities and microscopic suspended particles. In these situations, the medium *participates* in the lighting interactions as light travels between surfaces.

If there are very few particles within the medium, such as over short distances within clean air, the assumption that light travels unobstructed can be a reasonable approximation. However, even clean air scatters light (this is the reason the sky is blue) and as light travels over longer distances, such as in open outdoor scenes, more interactions with the particles are likely to occur.

In this chapter we expand our model for the behavior of light by considering effects caused by participating media. We relax the constraints imposed by Equation 2.9 and 2.10 by deriving expressions for how radiance changes as it travels through participating media. This

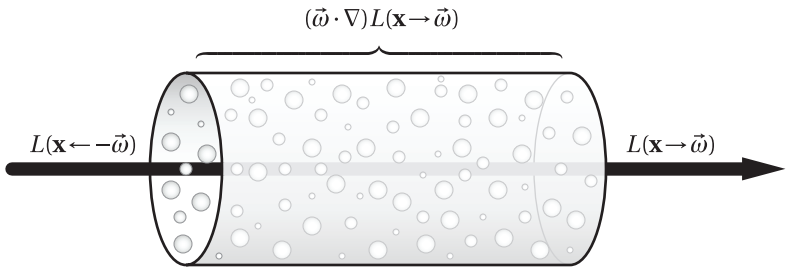

Figure 4.1: We treat participating media as a collection of microscopic scattering particles. When lights travels through the medium, a change of radiance, $(\vec{\omega}\cdot\nabla)L(\mathbf{x},\vec{\omega})$, may occur as the photons interact with the particles.

generalization allows us to simulate lighting within clouds, murky water, fog and any other medium which participates in the lighting interactions.

4.1 Assumptions About Scattering Media

In this dissertation, and most computer graphics applications, assumptions are made about the properties of the scattering media in order to more easily derive expressions about the behavior of light. In particular, we assume that the participating media can be modeled as a collection of microscopic particles (see Figure 4.1). Since the particles are microscopic and randomly positioned, we do not represent each individual particle in the lighting simulation. Instead, we consider the aggregate probabilistic behavior as light travels through the medium. Moreover, these particles are assumed to be spaced far apart relative to the size of an individual particle. This assumption implies that as a photon travels through the medium and interacts at a particle, this interaction is *statistically independent* from the outcome of subsequent interaction events.

4.2 Light Interaction Events

When a photon travels through a collection of microscopic particles, it may either miss all the particles and continue unaffected, or it may interact with some of the particles. The probability that an interaction does occur is related to the extinction coefficient, σ_t (units [1/m]), of the medium. This quantity depends on the density and size of the particles within the medium.

When an interaction occurs, two things may happen: the photon may be absorbed by the particle (by being converted to another form of energy, such as heat), or the photon may be scattered in another direction. The relative probabilities of these two events is given by the absorption, σ_a, and the scattering, σ_s, coefficients, and $\sigma_t = \sigma_a + \sigma_s$ is the extinction coefficient. Either of these two events lead to a change of radiance along the ray.

4.2.1 Extinction

Consider a thin beam of light traveling through a medium in direction $\vec{\omega}$ (see Figure 4.1). The number of photons entering this beam is proportional to the incident radiance $L(\mathbf{x} \to \vec{\omega})$ at the start of the beam $\mathbf{x}$. At each small step Δt along the beam, some fraction of the photons will interact with the medium and become absorbed. If the absorption coefficient within the segment is $\sigma_a(\mathbf{x} + t\vec{\omega})$, then a fraction $\sigma_a(\mathbf{x} + t\vec{\omega})\,\Delta t$ of the photons will be absorbed. Hence, the number of photons exiting this segment can be expressed as:

$$L((\mathbf{x} + t\vec{\omega}) \to \vec{\omega}) = L(\mathbf{x} \to \vec{\omega})\,(1 - \sigma_a(\mathbf{x} + t\vec{\omega})\,\Delta t). \tag{4.1}$$

By rearranging terms, we can determine the change in outgoing radiance between the start and end of this segment:

$$\frac{L((\mathbf{x} + t\vec{\omega}) \to \vec{\omega}) - L(\mathbf{x} \to \vec{\omega})}{\Delta t} = -\sigma_a(\mathbf{x})\,L(\mathbf{x} \to \vec{\omega}). \tag{4.2}$$

Taking the limit as $\Delta t \to 0$ computes the derivative. The differential change in radiance due to *absorption* as light moves along a direction $\vec{\omega}$ is:

$$\lim_{\Delta t \to 0} \left(\frac{L((\mathbf{x} + t\vec{\omega}) \to \vec{\omega}) - L(\mathbf{x} \to \vec{\omega})}{\Delta t} \right) = (\vec{\omega} \cdot \nabla_a)\,L(\mathbf{x} \to \vec{\omega}) = -\sigma_a(\mathbf{x})\,L(\mathbf{x} \to \vec{\omega}), \tag{4.3}$$

where ∇^a denotes the gradient due to absorption, and $(\vec{\omega} \cdot \nabla_a)$ computes the directional derivative along $\vec{\omega}$.

