

Skin Rendering: Reflectance and Integration

Morten S. Mikkelsen
Naughty Dog Inc., USA

May 28, 2009

Abstract

New and original approaches to model subsurface scattering, in computer graphics, have been introduced within the past decade. In this paper we are particularly interested in the application of these techniques in the context of achieving convincing looking human skin. The papers which we consider to be seminal, on this subject, base their analysis on several underlying theoretical principles which require an intimate knowledge of concepts used in optics. For this reason we provide, in our paper, a thorough walkthrough including all details necessary to comprehend the full analysis.

Ultimately, during the process of rendering, with subsurface scattering, an expensive integration step is performed. In this paper we discuss some of the proposals given by various authors.

Contents

1	Introduction	3
2	The Reflectance Profile	8
2.1	Radiance	9
2.2	The rendering equation	10
2.3	The two-term BSSRDF	12
2.4	The Diffusion Approximation	14
2.4.1	The Diffusion Equation	14
2.4.2	Approximate Boundary Condition	20
2.5	Approximation of the Diffuse BSSRDF	23
2.6	The Dipole Approximation	25
2.7	The Multipole	28
2.7.1	Reflectance of a Slab	29
2.7.2	Reflectance of Layers of Slabs	32
2.7.3	Combining Reflectance Profiles	34
2.8	Resulting Profiles	37
3	The Integration Step	42
3.1	Preliminaries	42
3.2	Two-pass integration	45
3.3	Mesh Unwrap-Based Distribution	48
3.4	Mesh Unwrap-Based Integration	49
3.5	Diffuse Color Maps	57
4	Conclusion	59

1 Introduction

The visual characteristics of a material are due to the interaction between light and the molecular structure and density of the medium. As an example metal distinguishes from a soft material because it does not exhibit any translucency due to *subsurface scattering*. This phenomenon occurs when light enters the medium and scatters around internally. The light will either be absorbed inside the medium or travel back out to the boundary of the medium and leave it while moving in some final outgoing direction. Technically, this light might scatter in the outer medium, let this be air, and return to the embedded medium. However, for simplicity it is most common in computer graphics to assume that light travels through air unobstructed and in straight lines. Given this approximation the problem, in the context of computer graphics, is reduced to determining the quantity of light which travels from a visible point, at the surface of a medium, and towards the observer. Thus we do not have to account for air as a separate medium. Nevertheless, even with this simplification one must still account for the complex path by which light travels within the medium. For the general case this involves a slow process in which the equation of transfer is solved numerically. For this reason only a few papers in graphics have taken this approach to achieve subsurface scattering. Instead, it is common, to proceed under the assumption that the reflectance model is determined by surface scattering. Any subsurface scattering is accounted for by a Lambertian component also known as a diffuse term. This term reflects light evenly in all outgoing directions.

If the surface of a medium is rugged and rough, at a micro scale level, then the reflected outgoing direction appears to be random which justifies the even distribution. Additionally, for a highly scattering medium light will follow, what appears to be a chaotic path inside the medium. Once it, finally, returns to the surface and leaves the medium all sense of the initial incident direction will be lost. This also agrees with the even distribution of the diffuse term and for this reason, and because of its efficiency, it is often used to account for both phenomena. Though this simplification is efficient in terms of performance, the approximation breaks down when applied to subsurface scattering in soft materials. This is the case because the model does not take into account that at the end of the traveled path the exit and entry locations, of the transmitted light, are generally not the same (see figure 1). To capture the true appearance of translucent materials light transferred, beneath the surface, between surface points must be taken into account. This

is true for many materials such as fruits, marble, ice, plants, skin, etc. In this paper we are particularly interested techniques for producing convincing looking human skin.

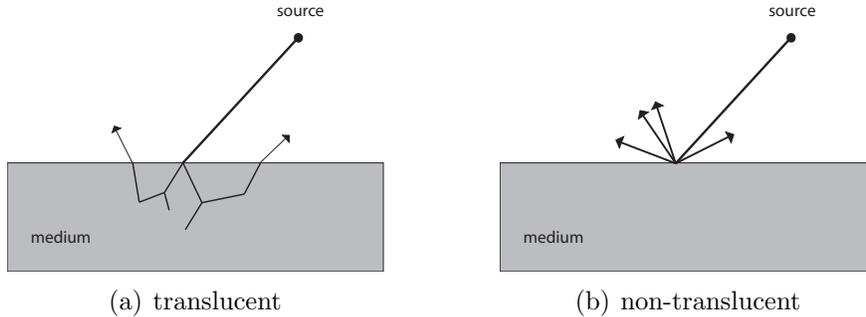


Figure 1: For a translucent material the point of exit, at the surface, is in general not the same as that of entry as shown in figure 1(a). For a non-translucent material they are considered to be approximately the same as shown in figure 1(b).

In 2001 Henrik Wann Jensen et al. [JMLH01] introduced a technique to computer graphics using an approximate model designed by Patterson et. al [PCW89] of the medical physics community. The approximation was designed to perform non-invasive measurements to be used during photodynamic therapy. Given a laser light incident on a surface, of human tissue, the method approximates the redistribution, due to internal multiple scattering, of light across this surface. Such a redistribution function is known as *reflectance* and describes the ratio between outgoing light at one surface point as a result of incoming light at another. In this case the other point is where the laser light is incident upon the surface. Jensen applies this model, in the context of computer graphics, by considering local irradiance as being injected into the medium by the laser light which is perpendicular to the surface. This implies an approximation that the incident direction, prior to entering the medium, has no effect on overall redistribution. For a highly scattering medium this is justified by the observation that after a few internal scattering events any sense of initial path is lost.

The motivation for using this reflectance profile is it allows us to determine outgoing radiance as an integral across the surface as opposed to numerically

solving the more general equation of transfer which is volumetric. Furthermore, the reflectance profile is equipped with an exponential fall-off which enables us to accelerate the integration process.

The reflectance profile is based on several simplifications and assumptions which will be covered in sections 2.4–2.6. However, the most critical simplification is that the derivation of the reflectance profile is based on a semi-infinite plane-parallel medium with constant material properties. This results in a single reflectance profile which is to be used at every shaded surface point. Integrating the reflectance profile over the entire surface plane gives us what is known as the *total diffuse reflectance*. Given an initial contribution of light injected into the medium the value of the total diffuse reflectance is the percentage, of this contribution, which will make it back up to the surface and leave the medium. Since the reflectance profile is wavelength dependent this means the total diffuse reflectance is also wavelength dependent. Furthermore, since it dictates the percentage of light reflected it relates directly to the concept of surface color used in computer graphics. However, since the material properties are constant this means we have one material color only which is not sufficient to capture the high frequency variation in color found in human skin. It is pointed out in [JMLH01] that volumetric variation in material properties would require a full participating media simulation. Instead, it is assumed that the shape of the reflectance profile remains constant and only varies across the surface by a scalar. This scalar is wavelength dependent and corresponds to storing the color in a texture mapped to the surface of the medium. The issue is discussed in further detail in section 3.5.

The claim is made in the more recent work by Craig Donner et al. [DJ05] that this remedy is insufficient. In their paper the shape of the resulting reflectance profile is constant as well, however, it is argued that the true shape of the reflectance profile, corresponding to human skin, requires a model which takes into account that the material properties change at different levels of depth below the surface. When this is not taken into account, using the method in [JMLH01], results become overly blurred or waxy. This is the report given in [DJ05] and to account for this they employ a similar, but more complex, approximate model based on multi-layered slabs. Each slab has constant material parameters and these are stacked on top of each other. Once again incident light is assumed to be injected, by a laser light, perpendicular to the configuration. This multi-layer model is borrowed from the optics community and is described in [CMZ97] and here in our paper

in section 2.7. It is common in medical physics to think of the human skin as composed of layers, corresponding to these slabs, and in [DJ05] these are referred to as the following: epidermis, upper dermis and bloody dermis. These define the skin from top to bottom where the initial two are slabs and the bloody dermis is considered semi-infinite. The constant material parameters of each of these three layers are given by a table in [DJ05] with a reference made to the medical physics book [Tuc00].

During the development of our paper we have implemented the methods, given in [JMLH01] and [DJ05], to produce their respective reflectance profiles. The results are shown in section 2.8 along with the reflectance profile obtained using open source software known as MCML. The program was written by Lihong Wang and Steven L. Jacques [WJ] and is also designed to produce reflectance profiles given a configuration of multi-layered media. Their software is based on an entirely different technique known as Monte Carlo simulation and we use their program to verify our implementation.

Once a reflectance profile has been established this is used to determine the outgoing radiance. This is done by an integration step performed across the surface. It was suggested by Jensen in 2002 [JB02] that this step could be accelerated by exploiting the rapid fall-off of the reflectance profile. This is done using an octree structure to allow for a more crude precision, during integration, at remote sample points. The method is explained in section 3.2.

Since then different authors have made attempts to solve this integral on the GPU. Examples are papers such as [DS03] and [SKP09] which both use a dual image-space representation. One as seen from the observer and the other as seen from the light source. Both suffer from disadvantages such as: the method assumes either a point or directional light source. Additionally, such a method suffers from sampling density mismatches between the two images which is a problem well known from shadow mapping and projective texturing in general. Finally, unlike Jensen’s octree algorithm the integration step of these methods is performed for every light source. A different approach to a GPU adaptation is suggested by Eugene d’Eon et al. in [dL07]. In this article the integration is performed by convolution in a 2D texture unwrap of the model. This is done by describing the reflectance profile using a weighted sum of Gaussians. Each Gaussian is a separable convolution which significantly accelerates the process. The details of this method are explained in section 3.4. The primary disadvantages to this method are: though the integration can be done as two full-screen passes it must be done per unwrap

times the amount of Gaussians used in the reflectance profile. In their paper [dL07] six Gaussians are used. Another problem is that the method only accounts for local subsurface scattering and not effects such as bright light transmitted through thin regions such as the ears of a human head. This is referred to in [dL07] as global subsurface scattering and is ultimately accounted for by resorting to a similar dual image-space representation. Thus the amount of necessary integration steps is now multiplied by the number of light sources in the scene.

2 The Reflectance Profile

In this section we will clarify and derive the underlying mathematical principles which lead to the reflectance profile used by Henrik Wann Jensen et al. [JMLH01] and that which is used in Craig Donner et al. [DJ05]. These are used to account for subsurface scattering during the rendering process.

Initially, we introduce in section 2.1 the basic terms associated with light. We follow up in section 2.2 with the *rendering equation*, used in computer graphics, to determine the quantity of light sent from a visible surface and in the direction of an observer. In order to do so a transfer function, which translates received light into reflected light, must be defined. In section 2.3 we discuss the separation of this transfer function into a term which accounts for subsurface scattering and one for surface scattering. To formulate the former Jensen turns to the diffusion approximation which describes density fluctuations. In this case density fluctuations of incoming light at points within a medium. The diffusion approximation is derived in section 2.4. Assuming we know the solution to the diffusion approximation, and sources are isotropic, then a simple formulation of the reflectance profile emerges which will be shown in section 2.5.

The formulation of the term, of the transfer function, which accounts for subsurface scattering is determined by observing the density fluctuations of light within a homogeneous semi-infinite plane-parallel medium. In this case the source is a laser beam which is perpendicular to the surface plane. To simplify the solution a common approximation, in optics, is applied which is known as the dipole approximation. This yields a nice analytical solution to the reflectance profile which is shown in section 2.6.

It is argued in [DJ05] that skin consists of multiple semi homogeneous layers and for this reason a switch is made to the multipole approximation which assumes the medium can be described by a stack of slabs such that each slab has finite thickness and is homogeneous. The multipole method is covered in section 2.7.

Finally, in section 2.8 we show the reflectance profiles we obtain given parameters defined in [JMLH01] and [DJ05] and compare these to a Monte Carlo simulation done using the software MCML.

2.1 Radiance

Light is the result of photons moving in free space and through sparse media. Fundamentally, in computer graphics, we need to determine how many photons collide with a visible surface for each unit of time. This value is known as *radiant flux*, Φ , and is measured in Watts. Henceforth we will, for brevity, refer to this simply as flux. A more discriminating way to measure the density of photons is *radiance*, $L(x, \vec{\omega})$. Radiance imposes restrictions by filtering out photons, at x , which are not within a differential frequency band centered at some given frequency level and also photons which are not within a differential solid angle around $\vec{\omega}$. It is measured in $Wm^{-2}sr^{-1}Hz^{-1}$ where sr is steradian which is the unit solid angle. The differential flux, from direction $\vec{\omega}$, flowing through an elementary area dA at x with the normal \vec{n} can at an infinitesimal level be determined by scaling the radiance by the cosine to the angle between $\vec{\omega}$ and \vec{n} which is scalar projection. From this we obtain the following equation which is also given by 7-1 in [Ish78]

$$\frac{d^2\Phi}{dAd\vec{\omega}} = (\vec{\omega} \bullet \vec{n})L(x, \vec{\omega}) \quad (1)$$

The cosine term diminishes the density of photons arriving within dA when the angle is large. An illustration of this is shown in figure 2.