At each step, the radiance may also be reduced due to photons being scattered into other directions. The probability of this happening within the segment is determined by the

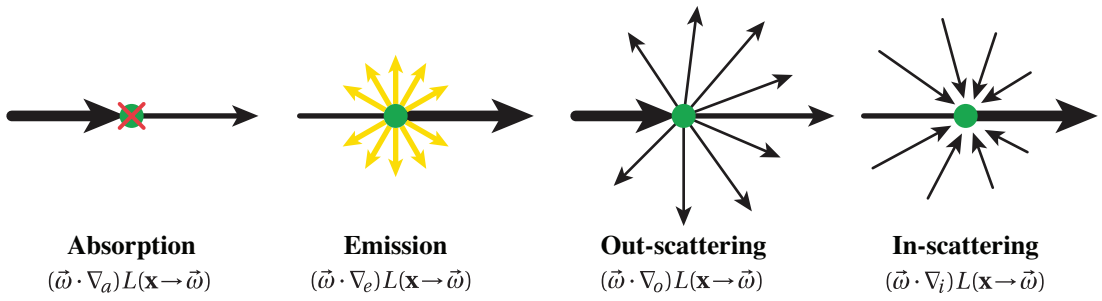

Figure 4.2: As light travels through a participating medium the radiance may change as a result of four different types of interactions: absorption, emission, out-scattering and in-scattering.

scattering coefficient σ_s of the medium. Following the same derivation we can express the change in radiance due to *out-scattering* as:

$$(\vec{\omega}\cdot\nabla_o)L(\mathbf{x}\rightarrow\vec{\omega}) = -\sigma_s(\mathbf{x})\,L(\mathbf{x}\rightarrow\vec{\omega}). \tag{4.4}$$

The net loss of radiance due to either absorption or out-scattering is called extinction, and it can be expressed as:

$$\begin{aligned}(\vec{\omega}\cdot\nabla_t)L(\mathbf{x}\rightarrow\vec{\omega}) &= -(\sigma_a(\mathbf{x})+\sigma_s(\mathbf{x}))\,L(\mathbf{x}\rightarrow\vec{\omega})\\ &= -\sigma_t(\mathbf{x})\,L(\mathbf{x}\rightarrow\vec{\omega}).\end{aligned} \tag{4.5}$$

The differential equation above can be solved to find the *beam transmittance*, or simply transmittance, denoted T_r. Transmittance gives the fraction of photons that can travel unobstructed between two points in the medium along a straight light. The remaining fraction accounts for photons that have either been absorbed or out-scattered in a different direction. By integrating Equation 4.5 we can express the transmittance as:

$$T_r(\mathbf{x}'\leftrightarrow\mathbf{x}) = e^{-\tau(\mathbf{x}'\leftrightarrow\mathbf{x})}. \tag{4.6}$$

The term in the exponent, τ, is called the *optical thickness* or *optical depth*, and comes about by

integrating the effect of extinction along a line segment:

$$\tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = \int_0^d \sigma_t(\mathbf{x} + t\vec{\omega})\, dt, \tag{4.7}$$

where $\mathbf{x} + t\vec{\omega}$ with $t \in [0, d]$ parametrizes the line segment between $\mathbf{x}$ and $\mathbf{x}'$. In a homogeneous medium, where σ_t is a constant, τ can be trivially computed as:

$$\tau(\mathbf{x}' \leftrightarrow \mathbf{x}) = d\sigma_t. \tag{4.8}$$

The transmittance is always between zero and one. When only absorption and out-scattering are considered, the transmittance can be used to compute the *reduced radiance*,

$$L(\mathbf{x} \leftarrow \vec{\omega}) = T_r(\mathbf{x}' \leftrightarrow \mathbf{x})\, L(\mathbf{x}' \rightarrow -\vec{\omega}), \tag{4.9}$$

which describes how radiance is reduced as it travels from $\mathbf{x}$ to $\mathbf{x}'$. This relation replaces Equation 2.10 if extinction effects are considered and is also known as *Beer's Law* [Bouguer, 1729].

The transmittance between a point and itself is always one: $T_r(\mathbf{x} \leftrightarrow \mathbf{x}) = 1$. Another useful property is that the transmittance between three points along a line is multiplicative:

$$T_r(\mathbf{x} \leftrightarrow \mathbf{x}'') = T_r(\mathbf{x} \leftrightarrow \mathbf{x}')\, T_r(\mathbf{x}' \leftrightarrow \mathbf{x}''), \tag{4.10}$$

for all points $\mathbf{x}'$ between $\mathbf{x}$ and $\mathbf{x}''$.

4.2.2 In-Scattering and Emission

In addition to loss of radiance, radiance may increase as it travels through the medium. The same process that causes the radiance to reduce along $\vec{\omega}$ due to out-scattering may cause it to increase due to *in-scattering* from other directions. In effect, scattering causes a loss of radiance along the incoming direction and a corresponding increase in radiance into the scattered direction. The change in radiance due to in-scattering can be derived in a similar fashion to

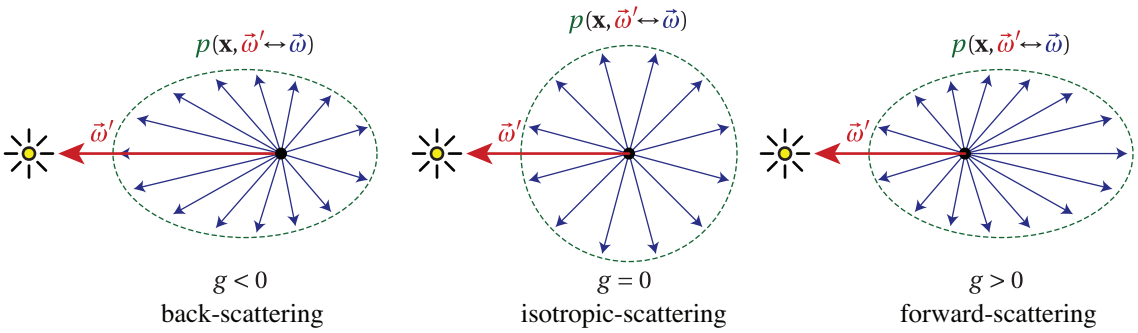

Figure 4.3: The phase function describes the angular distribution of light scattering at any point $\mathbf{x}$ within participating media. In the simplest case, light is scattered equally in all directions (middle). Many natural materials, however, scatter light preferentially in the backward (left) or forward (right) direction.

extinction and out-scattering above, but instead depends on L_i, the in-scattered radiance:

$$(\vec{\omega} \cdot \nabla_i) L(\mathbf{x} \rightarrow \vec{\omega}) = \sigma_s(\mathbf{x})\, L_i(\mathbf{x} \rightarrow \vec{\omega}), \tag{4.11}$$

We explain in-scattered radiance in more detail in the next section.

Media, such as fire, may also emit radiance, L_e, by spontaneously converting other forms of energy into visible light. This emission leads to a gain in radiance expressed as:

$$(\vec{\omega} \cdot \nabla_e) L(\mathbf{x} \rightarrow \vec{\omega}) = \sigma_a(\mathbf{x}) L_e(\mathbf{x} \rightarrow \vec{\omega}). \tag{4.12}$$

4.2.3 Medium Properties

The absorption σ_a and scattering coefficients σ_s, along with the phase function p discussed in the next section, fully define the local behavior of light in a participating medium. If the medium coefficients are constant throughout the medium, then the medium is said to be *homogeneous*. This can be an acceptable approximation in the case of uniform fog for instance. Most media is *heterogeneous*, which means that the absorption and scattering coefficients may vary throughout the medium.

Another descriptive property of the medium can be derived from the *unitless* ratio between the scattering and extinction coefficients, $\frac{\sigma_s}{\sigma_t}$. This is called the scattering *albedo*, and it gives the probability of a photon scattering at a particular location in the medium. An albedo of

zero means that the particles of the medium do not scatter light, while an albedo of one means that the particles do not absorb light. The albedo has a similar meaning as the average reflectivity of a surface.

4.3 In-Scattered Radiance and the Phase Function

4.3.1 In-Scattered Radiance

In-scattered radiance $L_i(\mathbf{x}\rightarrow\vec{\omega})$ represents all the light which is scattered into the thin beam along direction $\vec{\omega}$[1]. The computation of in-scatter radiance involves integrating incident radiance over the whole sphere of directions, $\Omega_{4\pi}$,

$$L_i(\mathbf{x}\rightarrow\vec{\omega}) = \int_{\Omega_{4\pi}} p(\mathbf{x},\vec{\omega}'\rightarrow\vec{\omega})\, L(\mathbf{x}\leftarrow\vec{\omega}')\, d\vec{\omega}', \tag{4.13}$$

where $p(\mathbf{x},\vec{\omega}'\rightarrow\vec{\omega})$ is the phase function describing the angular distribution of light scattering at a point in the medium (see Figure 4.3). Due to this function, in-scattered radiance is usually not isotropic (unlike emission and absorption), but depends on the outgoing direction.