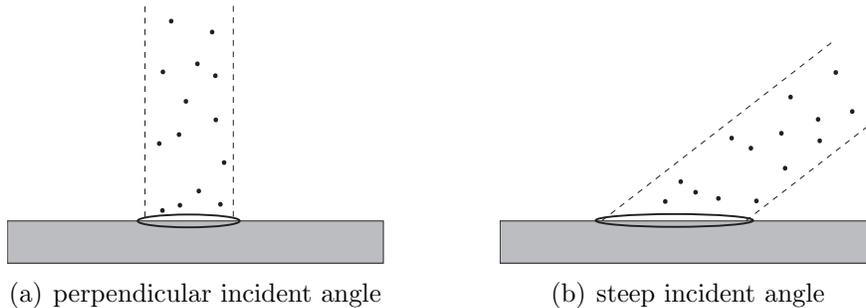


Figure 2: In figures 2(a) and 2(b) the same solid of incoming photons is shown. However, in 2(b) due to the incident angle the photons are projected onto a larger area which diminishes the density.

There exists an ambiguity in regards to radiance. It can either move along the path of $\vec{\omega}$ or it can be coming from $\vec{\omega}$. In this paper radiance received

from the direction of $\vec{\omega}$ is called *incoming radiance* $L(x, \vec{\omega})$ and radiance moving along $\vec{\omega}$ is *outgoing radiance* $L_o(x, \vec{\omega})$. Consequently, it follows that

$$L_o(x, -\vec{\omega}) = L(x, \vec{\omega}) \quad (2)$$

Similarly, the differential flux $\frac{d\Phi}{dA}$, per unit area, can be either incoming or outgoing relative to a surface orientation \vec{n} . These are known as irradiance, $I(x)$ and the radiosity $B(x)$ and are equivalent to integrating equation (1) over the hemisphere.

2.2 The rendering equation

The generally accepted reference model, in computer graphics, of global illumination is known as *the rendering equation* and was introduced in [Kaj86]. The rendering equation is an adaptation of the work in [SH81] on the phenomenon known as radiative heat transfer. It can be used to compute the outgoing radiance at a hard surface where the medium does not exhibit any subsurface scattering and it is given as the emitted radiance, L_e , plus an integral over all surfaces in the scene.

$$L_o(x \rightarrow x_r) = L_e(x \rightarrow x_r) + \int_S f_r(x_t \rightarrow x \rightarrow x_r) L_o(x_t \rightarrow x) V(x_t, x) G(x_t, x) dx_t$$

In this equation x_t represents transmitting locations in the scene and the integration is done with respect to this variable. The constant x is some light receiving surface location and the entire integral approximates how much of it is reflected from x in the direction of the surface point x_r . The function f_r is the *bidirectional reflectance distribution function* (BRDF), introduced by [NRH⁺77], which multiplied by the incident differential flux gives the corresponding differential outgoing radiance. The term $L_o(x_t \rightarrow x)$ represents the radiance sent, at x_t , towards x . The light is attenuated as it travels towards x and so to arrive at the incident flux at x the term is scaled by a geometric term

$$\vec{\omega}_i = \frac{x_t - x}{\|x_t - x\|}$$

$$G(x_t, x) = \frac{(\vec{\omega}_i \bullet \vec{n}) \cdot (-\vec{\omega}_i \bullet \vec{n}_t)}{\|x_t - x\|^2}$$

where \vec{n} and \vec{n}_t are the surface normals at x and x_t respectively. The denominator of the geometry term can be explained as in [Jen01] - i.e. imagine

photons resting on the surface of an expanding sphere. The density of photons will decrease proportionally to the area of the sphere. The two cosine factors in the numerator take the orientation of the transmitting and receiving surface locations into account as shown in figure 2. The geometry term should be recognized immediately by those familiar with the *form factor* used in radiosity based lighting methods. Finally, the term $V(x_t, x)$ is one when x_t and x are mutually visible and zero otherwise. This ensures that x only receives from the single closest transmitter x_t in every direction $\vec{\omega}_i$. The observation allows us to rewrite the integral, by substitution, where x_t is replaced by the first obtained hit when shooting from x along $\vec{\omega}_i$ and $dx_t = \frac{\|x_t - x\|^2}{(-\vec{\omega}_i \bullet \vec{n}_t)} d\vec{\omega}_i$. Thus, the integration is done over the hemisphere at x as follows

$$\begin{aligned} L_o(x \rightarrow x_r) &= L_o(x, \vec{\omega}_o) \\ &= L_e(x, \vec{\omega}_o) + \int_{\Omega_{2\pi}} f_r(x, \vec{\omega}_o, \vec{\omega}_i) L(x, \vec{\omega}_i) (\vec{\omega}_i \bullet \vec{n}) d\vec{\omega}_i \end{aligned} \quad (3)$$

However, to evaluate outgoing radiance, a significant limitation in this model is the assumption that light received at the position x is also reflected at x . Though this approximation is valid for many metals it fails to capture the soft appearance of translucent materials such as: skin, milk, etc. The ratio between outgoing radiance, at x_o , as a result of incoming differential flux, at x_i , is defined in [NRH⁺77] as the *bidirectional scattering surface reflectance distribution function* (BSSRDF). Given a known BSSRDF, $S(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o)$, the limitation is overcome by integrating the incoming radiance across the scene. This allows us to account for radiance received across the entire scene and subsequently transferred through the medium and exiting through a differential area at x_o in a direction within a differential solid angle around $\vec{\omega}_o$. If we omit the emission term, we arrive at the following double integral

$$L_o(x_o, \vec{\omega}_o) = \int_S \int_{\Omega_{2\pi}} S(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) L(x_i, \vec{\omega}_i) (\vec{\omega}_i \bullet \vec{n}_i) d\vec{\omega}_i dx_i \quad (4)$$

Like the BRDF, the BSSRDF S relates the differential outgoing radiance at x_o to the differential incident flux at x_i (see [NRH⁺77]).

$$dL_o(x_o, \vec{\omega}_o) = S(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) d\Phi_i(x_i, \vec{\omega}_i) \quad (5)$$

It follows from equation (1) that $L(x_i, \vec{\omega}_i) (\vec{\omega}_i \bullet \vec{n}_i)$ represents the flux received at x_i within a differential area, dA , from a direction within a differential angle

around $\vec{\omega}_i$. By using this observation in equation (4), and by substituting equation (5) into it, it follows that equation (4) is effectively summing up all contributions in the scene that lead to outgoing radiance at x_o in the direction $\vec{\omega}_o$.

2.3 The two-term BSSRDF

In practice, the task of determining the function of, S , is not a simple one. This issue is even a problem for the simpler BRDF form, f_r , and a common approximation in computer graphics is to split it into two terms, $f_r \simeq f_d + f_s$ where f_d represents diffuse reflection and f_s represents specular reflection. Diffuse reflection occurs when light bounces off a surface which at a micro-scale level has an uneven and rough terrain. This tends to make the reflected direction more random and spread the out-going radiance evenly across the hemisphere, $\Omega_{2\pi}$, which makes the observed result independent of the view-direction $\vec{\omega}_o$. Specular reflection represents the opposite phenomenon where light coming from a single direction is reflected at a locally flat surface into a single out-going direction and thus becomes view-dependent. Both diffuse and specular reflection represent ideal cases.

Initially, this line of thinking might seem like a dead-end in regards to the BSSRDF since S needs to take into account the entry location, x_i , the path along which the light has traveled inside the medium until it finally reaches its point of exit x_o . However, highly scattering soft materials actually agree very well with the approximation that the outgoing radiance at x_o is diffuse since light which is *transmitted* into such a medium scatters around almost chaotically and quickly begins to exhibit isotropic behavior within the medium. This observation is made in chapter 9 in [Ish78].

The portion of light which is not transmitted into the medium, at x_o , but is reflected at the surface might be controlled by a different term similar to the specular term in the approximate BRDF. Thus, we can approximate the BSSRDF by

$$S(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) \simeq S_d(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) + S_r(x_o, \vec{\omega}_i, \vec{\omega}_o) \quad (6)$$

where S_d is used to approximate light which has been transmitted into the medium at x_i , then scattered around multiple times and finally transmitted back out at x_o in the direction $\vec{\omega}_o$. The term S_r is used to determine the amount of light reflected, directly at the surface, in the direction $\vec{\omega}_o$. This approximation is also used in [DJ05].

This separation does not appear to simplify much since S_d still depends on all four parameters. However, its primary dependency on the incoming and outgoing directions $\vec{\omega}_i$ and $\vec{\omega}_o$ is mostly due to the different indices of refraction, when the light transmits in and out of the medium. When incoming radiance travels between media of different *indices of refraction*, μ_1 and μ_2 , a certain percentage is **reflected** at the boundary and the rest is **transmitted** into the other medium. In the remainder the percentage of reflected radiance function will be denoted $\rho_r(x, \vec{\omega})$ and the percentage of transmitted radiance $\rho_t(x, \vec{\omega})$. In both cases $\vec{\omega}$ is the direction of incoming radiance. The index of refraction roughly describes the density of the material. In more specific terms light traveling through a medium with index of refraction μ moves at $\frac{1}{\mu}$ times the speed of light in a vacuum. It is common in computer graphics to evaluate $\rho_r(x, \vec{\omega})$ using the Fresnel equations or an approximation by Schlick (see equations 2.29 and 2.30 in [Jen01]). The Fresnel equations are derived from the Maxwell equations based on ideal circumstances where both media are homogeneous, i.e., constant material properties and the *interface* between them is planar. In the case where air represents the outer medium and skin represents the embedded medium, the criteria of homogeneity is approximately met. Subsequently, the Fresnel equations are used, for skin, by Jensen in [JMLH01] to determine reflectance vs. transmittance. However, the surface of skin is by no means planar so it is decided in [DJ05] to replace the use of Fresnel with a diffuse transmission function. This is done by using a Torrance–Sparrow BRDF [TS67] to determine the reflectance at the surface.

$$\rho_r(x, \vec{\omega}) = \int_{\Omega_{2\pi}} f_r(x, \vec{\omega}', \vec{\omega})(\vec{\omega}' \bullet \vec{n}) d\vec{\omega}' \quad (7)$$

$$\rho_t(x, \vec{\omega}) = 1 - \rho_r(x, \vec{\omega}) \quad (8)$$

The second equation assumes that all light which is not reflected is transmitted. It should also be noted that due to energy conservation $\rho_r(x, \vec{\omega})$ cannot exceed 1.0.

A possible formulation for S_d can be determined by studying density fluctuations of incoming light inside a medium. This is implicitly described by the diffusion approximation which we will cover in the next section.

2.4 The Diffusion Approximation

In this section we cover the *diffusion approximation* which consists of two parts. That is a differential equation known as the *diffusion equation* and a supplement known as the *boundary condition*. The solution to the diffusion equation is the *fluence*, $\phi(x)$, which is the total incoming radiance at x . Since this is an isotropic solution the boundary condition is used as an additional constraint to account for the interface between adjacent media. The diffusion approximation is based on the observation that the light distribution in highly scattering media tends to become isotropic (see chap. 9 in [Ish78]). This is also implied by the solution, $\phi(x)$, which is isotropic. The diffusion equation is derived in section 2.4.1 and the boundary condition is covered in section 2.4.2.

2.4.1 The Diffusion Equation

To derive the diffusion equation we must start at the *equation of transfer* which is a volumetric differential equation of radiance. It belongs to *transport theory* which was initiated by Schuster in 1903. A thorough description of transport theory can be found in [Cha50].

$$(\vec{\omega} \cdot \nabla)L(x, \vec{\omega}) = -\sigma_t(x)L(x, \vec{\omega}) + \sigma_s(x) \int_{\Omega_{4\pi}} p(x, \vec{\omega}', \vec{\omega})L(x, \vec{\omega}')d\vec{\omega}' \quad (9)$$

The equation describes the rate of change in radiance. When a photon travels it may either move unaffected or interact with the medium in which it is traveling. If it does interact, it will either be absorbed or be scattered in a new direction. The quantity of photons scattered per unit traveled is given by $\sigma_s(x)$ and the quantity absorbed is given by $\sigma_a(x)$. The extinction coefficient is given as $\sigma_t(x) = \sigma_s(x) + \sigma_a(x)$. Thus, the rate of loss in radiance flowing in the direction $-\vec{\omega}$ at x is accounted for by the first term on the right side of equation (9). The second term accounts for gain in radiance due to in-scattering. This is done by integrating the radiance over all possible directions. During integration the radiance is scaled by $\sigma_s(x)$ and the phase function $p(x, \vec{\omega}', \vec{\omega})$. The former will give us the portion of radiance from $\vec{\omega}'$ scattered at x and the latter will give us the portion of this contribution which scatters specifically into the direction $\vec{\omega}$. The phase function is a probability density function and obeys $\int_{\Omega_{4\pi}} p(x, \vec{\omega}', \vec{\omega})d\vec{\omega}' = 1$. Though the scattering and absorption depend on the point x , they will, for brevity, be denoted σ_s

and σ_a . However, for a homogeneous medium σ_s and σ_a are constant for any point inside the medium.