The computation of in-scattered radiance closely parallels the computation of reflected radiance on surfaces. Note the similarity of in-scattered radiance in Equation 4.13 with reflected radiance in Equation 2.14. The distinction is that the BRDF is replaced with the phase function, there is no cosine foreshortening term (since no surface is present), and the integration is performed over the whole sphere of directions instead of the hemisphere.

4.3.2 Properties of the Phase Function

1. **Reciprocity.** Just like the BRDF, the phase function respects Helmholtz's law of reciprocity, so $p(\mathbf{x},\vec{\omega}'\rightarrow\vec{\omega}) = p(\mathbf{x},\vec{\omega}'\leftarrow\vec{\omega})$. To indicate this we use the following notation:

$$p(\mathbf{x},\vec{\omega}'\leftrightarrow\vec{\omega}). \tag{4.14}$$

[1] Note that in-scattered radiance represents exitant, not incident, radiance. Hence, the notation $L_i(\mathbf{x}\rightarrow\vec{\omega})$ is used.

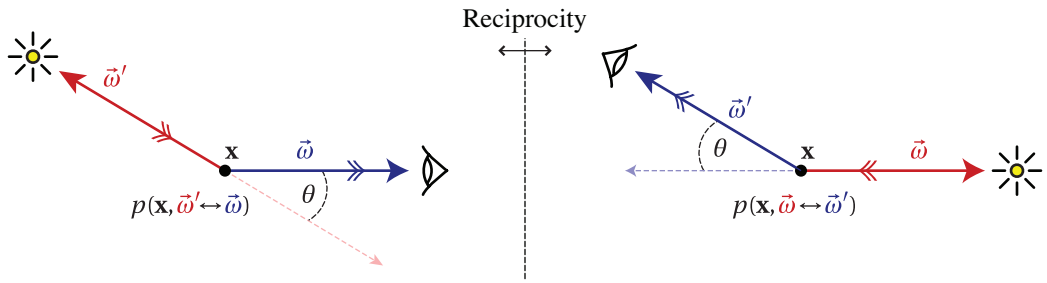

Figure 4.4: The phase function obeys Helmholtz's reciprocity principle: the value of the function remains unchanged if the direction of light is reversed. This is equivalent to swapping the location of the viewer and the light source (right), i.e., $p(\mathbf{x}, \vec{\omega}' \rightarrow \vec{\omega}) = p(\mathbf{x}, \vec{\omega} \rightarrow \vec{\omega}')$.

Note that in this dissertation the incident and outgoing directions for the phase function both point away from the evaluation location $\mathbf{x}$[2]. Additionally, the phase function often depends only on the angle θ between the incoming and outgoing directions, in which case it can be expressed simply as $p(\mathbf{x}, \theta)$, where $\theta = \cos^{-1}(-\vec{\omega} \cdot \vec{\omega}')$. An illustration of these concepts is provided in Figure 4.4.

2. **Normalization.** Unlike the BRDF the phase function has units of [1/sr]. It is also assumed normalized based on the derivations in this dissertation. This means that it always integrates to exactly one over the sphere of directions[3]:

$$\int_{\Omega_{4\pi}} p(\mathbf{x}, \vec{\omega}' \leftrightarrow \vec{\omega}) \, d\vec{\omega}' = 1, \quad \forall \vec{\omega}. \tag{4.15}$$

4.3.3 Examples of Phase Functions

Many different phase functions have been developed in the literature. We briefly explain some of the phase functions commonly used in computer graphics.

The Isotropic Phase Function

The equivalent to diffuse reflection is perfect isotropic scattering. The isotropic phase function is a constant and, based on the constraint in Equation 4.15, it is straightforward to show

[2] Some texts instead follow the convention that one vector points into $\mathbf{x}$ and the other vectors points away from $\mathbf{x}$.

[3] Some texts follow different derivations and instead define the phase function to integrate to the albedo, σ_s/σ_t, or to 4π.

that it equals

$$p_I(\mathbf{x}, \vec{\omega}' \leftrightarrow \vec{\omega}) = \frac{1}{4\pi}. \tag{4.16}$$

Isotropic media scatters light uniformly in all directions regardless of the incoming direction.

The Henyey-Greenstein Phase Function

One of the most commonly used non-isotropic phase functions is the Henyey-Greenstein phase function [1941], which was empirically derived to model the scattering of light by intergalactic dust. Due to its simplicity the Henyey-Greenstein phase function has also been successfully applied for simulating many other scattering materials, such as oceans, clouds, and skin.

The function only depends on the angle θ and is defined as:

$$p_{HG}(\mathbf{x}, \theta) = \frac{1 - g^2}{4\pi(1 + g^2 - 2g\cos\theta)^{1.5}}. \tag{4.17}$$

The single parameter, $g \in [-1, 1]$, determines the relative strength of forward and backward scattering.

The asymmetry parameter g has a physical meaning—it represents the average cosine of the scattered directions. Positive values of g give forward scattering, and negative values give backwards scattering. A value of $g = 0$ results in isotropic scattering. A convenient side-effect of this definition of g is that more complex scattering functions can easily be approximated by the Henyey-Greenstein function by computing g as:

$$g(\mathbf{x}) = \int_{\Omega_{4\pi}} p(\mathbf{x}, \vec{\omega}' \leftrightarrow \vec{\omega}) \cos\theta \, d\vec{\omega}'. \tag{4.18}$$

The Schlick Phase Function

Though the Henyey-Greenstein phase function is intuitive and flexible, the computation of the fractional exponent in the denominator is relatively costly. In computer graphics, the exact shape of the Henyey-Greenstein phase function is usually not essential. This led Blasi and

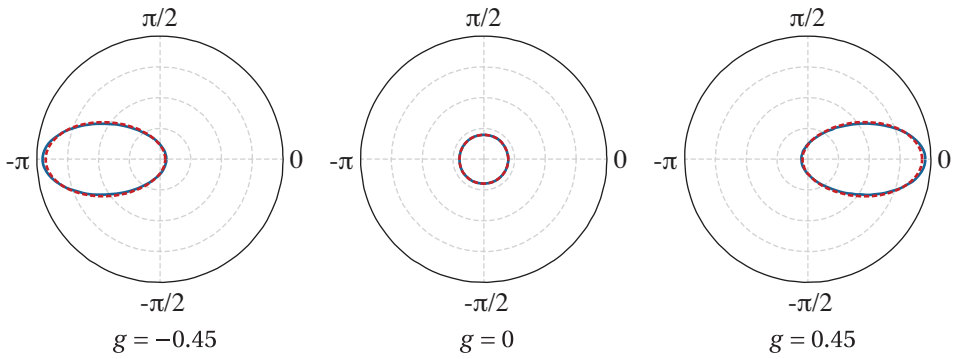

Figure 4.5: Polar plots visualizing the Henyey-Greenstein and Schlick phase functions as functions of θ. The Schlick phase function (red) closely approximates the shape of the Henyey-Greenstein phase function (blue) and is more efficient to evaluate.

colleagues to develop a similar phase function, which is more efficient to evaluate [1993]. The resulting Schlick phase function approximates the shape of the Henyey-Greenstein function as an ellipsoid:

$$p_S(\mathbf{x},\theta) = \frac{1-k^2}{4\pi(1+k\cos\theta)^2}, \tag{4.19}$$

where $k \in [-1,1]$ acts similarly to the g parameter and controls the preferential scattering direction. As with g, total backward scattering is obtained with $k=-1$, $k=1$ gives total forward scattering, and $k=0$ corresponds to isotropic scattering. Pharr and Humphreys [2004] found that for intermediate values of k, the polynomial equation

$$k \approx 1.55g - 0.55g^3 \tag{4.20}$$

can be used to give accurate approximations to the Henyey-Greenstein phase function. The Henyey-Greenstein and Schlick phase functions are shown in Figure 4.5 for three different values of g.

Rayleigh Scattering

In 1871, Rayleigh derived expressions for the scattering of light off molecules of air [1871]. Rayleigh scattering is an approximation for the behavior of light as it scatters in media composed