When studying the transport of radiance inside some medium it is common to divide the radiance into two parts (see section 7.4 in [Ish78]). The *reduced radiance* $L_{ri}(x, \vec{\omega})$ is radiance which has entered the medium but has not yet scattered or been absorbed. The other part, *diffuse radiance* $L_d(x, \vec{\omega})$, is radiance which was created inside the medium due to emission or in-scattering of reduced/diffuse radiance. An illustration is shown in figure 3. Given this separation reduced radiance has no gain due to in-scattering

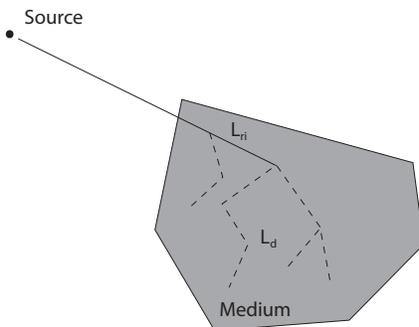


Figure 3: The source emits light which enters the sparse medium as shown here. The portion of this contribution which has survived absorption and has scattered once or more is referred to as diffuse radiance L_d . This portion of the radiance is illustrated by dashed lines. Prior to this first scattering event it is referred to as reduced radiance L_{ri} .

as opposed to diffuse radiance which has gain due to in-scattering of both diffuse and reduced radiance. We can think of in-scattered reduced radiance as emission $\varepsilon(x, \vec{\omega})$ in the transport equation of diffuse radiance.

$$\begin{aligned} \varepsilon(x, \vec{\omega}) &= \sigma_s \int_{\Omega_{4\pi}} p(x, \vec{\omega}', \vec{\omega}) L_{ri}(x, \vec{\omega}') d\vec{\omega}' & (10) \\ (\vec{\omega} \bullet \nabla) L_{ri}(x, \vec{\omega}) &= -\sigma_t L_{ri}(x, \vec{\omega}) \\ L(x, \vec{\omega}) &= L_d(x, \vec{\omega}) + L_{ri}(x, \vec{\omega}) \end{aligned}$$

This is because reduced radiance is converted into diffuse radiance only once inside the medium. From then on it remains diffuse radiance until it is either

absorbed or leaves the medium. The equation of transfer of diffuse radiance is now given by the following equation

$$(\vec{\omega} \bullet \nabla) L_d(x, \vec{\omega}) = -\sigma_t L_d(x, \vec{\omega}) + \sigma_s \int_{\Omega_{4\pi}} p(x, \vec{\omega}', \vec{\omega}) L_d(x, \vec{\omega}') d\vec{\omega}' + \varepsilon(x, \vec{\omega}) \quad (11)$$

In the following we will show how a specific two-term approximation, of $L_d(x, \vec{\omega})$, in equation (11) will allow us to reduce the equation into a simplified version known as the diffusion equation. The two terms used in the approximation are the fluence, which was conceptually introduced at the beginning of section 2.4, and the *vector irradiance*. These are given by the following equations

$$\begin{aligned} \phi(x) &= \int_{\Omega_{4\pi}} L_d(x, \vec{\omega}) d\vec{\omega} \\ \vec{E}(x) &= \int_{\Omega_{4\pi}} L_d(x, \vec{\omega}) \vec{\omega} d\vec{\omega} \end{aligned}$$

Since the fluence, $\phi(x)$, is the total incoming radiance at x this also implies that when divided by 4π it is equal to the average incoming radiance at x . The vector irradiance $\vec{E}(x)$ is the expected direction for incoming radiance. If we consider the normalized vector irradiance, $\vec{v}_e = \frac{\vec{E}(x)}{\|\vec{E}(x)\|}$, then the magnitude of $\vec{E}(x)$ is the irradiance from \vec{v}_e minus the irradiance from the opposite direction $-\vec{v}_e$. This is given by the following equation

$$\begin{aligned} \|\vec{E}(x)\| &= \vec{v}_e \bullet \vec{E}(x) \\ &= \int_{\Omega_{4\pi}} L_d(x, \vec{\omega}) \vec{v}_e \bullet \vec{\omega} d\vec{\omega} \end{aligned}$$

and by equation (1). We will now proceed by showing how a projection of the diffuse radiance onto the first two frequency bands of the *Real Spherical Harmonics* basis functions results in the two-term approximation involving $\phi(x)$ and $\vec{E}(x)$.

The diffuse radiance, at some point x , is a spherical function $L_d : S^2 \rightarrow \mathbb{R}$. Though radiance may, in reality, not be piecewise continuous we operate under the assumption that it is both bounded and piecewise continuous. This makes L_d Riemann integrable and so we can project it onto an orthonormal basis such as the Real Spherical Harmonics functions. These are organized in bands of increasing frequency such that each band, indexed by

$l \in \{0, 1, 2, \dots\}$, is represented by $2l + 1$ basis functions which are indexed by $m \in \mathbb{Z}$ such that $-l \leq m \leq l$. Subsequently, the basis functions $Y_l^m(\vec{\omega})$ are doubly indexed and projection, of a function $f : S^2 \rightarrow \mathbb{R}$, onto the first $L \in \mathbb{N}$ frequency bands is defined as the *partial sums* of the expanded spherical harmonics series

$$s_L(f)(\vec{\omega}') = \sum_{l=0}^{L-1} \sum_{m=-l}^l \left(\int_{\Omega_{4\pi}} f(\vec{\omega}) \cdot Y_l^m(\vec{\omega}) d\vec{\omega} \right) Y_l^m(\vec{\omega}') \quad (12)$$

A thorough introduction to Spherical Harmonics is beyond the scope of this paper. For additional information, the reader is referred to [Mac48]. Next, let the vector $\vec{\omega}$ be given by its components such that $\vec{\omega} = (\omega_1, \omega_2, \omega_3)$. Thus, the Cartesian version of the first two bands of basis functions is given as

$Y_l^m(\vec{\omega})$	$m = -1$	$m = 0$	$m = 1$
$l = 0$	$\frac{1}{2} \sqrt{\frac{1}{\pi}}$		
$l = 1$	$\frac{1}{2} \sqrt{\frac{3}{\pi}} \omega_2$	$\frac{1}{2} \sqrt{\frac{3}{\pi}} \omega_3$	$\frac{1}{2} \sqrt{\frac{3}{\pi}} \omega_1$

where $\vec{\omega}$ is a unit vector. Thus the two-term expansion obtained by the first two bands, for a low frequency approximation, is given by insertion of L_d into equation (12) with $L = 2$ which yields

$$\begin{aligned} L_d(\vec{\omega}') &\simeq s_2(L_d)(\vec{\omega}') \\ &= \frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \vec{E}(x) \bullet \vec{\omega}' \end{aligned} \quad (13)$$

We now proceed with our derivation of the diffusion equation. Similar to the definition of $\phi(x)$ and $\vec{E}(x)$ let the fluence and the vector irradiance of ε (see eq. (10)) be given as

$$\begin{aligned} Q_0(x) &= \int_{\Omega_{4\pi}} \varepsilon(x, \vec{\omega}) d\vec{\omega} \\ \vec{Q}_1(x) &= \int_{\Omega_{4\pi}} \varepsilon(x, \vec{\omega}) \vec{\omega} d\vec{\omega} \end{aligned}$$

When we integrate the left side of equation (11) we get the divergence of the vector irradiance

$$\begin{aligned} \int_{\Omega_{4\pi}} (\vec{\omega} \bullet \nabla) L_d(x, \vec{\omega}) d\vec{\omega} &= \nabla \bullet \int_{\Omega_{4\pi}} L_d(x, \vec{\omega}) \vec{\omega} d\vec{\omega} \\ &= \nabla \bullet \vec{E}(x) \\ &= \text{div}[\vec{E}(x)] \end{aligned}$$

The divergence of a vector field is the extent to which more flow is exiting or entering an infinitesimal volume at a point x . By integrating both sides of equation (11) we get the following formula for the divergence

$$\nabla \bullet \vec{E}(x) = -\sigma_a \phi(x) + Q_0(x) \quad (14)$$

As we see loss is determined by the absorption and gain by any external sources injecting reduced radiance into the medium.

In the following the approximation $s_2(L_d) \simeq L_d$ will be used in equation (11). As promised the process will ultimately lead to the diffusion equation. Initially, let the substitution be applied to the integral term of equation (11) and also assume the phase function p is given as a function of cosine to the angle between the in-scattered and out-scattered directions $\vec{\omega}'$ and $\vec{\omega}$. This is known as the phase angle.

$$\begin{aligned} \int_{\Omega_{4\pi}} p(\vec{\omega}' \bullet \vec{\omega}) s_2(L_d)(\vec{\omega}') d\vec{\omega}' &= \int_{\Omega_{4\pi}} p(\vec{\omega}' \bullet \vec{\omega}) \left(\frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \vec{E}(x) \bullet \vec{\omega}' \right) d\vec{\omega}' \\ &= \frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \int_{\Omega_{4\pi}} p(\vec{\omega}' \bullet \vec{\omega}) \vec{E}(x) \bullet \vec{\omega}' d\vec{\omega}' \\ &= \frac{\phi(x) + 3\vec{E}(x) \bullet \vec{\omega} \int_{\Omega_{4\pi}} p(\vec{\omega}' \bullet \vec{\omega}) \vec{\omega}' \bullet \vec{\omega} d\vec{\omega}'}{4\pi} \end{aligned} \quad (15)$$

The last step is the result of rotating the integral such that the vector $\vec{\omega}$ points down. The vectors $\vec{E}(x)$ and $\vec{\omega}'$ can thus be expressed in spherical coordinates (θ_1, φ_1) and (θ_2, φ_2) , given in this frame, where the domain is $[0; 2\pi[\times [0; \pi]$. Let ϕ be the angle between $\vec{E}(x)$ and $\vec{\omega}'$. Cosine to the angle is expressed in spherical coordinates as

$$\cos \phi = \cos \varphi_1 \cos \varphi_2 + \sin \varphi_1 \sin \varphi_2 \cos(\theta_1 - \theta_2)$$

During integration the second term is terminated since $\int_0^{2\pi} \cos \theta d\theta = 0$. This completes the last step in equation (15) since $\vec{E}(x) \bullet \vec{\omega} = \|\vec{E}(x)\| \cdot \cos \varphi_1$ and $\vec{\omega}' \bullet \vec{\omega} = \cos \varphi_2$. In addition to this we used that $\vec{\omega}$ and $\vec{E}(x)$ remain constant during integration. Next by defining the mean cosine as

$$g = \int_{\Omega_{4\pi}} p(\vec{\omega}' \bullet \vec{\omega}) (\vec{\omega}' \bullet \vec{\omega}) d\vec{\omega}'$$

we can write equation (15) as a simpler expression

$$\int_{\Omega_{4\pi}} p(\vec{\omega}' \bullet \vec{\omega}) s_2(L_d)(\vec{\omega}') d\vec{\omega}' = \frac{\phi(x) + 3g\vec{E}(x) \bullet \vec{\omega}}{4\pi}$$

And by defining the reduced scattering and extinction coefficients as

$$\begin{aligned}\sigma'_s &= \sigma_s \cdot (1 - g) \\ \sigma'_t &= \sigma'_s + \sigma_a\end{aligned}$$

we are finally ready to substitute $s_2(L_d)$ into both sides of equation (11).

$$\begin{aligned}(\vec{\omega} \bullet \nabla) \left(\frac{\phi(x) + 3\vec{E}(x) \bullet \vec{\omega}}{4\pi} \right) &= -\sigma_t \left(\frac{\phi(x) + 3\vec{E}(x) \bullet \vec{\omega}}{4\pi} \right) + \sigma_s \left(\frac{\phi(x) + 3g\vec{E}(x) \bullet \vec{\omega}}{4\pi} \right) + \varepsilon(x, \vec{\omega}) \\ &= -\frac{\sigma_a}{4\pi} \left(\phi(x) + 3\vec{E}(x) \bullet \vec{\omega} \right) - \sigma_s \frac{3}{4\pi} (1 - g) \vec{E}(x) \bullet \vec{\omega} + \varepsilon(x, \vec{\omega}) \\ &= -\frac{\sigma_a}{4\pi} \phi(x) - \sigma'_t \frac{3}{4\pi} \vec{E}(x) \bullet \vec{\omega} + \varepsilon(x, \vec{\omega}) \quad (16)\end{aligned}$$

The equation (16) can be further simplified. For any vector \vec{V} the following identities hold,

$$\begin{aligned}\int_{\Omega_{4\pi}} \vec{\omega}(\vec{\omega} \bullet \vec{V}) d\vec{\omega} &= \frac{4\pi}{3} \vec{V} \\ \int_{\Omega_{4\pi}} \vec{\omega}(\vec{\omega} \bullet \nabla(\vec{V} \bullet \vec{\omega})) d\vec{\omega} &= \vec{0}\end{aligned}$$

These are given in section 9 in [Ish78]. Scaling equation (16) by $\vec{\omega}$ and integrating over the sphere yields

$$\frac{\nabla\phi(x)}{3} = -\sigma'_t \vec{E}(x) + \int_{\Omega_{4\pi}} \varepsilon(x, \vec{\omega}) \vec{\omega} d\vec{\omega}$$

Let the diffusion constant be given as $D = \frac{1}{3\sigma'_t}$. Next we isolate $\vec{E}(x)$ and replace the integrated vector emission by the equivalent $Q_1(x)$.

$$\begin{aligned}\vec{E}(x) &= -\frac{\nabla\phi(x)}{3\sigma'_t} + \frac{1}{\sigma'_t} Q_1(x) \\ &= -D\nabla\phi(x) + 3DQ_1(x) \quad (17)\end{aligned}$$

Note that this expression (17) for the vector irradiance was derived based on the assumptions that $L_d = s_2(L_d)$ and that the phase function is a function of the phase angle. Now we substitute this result into the divergence, eq. (14), of $\vec{E}(x)$ which completes the diffusion equation

$$D\nabla^2\phi(x) = \sigma_a\phi(x) - Q_0(x) + 3D\nabla \cdot \vec{Q}_1(x) \quad (18)$$

The diffusion equation is a partial differential equation which describes density fluctuations. In this case fluctuations occurring in the average incoming radiance inside a medium. In other words the fluence $\phi(x)$, to be valid, should obey the diffusion equation.