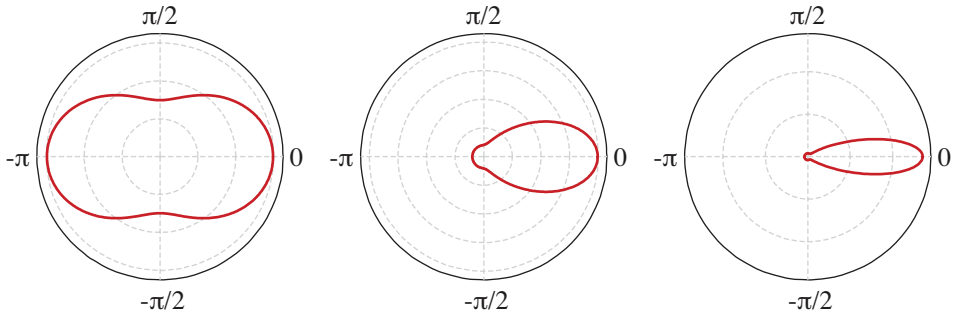

Figure 4.6: The Rayleigh scattering phase function (left) consists of a symmetric forward and backward scattering contribution. The Hazy (middle) and Murky (right) phase functions are approximations to the behavior of scattering in foggy atmospheres dictated by Lorenz-Mie theory.

of particles up to about a tenth of the wavelength of light. The model is derived by assuming a random distribution of very small specular spheres, and the resulting phase function is fairly simple:

$$p_R(\mathbf{x},\theta) = \frac{3}{16\pi}(1+\cos^2\theta). \tag{4.21}$$

The Rayleigh phase function is visualized in Figure 4.6.

Rayleigh scattering is also highly dependent on the wavelength of light. This is captured by a wavelength-dependent scattering coefficient,

$$\sigma_s = \frac{2\pi^5}{3}\frac{d^6}{\lambda^4}\left(\frac{n^2-1}{n^2+2}\right)^2, \tag{4.22}$$

where λ is the wavelength of light, and d and n are the diameter and refractive index of the particles, respectively. This wavelength dependence of Rayleigh scattering is what makes the sky blue and sunsets red.

Lorenz-Mie Theory

Rayleigh scattering is just an approximation that works when the size of the particles is small relative to the wavelength. When the particles are comparable to the wavelength of light, such as water droplets in fog, the more complicated Lorenz-Mie theory of scattering becomes applicable [Lorenz, 1890; Mie, 1908]. Lorenz-Mie theory can be used to derive phase functions

for a homogeneous collection of spherical particles where any ratio of diameter to wavelength is allowed. Mie scattering is based on more general theory derived from Maxwell's equations.

Nishita et al. [1987] provide two empirically derived approximations to the complicated Lorenz-Mie scattering functions for foggy atmospheres, one for "hazy" atmospheres, and one for "murky" atmospheres:

$$p_{MH}(\mathbf{x},\theta) = \frac{1}{4\pi}\left(\frac{1}{2} + \frac{9}{2}\left(\frac{1+\cos\theta}{2}\right)^8\right) \tag{4.23}$$

$$p_{MM}(\mathbf{x},\theta) = \frac{1}{4\pi}\left(\frac{1}{2} + \frac{33}{2}\left(\frac{1+\cos\theta}{2}\right)^{32}\right). \tag{4.24}$$

These phase functions are shown in Figure 4.6.

More recently, generalizations of Lorenz-Mie theory have been presented which can compute the scattering properties of a collection of non-spherical particles within an absorbing host medium [Frisvad et al., 2007].

4.4 The Volume Rendering Equation

4.4.1 The Radiative Transfer Equation

Given the four scattering events described in the previous sections we can begin to form a complete model of how light behaves in a participating medium. By combining Equations 4.5, 4.11, and 4.12, the total change in radiance along the ray at a position $\mathbf{x}$ can be expressed as

$$(\vec{\omega}\cdot\nabla)L(\mathbf{x}\rightarrow\vec{\omega}) = -\underbrace{\underbrace{\underbrace{\sigma_a(\mathbf{x})L(\mathbf{x}\rightarrow\vec{\omega})}_{\text{absorption}} - \underbrace{\sigma_s(\mathbf{x})\,L(\mathbf{x}\rightarrow\vec{\omega})}_{\text{out-scattering}}}_{\text{extinction}} + \underbrace{\sigma_a(\mathbf{x})\,L_e(\mathbf{x}\rightarrow\vec{\omega})}_{\text{emission}} + \underbrace{\sigma_s(\mathbf{x})\,L_i(\mathbf{x}\rightarrow\vec{\omega})}_{\text{in-scattering}}}. \tag{4.25}$$

This equation is known as the integro-differential form of the *radiative transfer*, or *radiative transport equation*, or simply the RTE [Chandrasekhar, 1960], and it incorporates the four possible types of interaction events that can occur within the medium (recall the illustration in Figure 4.2).

By integrating both sides of Equation 4.25 and using the rendering equation (Equa-

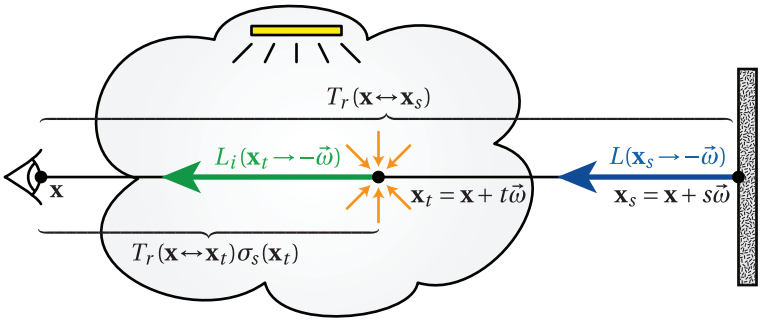

Figure 4.7: The radiance reaching the eye $L(\mathbf{x} \leftarrow \vec{\omega})$ is the sum of the reduced radiance from the nearest visible surface $L(\mathbf{x}_s \rightarrow \vec{\omega})$ and the accumulated in-scattered radiance $L_i(\mathbf{x}_t \rightarrow -\vec{\omega})$ along a ray.

tion 2.17) as a boundary condition, it is possible to obtain a purely integral equation for the radiance in the presence of participating media. The integral form of the RTE expresses the radiance as a sum of three terms:

$$\begin{aligned} L(\mathbf{x} \leftarrow \vec{\omega}) = & \underbrace{T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s)\, L(\mathbf{x}_s \rightarrow -\vec{\omega})}_{\text{reduced surface radiance}} + \\ & \underbrace{\int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\, \sigma_a(\mathbf{x})\, L_e(\mathbf{x} \rightarrow -\vec{\omega})\, dt}_{\text{accumulated emitted radiance}} + \\ & \underbrace{\int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t)\, \sigma_s(\mathbf{x}_t)\, L_i(\mathbf{x}_t \rightarrow -\vec{\omega})\, dt}_{\text{accumulated in-scattered radiance}}, \end{aligned} \tag{4.26}$$

where $L(\mathbf{x} \leftarrow \vec{\omega})$ describes the radiance that reaches $\mathbf{x}$ from direction $\vec{\omega}$, s is the depth of the medium from $\mathbf{x}$ in direction $\vec{\omega}$, $\mathbf{x}_t = \mathbf{x} + t\vec{\omega}$ with $t \in (0, s)$, and $\mathbf{x}_s = \mathbf{x} + s\vec{\omega}$ is a point on a surface past the medium (see Figure 4.7). The transmittance, T_r, models the reduction of radiance as it travels through the medium, as described in Section 4.2.1.

The first term in Equation 4.26 represents surface radiance entering at the backside of the medium, the second term incorporates the emission of radiance along the medium, and the last term is radiance scattered within the medium. The in-scattered radiance is given by Equation 4.13. In the remainder of this dissertation we will ignore the emission term since it is trivial to compute. With this simplification, the relationship between incident and excitant radiance can be expressed

as

$$L(\mathbf{x}\leftarrow\vec{\omega}) = \underbrace{T_r(\mathbf{x}\leftrightarrow\mathbf{x}_s)\,L(\mathbf{x}_s\rightarrow-\vec{\omega})}_{\text{reduced surface radiance}} + \underbrace{\int_0^s T_r(\mathbf{x}\leftrightarrow\mathbf{x}_t)\,\sigma_s(\mathbf{x}_t)\,L_i(\mathbf{x}_t\rightarrow-\vec{\omega})\,dt}_{\text{accumulated in-scattered radiance}}. \tag{4.27}$$

This equation replaces the relation in Equation 2.10 when media scattering and absorption are considered.

In computer graphics, this integral form of the radiative transfer equation is often referred to as the *volume rendering equation* because it is the equation that must be solved in order to render images with participating media. Unfortunately, the volume rendering equation is even more complex than the regular rendering equation. It describes radiance within participating media as a five-dimensional function over 3D positions in the volume and 2D directions over the sphere. The radiance at any point in the medium depends on the radiance of all other points in the medium as well as the radiance of all points on surfaces. In contrast, the regular rendering equation is only a four-dimensional subset of this expression (2D surface locations and 2D directions). This makes the volume rendering equation *extremely* costly to solve.

4.5 Methods for Solving the Volume Rendering Equation

Simulating light transport in participating media has been studied extensively in the literature. Many of the techniques used today were intially developed to study radiative heat transfer and neutron transport in the 1950s and 60s [Kourgnaoff, 1952; Chandrasekhar, 1960]. Mishra1 and Prasad [1998] provide an excellent survey of this early work. It was not until Blinn's pioneering paper in 1982 [Blinn, 1982] that volume effects were considered in computer graphics. Since that time a number of researchers have investigated ways to efficiently simulate volume scattering in the context of computer graphics. We briefly summarize the most significant work here and refer to the excellent survey by Cerezo et al. [2005] for a more comprehensive overview of solving the RTE for computer graphics.