2.4.2 Approximate Boundary Condition

Generally, it is not easy to determine the fluence $\phi(x)$. Even for the case of a homogeneous, but finite, medium there typically is no explicit solution. For a candidate $\phi(x)$ to be valid it must be a solution to the diffusion equation (18). For a trivial case such as an infinite homogeneous medium with an isotropic point light source there exists such a solution. Such a light source emits reduced radiance by the following equation

$$\varepsilon_{x_0}(x) = \frac{P_0}{4\pi} \cdot \delta(\|x - x_0\|)m^{-3}sr^{-1}$$

also given as equation 7-35 in [Ish78] where P_0 is the power. The Dirac δ function ensures emission only occurs at the position x_0 and obviously $\varepsilon_{x_0}(x)$ is defined to be independent of $\vec{\omega}$ which gives $Q_1(x) = 0$. By rearranging terms and dividing on both sides by D of the diffusion equation we get the following simplification.

$$\nabla^2\phi(x) - \sigma_{tr}^2\phi(x) = -3\sigma_t'Q_0(x)$$

where $\sigma_{tr} = \sqrt{3\sigma_a\sigma_t'}$ is the *effective transport coefficient*. Next, inserting $\varepsilon_{x_0}(x)$ into $Q_0(x)$ and substituting the result into this simplified diffusion equation gives

$$\nabla^2\phi(x) - \sigma_{tr}^2\phi(x) = -\frac{P_0}{D} \cdot \delta(\|x - x_0\|)$$

The final solution to such a differential equation is

$$\phi(r) = \frac{P_0}{D} \frac{e^{-\sigma_{tr}r}}{4\pi r} \quad (19)$$

where $r = \|x - x_0\|$. This is also the solution given by equation 9-53 in [Ish78] and by Jensen in [JMLH01].

For the case of a finite homogeneous medium in empty space it is more complicated. Diffuse radiance is the result of at least one scattering event within the medium. Diffuse radiance may leave the medium but since there is no scattering in empty space it does not return to it. In other words for a boundary point x , with the external and internal hemispheres $\Omega_{2\pi+}$ and $\Omega_{2\pi-}$, we have $L_d(x, \vec{\omega}) = 0$ when $\vec{\omega} \in \Omega_{2\pi+}$. This property is directionally dependent which implies the need to use a constraint. This is a supplement to the requirement that $\phi(x)$, which is isotropic, must be a solution to the diffusion equation (18). Such a constraint is known as a *boundary condition*. Strictly speaking the condition is only true for convex media. For additional details the reader is referred to [Ish78].

The criteria of the boundary condition states that diffuse radiance must come from within the medium. However, the diffusion equation is derived from a low-order spherical harmonics expansion of the diffuse radiance $L_d(x, \vec{\omega})$ given by equation (13) which has a simple angular distribution. Furthermore, the diffusion equation itself has an isotropic solution which does not take the boundary case into account. Because of the approximate representation the boundary condition cannot be satisfied in the exact form. Instead one must use an approximate boundary condition. The approximate boundary condition used in the diffusion approximation (see [Ish78]) is that the total **diffuse** flux received from outside the boundary must be zero.

$$\int_{\Omega_{2\pi}} L_d(x, \vec{\omega})(\vec{n} \bullet \vec{\omega})d\vec{\omega} = 0$$

However, if the finite medium is embedded in another medium (with another index of refraction) as opposed to just empty space. Then some of the outgoing diffuse radiance will bounce off the interface between the medium and the outside and back into the medium itself. Thus, instead of receiving no diffuse radiance at the surface from directions $\vec{\omega}$ in the upper hemisphere we receive whatever reflects back into the medium at the interface. This gives us the following equation.

$$\int_{\Omega_{2\pi+}} L_d(x, \vec{\omega})(\vec{n} \bullet \vec{\omega})d\vec{\omega} = F_{dr} \int_{\Omega_{2\pi-}} L_d(x, \vec{\omega})(-\vec{n} \bullet \vec{\omega})d\vec{\omega} \quad (20)$$

where F_{dr} is the diffuse reflectance.

$$F_{dr} = \frac{1}{\pi} \int_{\Omega_{2\pi}} F_r(\mu, \vec{n} \bullet \vec{\omega}') (\vec{n} \bullet \vec{\omega}') d\vec{\omega}' \quad (21)$$

and μ is the ratio of indices of refraction between the embedded medium and the outer medium. The above integral (21) is given, in [JMLH01], without the division by π . This is an error and will yield diffuse reflectance values in the range $[0; \pi]$. However, an approximate equation, given in [JMLH01], is used instead which does correspond to the correctly normalized equation (21).

The boundary condition can be expressed in terms of the fluence alone under the assumption that the reduced radiance $\varepsilon(x, \vec{\omega})$, at every x , is isotropic, i.e., the term is independent of $\vec{\omega}$. In this case $Q_1(x) = 0$ and equation (17) reduces to

$$\vec{E}(x) \simeq -D\nabla\phi(x) \quad (22)$$

We proceed by inserting the already assumed approximate form (13) into the boundary condition (20). If we do this on the left side we get

$$\begin{aligned} \int_{\Omega_{2\pi+}} \left(\frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \vec{E}(x) \bullet \vec{\omega} \right) (\vec{n} \bullet \vec{\omega}) d\vec{\omega} &= \frac{\phi(x)}{4} + \frac{3}{4\pi} \int_{\Omega_{2\pi+}} (\vec{n} \bullet \vec{\omega}) (\vec{E}(x) \bullet \vec{\omega}) d\vec{\omega} \\ &= \frac{\phi(x)}{4} + \frac{3}{4\pi} \vec{n} \bullet \vec{E}(x) \int_{\Omega_{2\pi+}} (\vec{n} \bullet \vec{\omega})^2 d\vec{\omega} \\ &= \frac{\phi(x)}{4} + \frac{3}{4\pi} \vec{n} \bullet \vec{E}(x) \frac{2\pi}{3} \\ &= \frac{\phi(x)}{4} + \frac{2}{4} \vec{n} \bullet \vec{E}(x) \end{aligned}$$

In the second step the same trick was used as in equation (15). Next we insert the dot product between approximation (22) and \vec{n} which gives us a new version, of equation (20), in terms of fluence

$$\phi(x) - 2D(\vec{n} \bullet \nabla\phi)(x) = F_{dr} \cdot (\phi(x) - 2D(-\vec{n} \bullet \nabla\phi)(x))$$

By rearranging terms we arrive at the final form

$$\phi(x) - 2DA(\vec{n} \bullet \nabla\phi)(x) = 0 \quad (23)$$

where $A = \frac{1+F_{dr}}{1-F_{dr}}$. This is the final approximate boundary condition for points x at the surface of the embedded finite medium. The equation is also given by Jensen on page 3 in [JMLH01].

2.5 Approximation of the Diffuse BSSRDF

In this section it will be shown how we can evaluate the diffuse BSSRDF, S_d , when all sources are isotropic. Thus approximation (22) is accurate and when the dot product between this and the inward normal $-\vec{n}$ is applied we get

$$\begin{aligned}
 D(\vec{n} \bullet \nabla \phi)(x) &\simeq -\vec{n} \bullet \vec{E}(x) \\
 &= \int_{\Omega_{4\pi}} L_d(x, \vec{\omega})(-\vec{n} \bullet \vec{\omega})d\vec{\omega} \\
 &= \int_{\Omega_{2\pi-}} L_d(x, \vec{\omega})(-\vec{n} \bullet \vec{\omega})d\vec{\omega} \tag{24}
 \end{aligned}$$

where x is at the surface of the medium. The switch in equation (24) from integrating over the entire unit sphere to the inner hemisphere is a result of the boundary condition which states that diffuse radiance is internal.

If we consider $L_o(x, \vec{\omega})$ the outgoing diffuse radiance then given equation (2) it follows that equation (24) is equal to $\int_{\Omega_{2\pi+}} L_o(x, \vec{\omega})(\vec{n} \bullet \vec{\omega})d\vec{\omega}$ which, relative to the normal \vec{n} , is the outgoing radiant flux area density $B(x)$ (see equation (1)).

$$D(\vec{n} \bullet \nabla \phi)(x) \simeq B(x) \tag{25}$$

Next imagine flux being injected into the medium from some differential area at x_i . Furthermore, assume we know the corresponding fluence $\phi(x)$ such that it approximates, specifically, how much of the flux emitted, from this area, at x_i is received at x as part, $dB(x_o)$, of the total outgoing radiant flux area density. We can now evaluate the reflectance, at x_o , as the following ratio

$$\begin{aligned}
 R(x_o, x_i) &= \frac{dB(x_o)}{d\Phi(x_i)} \\
 &\simeq \frac{D(\vec{n} \bullet \nabla \phi)(x_o)}{d\Phi(x_i)} \tag{26}
 \end{aligned}$$

This equation is given by H. Jensen, on page 3, in [JMLH01]. The ratio approximates the radiant exitance at x_o that occurs as a result of irradiance at x_i .

To approximate the diffuse BSSRDF, S_d , we need to approximate the outgoing radiance as opposed to radiant exitance at x_o . Furthermore, we

must also take into account that light generated outside the medium must make it through the interface between the medium and the outside. This is controlled by the transmittance $\rho_t(x_i, \vec{\omega}_i)$. Finally, the resulting radiant exitance at x_o must transmit, $\rho_t(x_o, \vec{\omega}_o)$, back out through the interface towards the eye of the camera. In section 2.3 the diffuse BSSRDF, S_d , is defined under the assumption that radiance moving inside the medium exhibits isotropic behavior. This assumption is based on the observation that in highly scattering media the light tends to distribute evenly in all directions. From this, and equation (1), it follows that the correspondence between diffuse radiosity and outgoing diffuse radiance is given by

$$\begin{aligned}
B(x_o) &\simeq \int_{\Omega_{2\pi+}} L_o(x_o)(\vec{n} \bullet \vec{\omega})d\vec{\omega} \\
&= L_o(x_o) \int_{\Omega_{2\pi+}} \vec{n} \bullet \vec{\omega}d\vec{\omega} \\
&= \pi \cdot L_o(x_o)
\end{aligned} \tag{27}$$

Now finally given equations (5) and (26) we arrive at the following diffuse BSSRDF

$$S_d(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) = \frac{1}{\pi} \rho_t(x_o, \vec{\omega}_o) R(x_o, x_i) \rho_t(x_i, \vec{\omega}_i) \tag{28}$$

This equation is also given by H. Jensen, on page 3, in [JMLH01]. To summarize equation (28) the procedure is as follows. Initially, the incoming radiance from some direction $\vec{\omega}_i$ must transmit through the interface. Once it has made it through, $\rho_t(x_i, \vec{\omega}_i)$, there is an implicit assumption that sense of direction is lost immediately given that $R(x_o, x_i)$ ignores both $\vec{\omega}_i$ and $\vec{\omega}_o$. In other words, for simplicity, light injected into the medium at x_i from any direction $\vec{\omega}_i$ is, past the point of transmittance, entirely diffuse. In reality it seems unlikely that we can completely dismiss direction of the incoming radiance past the point of transmittance but to simplify the BSSRDF it is a reasonable approximation made by Jensen. Finally, once the radiance makes it to x_o it has to transmit back out through the interface $\rho_t(x_o, \vec{\omega}_o)$ in the right direction. An additional assumption made, in the beginning of this section, to create approximation (26) was that any existing emitters are isotropic. The only missing element to finalize S_d is determining the fluence $\phi(x)$ resulting from diffuse radiance injected at some x_i . This is required to complete equation (26) and once we have this we have an explicit function to approximate the BSSRDF.

2.6 The Dipole Approximation

Even for a homogeneous medium the case of an arbitrary finite topology typically does not have an analytical solution $\phi(x)$ that obeys both the diffusion equation (18) and the boundary equation (23). For the case of subsurface reflection it is pointed out in [JMLH01] that this can be modeled as a semi-infinite plane-parallel medium. In other words a plane separating \mathbb{R}^3 such that all points on one side belong to one medium and all points on the other side belong to another medium. The separating plane is the interface between the two media. Both media are homogeneous but do not have the same material properties. Even under such simplified circumstances the solution for a point light, given by equation (19), does not obey the boundary condition (23). Furthermore, it is more practical to define the diffuse BSSRDF, S_d , in a way which is independent of the actual point lights and their specific locations in the scene. This is possible due to the approximation that the transmitted radiance, at x_i , quickly becomes isotropic. This was mentioned in section 2.5. Since we can ignore the original incoming direction $\vec{\omega}_i$, past transmittance, we can consider the total irradiance received at x_i as having been injected by a simple vertical cylindrical beam over x_i . Subsequently, we need to determine the resulting fluence $\phi(x)$ given the imaginary source beam over x_i . This will allow us to determine approximation (26) and thus complete the approximation of the diffuse BSSRDF given by equation (28).

Ironically, even this simplification is not adequate to derive a simple result for the fluence. Jensen points out that the result involves an infinite sum of Bessel functions [JMLH01]. Furthermore, he seeks a simple solution that does not involve infinite sums or numerical solutions. For this reason Jensen turns to the work of Eason [EVNT78] and Farrell et al. [FPW92]. Their method approximates the volumetric source distribution, resulting from the cylindrical beam, using two point lights. These two point lights are positioned like a dipole. The first point light emits positive radiance, inside the medium of interest, below the plane. It is referred to as the real light source and is placed at a distance $z_r = \frac{1}{\sigma_t}$ below x_i . This distance is known as *one mean free path* and represents the average distance a particle will travel before its first interaction inside the medium. The reason for this is that σ_t is the sum of absorbed and scattered photons per unit traveled. Subsequently, the reciprocal is the average distance traveled before a photon gets absorbed or scattered. Thus, in incremental terms, the first time photons injected by the laser beam, interact with the medium is approximately at

this distance underneath the interface. By replacing the laser beam by the real light source, which is isotropic, there is an additional assumption that the first bounce, at the specified distance underneath the interface, is evenly distributed. Effectively, the real source is emulating the redistribution of flux, passed by the laser beam, past the first bounce inside the medium. The power P_0 of the real source is inherited from the laser beam but since a certain portion of the flux will be absorbed while traveling one mean free path below x_i the power of the real source is adjusted by scaling by the reduced albedo $\alpha' = \frac{\sigma_s}{\sigma_t}$. In other words, any flux that does not get absorbed is scattered in a random direction. The second light source emits negative radiance and is referred to as the virtual light source. It is placed at a slightly greater distance $z_v = z_r + 4AD$ above the plane at x_i . The reason the second light source is needed is because of the boundary condition (23). The dipole model does not satisfy it exactly but it does satisfy a common approximation of it known as *the extrapolated boundary condition* (see figure 4). This approxi-

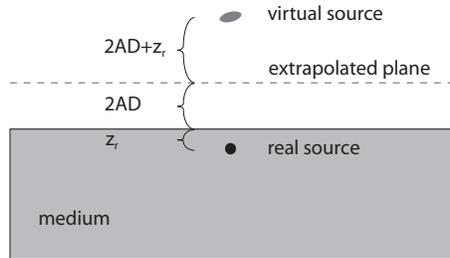


Figure 4: This figure shows the dipole configuration. The real source is placed at a distance z_r below the surface. The extrapolated plane is at $2AD$ above the surface and, finally, the virtual source mirrors the real source in the extrapolated plane. The two sources cancel each other out at points in this plane.

mate condition states that the fluence $\phi(x)$ must be zero at an extrapolated flat surface outside the medium at a distance $z_b = 2AD$ from the physical boundary of the medium. For a semi-infinite slab such as ours this condition can be achieved with two isotropic point sources. The extrapolated bound-

ary condition is a sensible replacement for (23) assuming a linear behavior of the fluence around the geometrical boundaries [CMZ97].

For any point x we can express the distances to the real and virtual light source as

$$\begin{aligned} d_r(x) &= \|(x_i - z_r \cdot \vec{n}) - x\| \\ d_v(x) &= \|(x_i + z_v \cdot \vec{n}) - x\| \end{aligned}$$

And now given equation (19) we arrive at the dipole fluence

$$\phi(x) = \frac{\alpha' P_0}{4\pi D} \left(\frac{e^{-\sigma_{tr} d_r}}{d_r} - \frac{e^{-\sigma_{tr} d_v}}{d_v} \right) \quad (29)$$

which is the approximation of the fluence of the vertical cylindrical beam. The function is clearly a solution to the diffusion equation since equation (19) is. Furthermore, given the chosen distances, z_r and z_v , the fluence $\phi(x)$ is zero at the extrapolated boundary. Finally, this is inserted into approximation (26). Applying $\vec{n} \bullet \nabla$ to the real light source yields

$$\begin{aligned} (\vec{n} \bullet \nabla) \frac{e^{-\sigma_{tr} d_r(x)}}{d_r(x)} &= \vec{n} \bullet \left(\frac{1}{d_r(x)} \nabla e^{-\sigma_{tr} d_r(x)} + e^{-\sigma_{tr} d_r(x)} \nabla \frac{1}{d_r(x)} \right) \\ &= \vec{n} \bullet \left(-\sigma_{tr} \frac{(x_i - z_r \cdot \vec{n}) - x}{d_r^2(x)} e^{-\sigma_{tr} d_r(x)} - e^{-\sigma_{tr} d_r(x)} \frac{(x_i - z_r \cdot \vec{n}) - x}{d_r^3(x)} \right) \\ &= -e^{-\sigma_{tr} d_r(x)} \frac{d_r(x) \sigma_{tr} + 1}{d_r^3(x)} \vec{n} \bullet ((x_i - x) - z_r \cdot \vec{n}) \\ &= z_r e^{-\sigma_{tr} d_r(x)} \frac{d_r(x) \sigma_{tr} + 1}{d_r^3(x)} \end{aligned}$$

A similar expression can be evaluated for the virtual source. In the second step we used $\nabla d_r(x) = \frac{(x_i - z_r \cdot \vec{n}) - x}{d_r(x)}$ and $\nabla \frac{1}{d_r(x)} = -\frac{(x_i - z_r \cdot \vec{n}) - x}{d_r^3(x)}$. Additionally, in the last step $\vec{n} \bullet (x_i - x) = 0$ was used, which is true, when x is located at the surface. In this case, given $r = \|x - x_i\|$, we can express the distances to the real and virtual light source as

$$\begin{aligned} d_r(r) &= \sqrt{r^2 + z_r^2} \\ d_v(r) &= \sqrt{r^2 + z_v^2} \end{aligned}$$

Now assuming the cylindrical beam injects radiance by the power $d\Phi(x_i) =$

P_0 we can finally evaluate approximation (26)

$$\begin{aligned}
R(x, x_i) &\simeq \frac{D(\vec{n} \bullet \nabla \phi)(x)}{d\Phi(x_i)} \\
&= \frac{\alpha'}{4\pi} \left((d_r \sigma_{tr} + 1) \frac{z_r e^{-\sigma_{tr} d_r}}{d_r^3} + (d_v \sigma_{tr} + 1) \frac{z_v e^{-\sigma_{tr} d_v}}{d_v^3} \right) \quad (30)
\end{aligned}$$

Equation (30) is not exactly the same as equation (4) in [JMLH01]. The real term in [JMLH01] is missing a scale by z_r which is clearly an error. Furthermore, they have a division by σ'_t which is not in the derived equation here. In the more recent work by C. Donner and H. Jensen [DJ05] the scale by z_r is included and the division by σ'_t has been removed.

Up until now we have assumed a semi-infinite medium. For such a case light entering the medium will either return to the surface or be absorbed. In the next section we will investigate the scenario for a flat medium of limited thickness known as a slab. For such a case light which has transmitted through to the bottom of the slab may not return to the surface. Furthermore, we investigate a configuration similar to that which we have explained in this section but using stacked slabs.

2.7 The Multipole

In this section we derive and explain the multipole model which is described in Daniele Contini et al. [CMZ97] and also used by Craig Donner et al. [DJ05] to model human skin in the context of computer graphics. The medium is approximated using a configuration of stacked slabs and again the reflectance of an incident laser light is determined. Since a slab has more than one boundary a single dipole is no longer sufficient to satisfy the condition. Furthermore, diffuse radiance can transmit through to the bottom of the slab and exit there. The reflectance and transmittance of a single slab are derived in section 2.7.1. To determine the reflectance or transmittance of a medium of stacked slabs it is necessary to combine profiles in a specific convolution series. This will be explained in section 2.7.2. Finally, in section 2.7.3 we show that the Hankel transformation can be used to transform profiles into frequency space where we do the convolution.

2.7.1 Reflectance of a Slab

In the following a medium of finite thickness s and planar at the top and bottom will be known as a slab. Similar to the previous section a thin laser beam is vertical to the medium and incident at the surface at x_i . We can consider everything above and below the slab as two semi-infinite media. Let the upper semi-infinite medium have index of refraction μ_1 , the slab μ_2 , and the lower semi-infinite medium μ_3 . Let A_0 represent $A(\frac{\mu_2}{\mu_1})$ at the top and A_s represents $A(\frac{\mu_2}{\mu_3})$ at the bottom. Since we are dealing with internal reflection A_0 and A_s are both evaluated with μ_2 in the numerator, i.e., the slab relative to the adjacent medium.

Now similar to section 2.6 the redistribution, after the first bounce, of flux injected by the laser beam can be approximated by a real point source one mean free path $\ell = \frac{1}{\sigma_t}$ below x_i . Once again, to satisfy the extrapolated boundary condition, the virtual light is placed at the distance $\ell + 4A_0D$ above the surface. However, there is a problem. The medium now has an additional boundary which is the bottom of the slab. Given the definition of the extrapolated boundary condition the flux must be zero at a distance $2A_sD$ below the slab. To satisfy this condition the contribution of the inserted dipole must be canceled out by mirroring the dipole about the lower extrapolated boundary plane. Only the position of the second dipole is mirrored. The orientation of polarity remains the same. The second dipole will satisfy the extrapolated boundary condition below the slab but now the condition is compromised above the slab. Subsequently, the second dipole has to be mirrored about the upper extrapolated boundary plane and so on forming an infinite sum of dipoles. The configuration and the first two dipoles are shown in figure 5.

In the following assume \vec{n} is aligned with the Z axis in the positive direction and that position x_i is at origo. Mirroring about some arbitrary coordinate $k_z \in \mathbb{R}$ is done using a simple formula $-(z - k_z) + k_z = 2k_z - z$. We can reference each dipole by an index $i \in \mathbb{Z}$ such that $i < 0$ is a dipole mirrored about the lower boundary at coordinate $k_l = -(s + 2A_sD)$ and $i > 0$ is a dipole mirrored about the upper boundary at coordinate $k_u = 2A_0D$. The dipole $i = 0$ is the initial dipole. Let $j \in \mathbb{N}$ such that the coordinates of

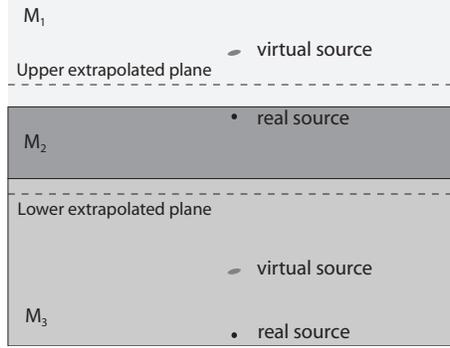


Figure 5: In this figure a slab is shown, and similar to figure 4, a dipole is used to represent the fluence such that the extrapolated boundary condition is satisfied above the surface. This dipole is then mirrored by the lower extrapolated boundary plane to satisfy the boundary condition there. This second dipole must now be mirrored by the upper extrapolated plane and so on.

the dipoles are determined by the following recursive formula.

$$\begin{aligned}
 z_{r,j} &= -z_{v,-j} + 2k_u \\
 z_{v,j} &= -z_{r,-j} + 2k_u \\
 z_{r,-j} &= -z_{v,j-1} + 2k_l \\
 z_{v,-j} &= -z_{r,j-1} + 2k_l
 \end{aligned}$$

Where $z_{r,i}$ and $z_{v,i}$ are the coordinates of the real and virtual components of dipole i . By substituting the fourth equation into the first and the second equation into the third we arrive at

$$\begin{aligned}
 z_{r,j} &= z_{r,j-1} + 2(k_u - k_l) \\
 z_{r,-j} &= z_{r,-(j-1)} - 2(k_u - k_l)
 \end{aligned}$$

which are recursive formulas expressing the coordinates of the real lights after two mirroring steps. For this reason the second equation is only true for $j > 1$ as opposed to the first equation which is true for $j > 0$. This is because the first mirror operation is applied about the lower extrapolated boundary. Similar equations are true for the virtual components. Now given

this recursive formula, of the form $x_{j+1} = x_j + k$, we can rewrite the real and virtual coordinates as

$$\begin{aligned} z_{r,i} &= z_{r,0} + i \cdot 2(k_u - k_l) \\ z_{v,i} &= z_{v,0} + i \cdot 2(k_u - k_l) \end{aligned}$$

where $z_{r,0} = -\ell$ and $z_{v,0} = \ell + 4A_0D$. These equations are true for $i \in \mathbb{Z}$ and express the displacement of the dipole as two times the distance to the extrapolated boundaries. By substitution we arrive at the final formulations

$$z_{r,i} = -\ell + i \cdot 2(2A_0D + s + 2A_sD) \quad (31)$$

$$z_{v,i} = \ell + 4A_0D + i \cdot 2(2A_0D + s + 2A_sD) \quad (32)$$

These formulas are the equivalent of those in [DJ05] and [CMZ97] but negated due to the Z-axis facing down in their work as opposed to up as it is here. In [CMZ97] $\mu_1 = \mu_3$ so in their evaluation $A_0 = A_s$.