4.5.1 Deterministic Methods

Some of the earliest research for participating media rendering extended discrete radiosity methods to handle isotropic volume scattering. Rushmeier and Torrance [1987] used the *zonal method* [Hottel and Sarofim, 1967] which discretizes the volume rendering equation and solves it over a collection of finite elements or *zones*. This early approach simulated volume-volume, surface-volume, as well as surface-surface interactions but were impractical for all but the simplest scenes due to the high storage costs and $O(n^7)$ computation complexity. As with surface-base radiosity techniques, some of these issues were later addressed by using clustering and hierarchical computation [Bhate, 1993; Sillion, 1995].

Deterministic methods were later augmented to handle anisotropic scattering. P_N methods use a series approximation to replace the integro-differential RTE with a set of partial differential equations, which are expanded using N spherical harmonics terms. This approach was introduce to graphics in 1984 by Kajiya and Herzen but had been used in radiative heat transfer and neutron transport fields much earlier [Chandrasekhar, 1960; Kourgnaoff, 1952] and was originally introduced as early as 1917 by Jeans. Bhate and Tokuta [1992] extended the zonal method with P_N approximations by using spherical harmonics instead of constant basis functions to simulate anisotropic scattering. The approach was even more computationally expensive and required the computation of $O(n^7 + N^2 n^6)$ form factors. An alternate approach is the *discrete ordinates* method, which was first proposed by Chandrasekhar [1960] for simulating stellar and atmospheric radiation. This method subdivides the radiosity leaving each zone into a collection of M discrete directional bins and reduces the complexity to $O(n^7 + Mn^6)$. Unfortunately, discrete ordinates suffer from the "ray effect" [Lathrop, 1968] since light is propagated in discrete directions instead of over the whole solid angle of the bins. Max [1994] reduced the ray effect by approximating the spread through the whole bin and was able to improve efficiency to $O(Mn^3 \log n + M^2 n^3)$ per iteration.

One major advantage of radiosity techniques is that the solution is computed in world-space. This allows the algorithms to exploit spatial coherence by computing the solution sparsely. Unfortunately, volumetric finite element methods share the drawbacks of surface-based radiosity

techniques. Scenes with complex geometry and volume densities cannot easily be handled since the computation of light transport it strictly coupled to the geometric representation of the scene.

4.5.2 Stochastic Methods

Some of the most popular techniques to date are based on stochastic sampling and Monte Carlo integration. Rushmeier [1988] extended the zonal method by using a Monte Carlo pass to account for one bounce of anisotropic scattering. Extensions of path tracing to simulate volumetric scattering in participating media have been suggested by Hanrahan and Krueger [1993] and Pattanaik and Mudur [1993]. Extensions to bidirectional path tracing and Metropolis light transport have also been proposed [Lafortune and Willems, 1996; Pauly et al., 2000]. Monte Carlo approaches are attractive because of their sound underlying theoretical framework and their generality. They are unbiased and guaranteed to converge to the exact solution. In addition, it is straightforward to include heterogeneous media, anisotropic phase functions, and scattering from surfaces. The downside of these approaches is that they suffer from noise that can only be overcome with a huge computational effort.

One strategy to solve this issue is to make simplifying assumptions about the participating media. For example, homogeneous media with a high scattering albedo can be modeled accurately using a diffusion approximation [Stam, 1995; Jensen et al., 2001b], which leads to very efficient rendering algorithms. Premoze et al. [2004], under the assumption that the medium is tenuous and strongly forward scattering, use a path integral formulation to derive efficient rendering algorithms. Sun et al. [2005] render single scattering in real time, but without shadowing effects.

In contrast, photon mapping [Jensen and Christensen, 1998] improves the efficiency of path-tracing without making additional assumptions about the properties of the medium being rendered. Similar to Monte Carlo methods, photon mapping handles isotropic, anisotropic, homogeneous, and heterogeneous media of arbitrary albedo. A disadvantage of photon mapping is that it introduces bias to the solution of the radiative transfer equation. In practice, however, this bias is preferable to the noisy solutions of pure Monte Carlo methods. We cover the photon

mapping method in more detail in Chapter 7.

This dissertation focuses on biased Monte Carlo solutions to the radiative transport equation. These techniques exploit the spatial coherence benefits associated with finite element methods while still retaining the generality of Monte Carlo based approaches.