In section 2.6 the location of the real and the virtual source was given as distances z_r and z_v . Since they are now determined by signed coordinates the distance to dipole i is determined by

$$\begin{aligned} d_{r,i}(x) &= \|(x_i + z_{r,i} \cdot \vec{n}) - x\| \\ d_{v,i}(x) &= \|(x_i + z_{v,i} \cdot \vec{n}) - x\| \end{aligned}$$

Thus, as in section 2.6, by assuming x is at the surface we can sum up contributions from $2n + 1$ dipoles which gives the following reflectance

$$R(\|x - x_i\|) \simeq \sum_{i=-n}^n \frac{\alpha'}{4\pi} \left((d_{v,i}\sigma_{tr} + 1) \frac{z_{v,i}e^{-\sigma_{tr}d_{v,i}}}{d_{v,i}^3} - (d_{r,i}\sigma_{tr} + 1) \frac{z_{r,i}e^{-\sigma_{tr}d_{r,i}}}{d_{r,i}^3} \right)$$

The equation approximates how much of the contribution from the laser beam gets transferred through the medium and over to the point x at the surface. Similarly, we can evaluate how much gets transferred to the bottom, i.e., transmittance. If we consider x'_i the equivalent location of x_i but at the bottom of the slab, s units below, then we can rewrite the distance formulas as

$$\begin{aligned} d'_{r,i}(x') &= \|(x'_i + (z_{r,i} + s) \cdot \vec{n}) - x'\| \\ d'_{v,i}(x') &= \|(x'_i + (z_{v,i} + s) \cdot \vec{n}) - x'\| \end{aligned}$$

Similar to before, by assuming x' is in the same plane as x'_i , we go through the same procedure as in section 2.6 and arrive at the transmittance

$$T(\|x' - x'_i\|) \simeq \sum_{i=-n}^n \frac{\alpha'}{4\pi} \left((d'_{r,i}\sigma_{tr} + 1) \frac{(z_{r,i} + s) e^{-\sigma_{tr}d'_{r,i}}}{d'^3_{r,i}} - (d'_{v,i}\sigma_{tr} + 1) \frac{(z_{v,i} + s) e^{-\sigma_{tr}d'_{v,i}}}{d'^3_{v,i}} \right)$$

The main difference is the fact that the contributions are negated because at the bottom the inward normal is \vec{n} and not $-\vec{n}$ like it is at the top. This affects equation (24) which means we need to apply $(-\vec{n} \bullet \nabla)$ instead during evaluation.

We can simplify the distance formulas since $\|x - x_i\| = \|x' - x'_i\|$. If we refer to this distance as r then the distance formulas for x and x' , at the surface and bottom respectively, can be written as

$$\begin{aligned} d_{r,i}(r) &= \sqrt{r^2 + z_{r,i}^2} \\ d_{v,i}(r) &= \sqrt{r^2 + z_{v,i}^2} \\ d'_{r,i}(r) &= \sqrt{r^2 + (z_{r,i} + s)^2} \\ d'_{v,i}(r) &= \sqrt{r^2 + (z_{v,i} + s)^2} \end{aligned}$$

which implies that $R(r)$ and $T(r)$ are both functions of r . It should also be noted that this multi-dipole evaluation of $R(r)$ is identical to the original equation (30) when $n = 0$. The contributions made by the real components are negated but this is because $z_{r,0} = -z_r$ since z_r was defined as a distance in section 2.6.

2.7.2 Reflectance of Layers of Slabs

Materials of higher complexity can be approximated by constructing a model of multiple slabs stacked on top of each other. Each slab has constant scattering, absorption and index of refraction values which may differ from those of the other slabs. Reflectance and transmittance functions $R(r)$ and $T(r)$ are evaluated, using these values, for each slab including those of the adjacent upper and lower slab. Evaluation of $R(r)$ and $T(r)$ for a stacked slab is done the same way as explained in section 2.7.1 where the adjacent layers are semi-infinite. Similar to the model given in section 2.7.1, for the case of stacked slabs, it makes sense to consider everything above the top

slab a semi-infinite medium and also everything below the bottom slab is a semi-infinite medium.

If $\tilde{T}(x) = T(\|x\|)$ then for a single slab, when incident flux at the surface x_i is known, the transmitted exitant flux at the bottom x' can be determined

$$\begin{aligned} B(x') &= \int_{\mathbb{R}^2} \tilde{T}(x - x_i) I(x_i) dx_i \\ &= \left(\tilde{T} * I \right) (x) \end{aligned} \tag{33}$$

by convolution of \tilde{T} and the incident flux per unit area at the surface. The integration is done across the surface plane of the slab \mathbb{R}^2 . Note that x' is the point at the bottom of the slab which corresponds to x at the surface. Assuming all layers are at least a few mean free paths in thickness we can assume all interactions between adjacent layers occur due to multiple scattering which will make the diffusion theory applicable. Subsequently, for two consecutive slabs with transmittance profiles T_1 and T_2 we can evaluate the exitant flux at the bottom of the second slab by $B(x'') = \left(\tilde{T}_2 * \left(\tilde{T}_1 * I \right) \right) (x)$. Because convolution is associative we can convolve the transmittance profiles prior to knowing the incident flux

$$\left(\tilde{T}_2 * \tilde{T}_1 \right) * I = \tilde{T}_2 * \left(\tilde{T}_1 * I \right)$$

However, $\tilde{T}_{12} = \tilde{T}_2 * \tilde{T}_1$ does not account for the full transmittance profile through both slabs. The reason is that diffuse radiance transmitted through the first slab and into the second can scatter around there and then make it back up to the surface of the second slab, which is determined by R_2 . Then it might scatter around in the first slab, again, but then make it back to the bottom which must be determined by a reflectance profile evaluated for the bottom of the first slab. A reflectance profile for the bottom of a slab is done simply by turning the stack of slabs upside down and then proceeding the same way as for the surface of a slab. This presents a need to distinguish between reflectance at the surface $R^+(r)$ and at the bottom $R^-(r)$ of a slab. Additionally, let $T^+(r)$ represent transmittance from surface to bottom and $T^-(r)$ from bottom to surface. Because convolution is distributive and since diffuse radiance can loop around between the first and second slab an

arbitrary amount of times we can evaluate the combined transmittance as

$$\begin{aligned}\tilde{T}_{12}^+ &= \tilde{T}_2^+ * \left(\delta + \tilde{R}_1^- * \tilde{R}_2^+ + \tilde{R}_1^- * \tilde{R}_2^+ * \tilde{R}_1^- * \tilde{R}_2^+ + \dots \right) * \tilde{T}_1^+ \\ \tilde{T}_{12}^- &= \tilde{T}_1^- * \left(\delta + \tilde{R}_2^+ * \tilde{R}_1^- + \tilde{R}_2^+ * \tilde{R}_1^- * \tilde{R}_2^+ * \tilde{R}_1^- + \dots \right) * \tilde{T}_2^-\end{aligned}$$

where we have used that $f * \delta = f$. For clarity the convolution sequence is given here right to left though the terms can be refactored since convolution is also commutative. For instance the series of convolved reflectance profiles in \tilde{T}_{12}^- and in \tilde{T}_{12}^+ are really identical. Similar combined reflectance profiles are given by

$$\begin{aligned}\tilde{R}_{12}^+ &= \tilde{T}_1^- * \left(\delta + \tilde{R}_2^+ * \tilde{R}_1^- + \tilde{R}_2^+ * \tilde{R}_1^- * \tilde{R}_2^+ * \tilde{R}_1^- + \dots \right) * \tilde{R}_2^+ * \tilde{T}_1^+ + \tilde{R}_1^+ \\ \tilde{R}_{12}^- &= \tilde{T}_2^+ * \left(\delta + \tilde{R}_1^- * \tilde{R}_2^+ + \tilde{R}_1^- * \tilde{R}_2^+ * \tilde{R}_1^- * \tilde{R}_2^+ + \dots \right) * \tilde{R}_1^- * \tilde{T}_2^- + \tilde{R}_2^-\end{aligned}$$

The reflectance and transmittance profile of stacks with more than two slabs can be determined by iteratively combining current profiles with those of the next slab in the stack. The combined reflectance profile of a stack is, unlike the reflectance profile of a semi-infinite medium (see equation (30)), not a simple analytical function. For this reason [DJ05] does a discrete frequency space evaluation which will be covered in the next section.

2.7.3 Combining Reflectance Profiles

The reflectance and transmittance profiles derived in section 2.7.2, for a stack of slabs, are comprehensive convolution sequences. Because of this Craig Donner [DJ05] decides to use the classic *convolution theorem* which states that

$$\mathcal{F}\{f * g\} = k \cdot \mathcal{F}\{f\} \cdot \mathcal{F}\{g\}$$

where $\mathcal{F}\{f\}$ denotes the Fourier transform of f and k depends on the chosen normalization of the Fourier transform. For some chosen disk size $d > 0$ the Fourier transformed profiles are given by

$$\begin{aligned}\mathcal{R}(u, v) &= \mathcal{F}\{\tilde{R}\} = \int_{x^2+y^2 < d^2} \tilde{R}(x, y) e^{-2\pi i(ux+vy)} dx dy \\ \mathcal{T}(u, v) &= \mathcal{F}\{\tilde{T}\} = \int_{x^2+y^2 < d^2} \tilde{T}(x, y) e^{-2\pi i(ux+vy)} dx dy\end{aligned}$$

In the typical continuous case $d = \infty$ but to do numerical evaluation it is assumed that values of significant contribution are within a finite domain. For this reason we denote the upper limit by d which is more general. It follows that, in frequency space, the convolution sequences of the combined profiles is evaluated using

$$\begin{aligned}\mathcal{R}_{12}^+ &= \mathcal{T}_1^- \cdot \left(1 + \mathcal{R}_2^+ \mathcal{R}_1^- + (\mathcal{R}_2^+ \mathcal{R}_1^-)^2 + (\mathcal{R}_2^+ \mathcal{R}_1^-)^3 + \dots\right) \cdot \mathcal{R}_2^+ \cdot \mathcal{T}_1^+ + \mathcal{R}_1^+ \\ \mathcal{R}_{12}^- &= \mathcal{T}_2^+ \cdot \left(1 + \mathcal{R}_1^- \mathcal{R}_2^+ + (\mathcal{R}_1^- \mathcal{R}_2^+)^2 + (\mathcal{R}_1^- \mathcal{R}_2^+)^3 + \dots\right) \cdot \mathcal{R}_1^- \cdot \mathcal{T}_2^- + \mathcal{R}_2^- \\ \mathcal{T}_{12}^+ &= \mathcal{T}_2^+ \cdot \left(1 + \mathcal{R}_1^- \mathcal{R}_2^+ + (\mathcal{R}_1^- \mathcal{R}_2^+)^2 + (\mathcal{R}_1^- \mathcal{R}_2^+)^3 + \dots\right) \cdot \mathcal{T}_1^+ \\ \mathcal{T}_{12}^- &= \mathcal{T}_1^- \cdot \left(1 + \mathcal{R}_2^+ \mathcal{R}_1^- + (\mathcal{R}_2^+ \mathcal{R}_1^-)^2 + (\mathcal{R}_2^+ \mathcal{R}_1^-)^3 + \dots\right) \cdot \mathcal{T}_2^-\end{aligned}$$

It is pointed out in [DJ05] that when $|\mathcal{R}_2^+ \mathcal{R}_1^-| < 1$ the geometric series can be simplified.

$$\frac{1}{1 - \mathcal{R}_2^+ \mathcal{R}_1^-} = 1 + \mathcal{R}_2^+ \mathcal{R}_1^- + (\mathcal{R}_2^+ \mathcal{R}_1^-)^2 + (\mathcal{R}_2^+ \mathcal{R}_1^-)^3 + \dots$$

Donner uses the discrete *fast Fourier transform* FFT in [DJ05] to create samples of the combined profiles. In a more recent paper [DJ06a] by Donner the transformation is done using the *Hankel transform* which will be shown in the following. The Hankel transform $\mathcal{H}\{f\}$ of order zero is defined by

$$F_0(k) = \int_0^\infty f(r) J_0(k \cdot r) r dr \quad (34)$$

where f is a function of a single variable and $J_0(x)$ is the *Bessel function* of order zero. The Bessel functions (of the first kind) of order n can be expressed as the integral

$$J_n(z) = \frac{1}{2\pi i^n} \int_0^{2\pi} e^{iz\cos\theta} e^{in\theta} d\theta$$

From this it follows that $J_0(x)$, for $x \in \mathbb{R}$, is equal to

$$\begin{aligned}J_0(x) &= \frac{1}{2\pi} \int_0^{2\pi} e^{ix\cos\theta} d\theta \\ &= \frac{1}{2\pi} \int_0^{2\pi} e^{-ix\cos\theta} d\theta\end{aligned} \quad (35)$$

where the last step follows because the imaginary component of the result is zero and because $J_0(x)$ is an even function on the real domain. Now, given the following definitions

$$\begin{aligned} x &= r \cos \theta & y &= r \sin \theta & r &= \sqrt{x^2 + y^2} \\ u &= q \cos \phi & v &= q \sin \phi & q &= \sqrt{u^2 + v^2} \end{aligned}$$

we can rewrite the Fourier transformed reflectance profile

$$\begin{aligned} \mathcal{F}\{\tilde{R}\} &= \int_{x^2+y^2 < d^2} \tilde{R}(x, y) e^{-2\pi i(ux+vy)} dx dy \\ &= \int_0^{2\pi} \int_0^d R(r) e^{-2\pi i q r \cos(\theta-\phi)} r dr d\theta \\ &= \int_0^d \int_{0-\phi}^{2\pi-\phi} R(r) r e^{-2\pi i q r \cos \theta} d\theta dr \\ &= \int_0^d R(r) \int_0^{2\pi} e^{-i2\pi q r \cos \theta} d\theta r dr \\ &= 2\pi \int_0^d R(r) J_0(2\pi q r) r dr \end{aligned}$$

The resulting version of the Fourier transform of a radially symmetric function already appears to be very similar to the Hankel transform (see equation (34)). The conversion from Fourier to Hankel transform can be completed using substitution with $r' = 2\pi r$ and $dr' = 2\pi dr$. If we additionally define the function $g(r') = \frac{r'}{2\pi}$ we arrive at the result

$$\begin{aligned} \mathcal{F}\{\tilde{R}\} &= \frac{1}{2\pi} \int_0^{2\pi d} R\left(\frac{r'}{2\pi}\right) J_0(qr') r' dr' \\ &= \frac{1}{2\pi} \int_0^{2\pi d} (R \circ g)(r') J_0(qr') r' dr' \\ &= \frac{\mathcal{H}\{R \circ g\}}{2\pi} \end{aligned} \tag{36}$$

In other words, the Fourier transformed reflectance profile \tilde{R} is equal to the Hankel transformation of $R(\frac{r'}{2\pi})$ divided by 2π . The same principle applies to the transmittance \tilde{T} since it is also a radially symmetric function. Note that if the Fourier transformation is done over the domain with disc size

d then the corresponding Hankel transformation is done with a domain of disc size $2\pi d$. This is a result of the substitution step which is significant for $d < \infty$. Another useful consequence of equation (36) is that the Fourier transformation of a real radially symmetric function is again a real radially symmetric function.

2.8 Resulting Profiles

In this section, reflectance profiles generated using the dipole- and multipole methods are tested and compared to evaluation of similar profiles done using Monte Carlo simulation. The software used for this simulation is called (MCML) Monte Carlo for Multi-Layered media and was written by Lihong Wang and Steven L. Jacques [WJ]. The software allows the user to configure a model of stacked slabs by thickness, refractive index and scattering- and absorption coefficients for each slab. The simulation evaluates reflectance and transmittance profiles based on an infinitely narrow, vertically incident, photon beam as the light source. We have configured MCML to use 1000000 photons for its evaluation. The multipole model approximates the same configuration but will only take reflectance due to multiple scattering into account. At the point x_i , where the beam is incident at the surface, first and second order scattering will dominate results and subsequently the multipole (and dipole) model is predicted to produce results less than that of MCML close to the center. All results in the following are given in millimeters.

Slab	σ_a			σ'_s			μ	thickness
	R	G	B	R	G	B		
skin	0.032	0.17	0.48	0.74	0.88	1.01	1.3	∞

Table 1: This table contains, for a skin sample, the absorption- and reduced scattering parameter reported in [JMLH01].

In [JMLH01] measurements of σ_a and σ'_s are made for various materials such as: ketchup, chicken, milk, skin, etc. This is done by focusing a vertical beam of white light on the surface of a sample. An HDR picture is taken to record the radiant exitance across the surface resulting from the beam of light. The reduction in radiance caused by transmittance in and out, see eq.

(28), through the interface is canceled out in [JMLH01] by exploiting the fixed constellation of the camera. Jensen is using the dipole approximation and subsequently equation (30) is used in a least-squares fit to reverse the best match for σ_a and σ'_s . To further stabilize results Jensen computes the total diffuse reflectance R_{tot} from the picture and then uses this as an additional constraint during fitting.

$$R_{tot} = \int_{\mathbb{R}^2} R(\|x_i - x\|) dx \quad (37)$$

Results for fitting of two skin samples are given in [JMLH01]. Table 1 in this paper is the reported result of the first skin sample. A comparison between evaluation made by MCML and equation (30), with a logarithmic distribution along the Y -axis, is shown in figure 7(a). The dipole approximation, in this case, appears to agree very well with MCML and as predicted equation (30) underestimates reflectance close to x_i . Figure 7(b) shows a close-up of a linear distribution along the Y -axis and though the graphs do not align exactly there is a distinct conceptual match. This match is further improved as the distance from the center is increased. Ultimately, the diffuse BSSRDF, S_d , is to be used in equation (4). To get a clearer picture of where the most significant contributions, due to subsurface scattering, are received from (or passed on to) we need to take into account that contributions are weighted by a differential surface area. Given that $\int_{\mathbb{R}^2} R(\|x_i - x\|) dx = 2\pi \int_0^\infty R(r) r dr$ it makes sense to take a look at $rR(r)$ which is shown in figure 6(a). The graph shows a peak near the center but at the actual center contributions due to multiple scattering are negligible. Furthermore, the graph reveals that reflectance values for green and blue tend to zero faster than those of red. This agrees very well with the profile of Caucasian skin which tends to exhibit red tones within darker regions such as wrinkles and shadows.

In [DJ05] the idea of modeling human skin with a semi-infinite medium is challenged. In optics skin is modeled as multiple layers with different material properties. A typical model, see [Tuc00], divides the skin into two primary layers. The first is the *epidermis* and the second is the *dermis*. In [DJ05] the dermis is further divided into two separate layers the upper dermis and the bloody dermis. The values used in [DJ05] are given here in table 2. Furthermore, the combined reflectance profile is now evaluated using the multipole method explained in section 2.7. As in [DJ06a] the Hankel transformation of

Slab	σ_a			σ_s			μ	\mathbf{g}	thickness
	R	G	B	R	G	B			
epidermis	2.1	2.1	5.0	48.0	60.0	65.0	1.4	0.0	0.03
upper dermis	0.16	0.19	0.3	32.0	40.0	46.0	1.34	0.25	0.05
bloody dermis	0.085	1.0	25.0	4.5	4.7	4.8	1.4	0.8	∞

Table 2: This table contains the material properties for the multi-slab model of human skin used in [DJ05].

order zero is used to transform reflectance and transmittance profiles of individual slabs into frequency space before convolution. This transformation step is performed using equation (36). Specifically, for the implementation used in this paper discrete transformation is done using GSL – The GNU Scientific Library.

The table values are also fed to MCML and a comparison between the two resulting reflectance profiles is shown in figure 7(c) with a logarithmic distribution along the Y -axis. Again there appears to be a really good correspondence between the Monte Carlo simulation and that of the (multi) dipole approximation. For the case of a linear distribution along the Y -axis, in figure 7(d), there is still a close match between that of MCML and the multipole method. Unfortunately, when switching to 6(b), which shows $rR(r)$, a problem is revealed. Reflectance, due to multi scattering, is focused at the center to an extent which makes all other contributions across the surface insignificant. Specifically, 99% of the total diffuse reflectance is spent within the range $r \in [0; 0.01]$. In other words the resulting profile is similar to that of a BRDF which does not take subsurface scattering into account. In [DJ05] the values in table 2 are credited directly to the work of Valery Tuchin and a reference is made to the book [Tuc00]. Since a BRDF like behavior is clearly an incorrect profile for skin we have since then acquired a copy of [Tuc00]. We were not able to find table 2 in this book so we made contact with Craig Donner who confirmed the values are in fact only based on the tables of [Tuc00] and were reworked and adapted during development of their paper. There is no plot for the reflectance profile shown in [DJ05]. For this reason we tried to compare results to a different example by Craig Donner given in [DJ06a]. The example consists of two slabs and material parameters are given in figure 4a in [DJ06a]. The same parameters are given here in table 3 and graphs for logarithmic and linear distribution along the Y -axis are given

Slab	σ_a	σ'_s	μ	thickness
first slab	0.005	1.3	1.4	6.0
second slab	0.01	1.0	1.4	∞

Table 3: This table contains the material properties of two slabs used to generate fig 4a in [DJ06a].

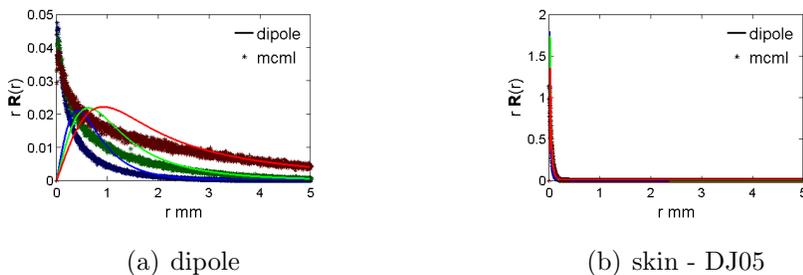
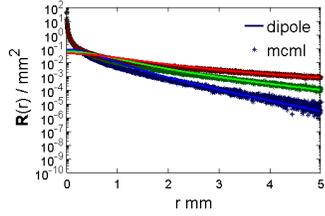


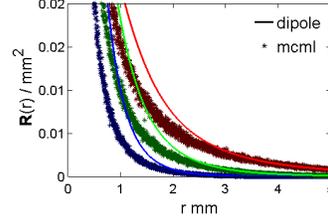
Figure 6: In figures 6(a) and 6(b) we see the reflectance multiplied by the distance $rR(r)$. These were made from tables 1 and 2 respectively.

in figures 7(e) and 7(f) respectively. Once again results between MCML and the multipole method agree very well. Additionally, this time the graph given in [DJ06a] in figure 4a agrees with the result produced here.

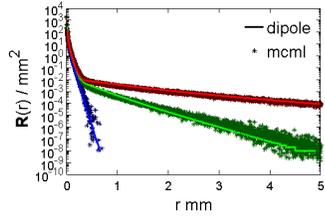
Alas, this example is not that of a skin profile. Since the reflectance profile produced by our multipole implementation matched that of MCML a possibility is that the table, as printed in the paper [DJ05], could contain errors. Though the epidermis and upper dermis in table 2 are thin a possible issue is the specified absorption and scattering parameters are large relative to those given in [JMLH01]. Both papers define values relative to units in millimeters. Large values will increase the effective transport coefficient σ_{tr} (given in section 2.4.2). A large coefficient for σ_{tr} will force reflectance and transmittance to rapidly decline with distance (see eq. (30)). Regardless, we have decided to abandon table 3 in [DJ05].



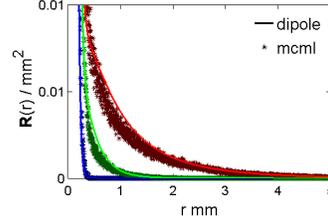
(a) dipole logarithmic



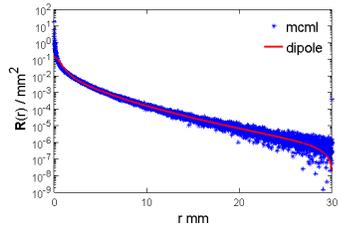
(b) dipole linear



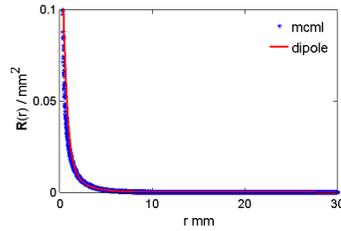
(c) skin – multipole logarithmic



(d) skin – multipole linear



(e) multipole logarithmic



(f) multipole linear

Figure 7: In figures 7(a), 7(c) and 7(e) we see the reflectance profiles resulting from tables 1, 2 and 3 respectively. These are presented with a logarithmic distribution along the Y-axis. The same profiles are shown with a linear distribution in figures 7(b), 7(d) and 7(f).

3 The Integration Step

In section 2 we focused on the reflectance profile. Once such a profile is obtained it must be used in equation (4), during rendering, to determine outgoing radiance. To do this we must numerically compute an integration step across the scene. The reflectance profile is derived based on the assumption that the medium consists of a configuration of stacked slabs which implies an entirely flat surface. An arbitrary scene to be rendered is generally not flat but for the process to remain practical this is ignored during integration in [JMLH01] and [DJ05].

Another significant subtlety is the reflectance profile was derived based on the assumption that sources are isotropic. However, this assumption was only applied to the initial internal scattering event which redistributes transmitted irradiance. Subsequently, we can use the reflectance profile with any type of light including area light sources.

In this section we will focus on different approaches to solving the integration. First we give the final simplified version of equation (4) which is obtained by inserting the derived transfer function. Next since integration across the entire scene, for every pixel, is ineffective integration is restricted to a finite range. These details are discussed in section 3.1.

The integration step is further simplified by exploiting that the derivative of the reflectance profile (and the fall-off) tends to zero. This is done by using an octree of sample points to integrate at a lower resolution as distance becomes greater. This approach was introduced by Jensen et al. [JB02] and is described here in section 3.2.

A more recent approach by Eugene d'Eon et al. [dL07] completes the integration step on a GPU at an interactive framerate. This is done by assuming the reflectance profile is given as a sum of Gaussians, which are separable, and by performing the integration in a 2D texture unwrap. We describe their method here in section 3.4. and suggest an improvement to stabilize results near the boarder of the unwrap.

3.1 Preliminaries

By substituting (28) into (6) we arrive at the final BSSRDF

$$S(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) \simeq \frac{1}{\pi} \rho_t(x_o, \vec{\omega}_o) R(\|x_o - x_i\|) \rho_t(x_i, \vec{\omega}_i) + S_r(x_o, \vec{\omega}_i, \vec{\omega}_o) \quad (38)$$

where $R(\|x\|)$ is the previously, assumed to be, radially symmetric reflectance profile. If we assume the second term is defined as,

$$S_r(x_o, \vec{\omega}_i, \vec{\omega}_o) = \delta(\|x_o - x_i\|) \cdot f_r(x_o, \vec{\omega}_i, \vec{\omega}_o)$$

where f_r is some chosen BRDF, then by substituting (38) into equation (4) we get the following approximation

$$L_o(x_o, \vec{\omega}_o) \simeq \frac{\rho_t(x_o, \vec{\omega}_o)}{\pi} \int_S R(\|x_o - x_i\|) \int_{\Omega_{2\pi}} \rho_t(x_i, \vec{\omega}_i) L(x_i, \vec{\omega}_i) (\vec{\omega}_i \bullet \vec{n}_i) d\vec{\omega}_i dx_i + \int_{\Omega_{2\pi}} f_r(x_o, \vec{\omega}_i, \vec{\omega}_o) L(x_i, \vec{\omega}_i) (\vec{\omega}_i \bullet \vec{n}_i) d\vec{\omega}_i \quad (39)$$

Preferably, f_r should be the same BRDF as that which is used in equations (7) and (8) for diffuse reflectance and transmittance.

It is now clear that the inner integral of the first term of equation (39) collects the total transmitted irradiance at every point x_i . This is the flux per area unit transmitted through the interface. Next the outer integral weights every such contribution using the reflectance profile which converts incoming differential flux into outgoing. The division by π converts the radiosity into diffuse outgoing radiance and finally the result is transmitted back out through the interface at x_o in the direction $\vec{\omega}_o$.

Integrating across the entire scene for every visible point x_o is not a realistic option. However, the reflectance profile, $R(\|x\|)$, is a monotonically decreasing function and is subject to exponential decay. This is given by equation (30) for the dipole approximation and in section 2.7.1 for the multipole. Furthermore, the decay constant is given by the transport coefficient σ_{tr} which was defined in section 2.4.2.

When using the dipole approximation the reflectance function is analytical and a practical approach to determining the effective integration range d_{max} is choosing this range such that

$$\frac{2\pi}{R_{tot}} \int_0^{d_{max}} R(r) r dr = 1 - \varepsilon$$

for some chosen ε . Note that the left side of this equation is the ratio between the diffuse reflectance within a radius of d_{max} and the total diffuse reflectance. This ratio is in the range zero to one which means $0 < \varepsilon < 1$. We have used

Maple 12 to solve the integral. After substitution of d_{max} by s and isolating ε we are given this equation.

$$\varepsilon(s) = \frac{z_r d_v(s) e^{-\sigma_{tr} d_r(s)} + z_v d_r(s) e^{-\sigma_{tr} d_v(s)}}{d_r(s) d_v(s) (e^{-\sigma_{tr} z_v} + e^{-\sigma_{tr} z_r})} \quad (40)$$

What we need is the inverse function which does not have an analytical solution. However, the function is monotonically decreasing so for any choice of $s \in]0; \infty[$ there is a unique $\varepsilon \in]0; 1[$ and vice versa. Thus we can determine s from a user-defined error ε_0 in Maple using `fsolve($\varepsilon_0 = \varepsilon(s)$)`. This is equivalent to finding the root of the function $g(s) = \varepsilon_0 - \varepsilon(s)$ which can be done using an algorithm such as the Newton-Raphson method or the Secant method.

For a discrete reflectance profile it is possible to search for d_{max} by iteratively generating the table for $R(r)$ over a range which is increased by some delta per iteration. The search for d_{max} is terminated once a certain amount of iterations have been performed or once the changes to d_{max} have become insignificant. Note that the reflectance profile is generated only once per material. The range d_{max} on the other hand can be used during evaluation of equation (39) to reduce the execution time significantly. Instead of integrating across the entire scene per x_o we only have to integrate over the intersection S between all surfaces of the scene and a ball at x_o with radius d_{max} .

The reflectance profile is generated under the assumption that the medium has a planar surface. This limitation is removed in both [JMLH01] and [DJ05] by simply passing the distance $r = \|x_o - x_i\|$ to $R(r)$ regardless of orientation relative to the surface normal at x_o . This solution is technically incorrect but preserves practicality of the approach. Conceptually, the reflectance is reduced at an exponential speed, by distance to the source, whether the surface is exactly planar or not. For a thin medium, such as paper, it is suggested in [DJ05] to use a simple remedy such as

$$t = \frac{1 - (\vec{n}_i \bullet \vec{n}_o)}{2}$$

$$P(r) = (1 - t)R(r) + tT(r)$$

where $P(r)$ is a blend between the transmittance and reflectance based on the angle between the surface normals at x_i and x_o . The combined profile $P(r)$ is subsequently used instead of $R(r)$ in equation (39).

Equation (39) can generally not be solved analytically. Some form of discretization needs to take place followed by numerical integration. In the following different ways to complete this task will be described.

3.2 Two-pass integration

Adjacent pixels often represent points on a surface within a close proximity of one another. Specifically, a point, x_o on a surface to be shaded is very likely to be within d_{max} distance of other points to be shaded. The potential of a frequent, and significant, overlap between points to be shaded motivates the idea of caching results for the inner integral in equation (39). This is done for a chosen distribution of N sample points. Once this is done the outer integral can be evaluated in a subsequent pass. To evaluate the outer integral the differential area at the surface associated with each sample point x_i must be known. If the sample points have a reasonably even distribution then this differential area is trivially

$$dx_i = \frac{A_{tot}}{N}$$

where A_{tot} is the total surface area of a mesh and N is the number of sample points distributed across the mesh. If the distribution of sample points is not even then it is possible to determine dx_i using a Voronoi diagram. However, this also makes the procedure more complicated.

The two-pass approach is used by [JB02] to accelerate the integration. In the first pass sample points are evenly distributed using Turk’s point repulsion algorithm [Tur92]. Let triangle $j \in \mathbb{N}$ be denoted T_j and let the area of this triangle be known as A_j . For a total budget of N sample points the algorithm works the following way

1. Randomly assign each sample point to a triangle, T_j , such that the probability of choosing T_j is $P(x_i \in T_j) = \frac{A_j}{A_{tot}}$.
2. Randomly position each point x_i assigned to T_j inside the triangle.
3. Refine results by applying repulsion between all points distributed across the mesh.

The first step is achieved by creating a random positive number in the range $[0; A_{tot}]$ and then by doing a binary search through a list of partial sums

of the triangle areas in the model. The second step requires the ability to generate evenly distributed points inside a triangle. This can be achieved using the following equations also given by Turk in Graphics Gems [Tur93]

$$\begin{aligned} u &= 1 - \sqrt{\xi_1} \\ v &= \xi_2 \cdot \sqrt{\xi_1} \\ p_{rnd} &= (1 - u - v) \cdot p_0 + u \cdot p_1 + v \cdot p_2 \end{aligned}$$

where $\xi_1, \xi_2 \in [0; 1]$ are two randomly generated numbers and p_0, p_1 and p_2 are the three vertices of the triangle. In the final step the distribution of sample points is refined using relaxation by repelling neighboring points. Turk lets the radius of repulsion be given by $r = 2\sqrt{dx_i}$. Two points which have a distance of r or more between them do not affect each other. The repulsion force is chosen by Turk such that it increases linearly as distance decreases below r . For more information about the repulsion algorithm the reader is referred to [Tur92].

The choice of the number of sample points, N , is made in [JB02] based on the mean free path ℓ . Since this is the average distance until the next interaction inside the medium it makes sense to consider this the maximum distance any point should have to its closest neighbor. Subsequently, the number of samples is chosen as

$$N = \frac{A_{tot}}{\pi\ell^2}$$

where the denominator is the area of a small disk with radius ℓ . This can be thought of as the differential area dx_i . For every resulting sampling point the transmitted irradiance

$$I(x_i) = \int_{\Omega_{2\pi}} \rho_t(x_i, \vec{\omega}_i) L(x_i, \vec{\omega}_i) (\vec{\omega}_i \bullet \vec{n}_i) d\vec{\omega}_i$$

determined by the inner integral, in equation (39), is stored along with the position of the sampling point. If a nonuniform distribution is used then the differential area dx_i must be stored too.

In the second pass when the outer integral is processed for a pixel, corresponding to some point x_o on a surface, all sample points, x_i , within a distance of d_{max} are collected. The transmitted irradiance of each sample is scaled by $R(\|x_o - x_i\|)dx_i$ and accumulated. Once all contributions have been

accumulated the final result is scaled by $\frac{\rho_t(x_o, \vec{\omega}_o)}{\pi}$. To accelerate the process of collecting all sample points within the proper range it is appropriate to use a hierarchical structure. In [JB02] an octree is used with up to eight samples in a leaf node to make the implementation more efficient. However, collecting all sample points within a distance of d_{max} , even with the octree, is still a costly process. For this reason Jensen takes advantage of the fact that the reflectance profile is essentially made up of a reciprocal exponential curve e^{-kx} . In other words, as x increases, change in reflectance value slows down exponentially. In fact values tend to zero. From this we can conclude as distance to x_o increases we can approximate by applying the same reflectance value to a wider range of sample points. This allows us to preaccumulate the transmitted irradiance weighted by the differential area for all sample points within the boundary of a node. This represents the accumulated radiant flux contained within node $k \in \mathbb{N}$ and we shall denote this quantity

$$\Phi_k = \sum_{i=1}^{N_k} I(x_i) dx_i$$

where N_k is the number of sample points within the boundary. Thus for some appropriately chosen center P_k relative to the N_k sample points we can approximate the outer integral using

$$R(\|x_o - P_k\|) \cdot \Phi_k \simeq \sum_{i=1}^{N_k} R(\|x_o - x_i\|) \cdot I(x_i) dx_i \quad (41)$$

The center P_k can be chosen in a number of different ways. One option is to determine the smallest possible bounding sphere for the N_k sample points. Another is to simply average the sample points and a third option is to crudely use the center of the node itself. In [JB02] yet another option is chosen where the center is a weighted average of the sample points

$$P_k = \frac{\sum_{i=1}^{N_k} x_i I(x_i)}{\sum_{i=1}^{N_k} I(x_i)}$$

where weights are determined by the transmitted irradiance. In other words P_k is chosen such that it is closer to the brightest contributions.

The process of evaluating the outer integral now becomes recursive. The octree is traversed from the top node and for every node visited we check

if it is small enough relative to its distance to x_o . If so we settle with the evaluation provided by approximation (41) and if not the child nodes are evaluated recursively. If it is a leaf node then the exact accumulation given on the right side of approximation (41) is performed. In [JB02] a simple criteria is used to determine whether or not to recurse. Let the value $d_k \in \mathbb{R}$ represent the chosen size of node k . If the ratio between this size and the distance is less than some user-defined value ε then recursion is ended and the node is evaluated using approximation (41).

$$\frac{d_k}{\|x_o - P_k\|} < \varepsilon \quad (42)$$

A logical option when determining a suitable value for d_k is the radius of a minimum bounding sphere containing the N_k sample points. However, in [JB02] another simple evaluation is used based on the total area sum of all differential areas. This way it is taken into account how large a surface area the sample points represent.

$$A_k = \sum_{i=1}^{N_k} dx_i$$

$$d_k = \sqrt{A_k}$$

In [JB02] this is justified by how $\frac{A_k}{\|x_o - P_k\|^2} < \varepsilon^2$ represents a limit on the approximate solid angle containing the surface area within the node. Either way, whether or not, to recurse is determined by (42). To summarize, the values stored with every node are Φ_k , A_k , P_k and d_k and for leaf nodes up to eight sample points are stored including differential areas dx_i when a nonuniform distribution is used.

3.3 Mesh Unwrap-Based Distribution

For a mesh supplied with a mesh unwrap it is possible to distribute samples according to the unwrap. This is done by resolving for every used texel in the texture map the corresponding surface position(s). If the unwrap is generated such that it is approximately conformal to the surface of the mesh then the samples can be considered evenly distributed. Generally, this is not the case and dx_i will have to be evaluated for each sample point. For a triangular

mesh it is shown in [Mik08] how a triangle given as a parametrization $\sigma(s, t)$ of the texture map sampling coordinate (s, t) has first order derivatives

$$\sigma_s = \frac{(t_3 - t_1)_y \cdot (p_2 - p_1) - (t_2 - t_1)_y \cdot (p_3 - p_1)}{(t_2 - t_1)_x \cdot (t_3 - t_1)_y - (t_2 - t_1)_y \cdot (t_3 - t_1)_x} \quad (43)$$

$$\sigma_t = \frac{-(t_3 - t_1)_x \cdot (p_2 - p_1) + (t_2 - t_1)_x \cdot (p_3 - p_1)}{(t_2 - t_1)_x \cdot (t_3 - t_1)_y - (t_2 - t_1)_y \cdot (t_3 - t_1)_x} \quad (44)$$

where p_1, p_2 and p_3 are the vertices of the triangle and t_1, t_2 and t_3 are the texture coordinates of the unwrap assigned to the triangle. The absolute value of the denominator is actually equal to two times the area, in the texture, spanned by the texture coordinates

$$\begin{aligned} A_{map} &= \left| \det \left[\begin{array}{c|c} t_2 - t_1 & t_3 - t_1 \end{array} \right] \right| \\ &= \left| (t_2 - t_1)_x \cdot (t_3 - t_1)_y - (t_2 - t_1)_y \cdot (t_3 - t_1)_x \right| \end{aligned}$$

The alignment of σ_s and σ_t shows how horizontal and vertical lines, in the texture, map to the surface and the magnitudes $\|\sigma_s\|$ and $\|\sigma_t\|$ are equal to units traveled across the surface per unit in the texture in the corresponding direction. We can determine two times the area of the triangle as $A_{tri} = \|(p_2 - p_1) \times (p_3 - p_1)\|$ and from this it follows that we can determine the surface to texture ratio as

$$\begin{aligned} dx_i &= \|\sigma_s \times \sigma_t\| \\ &= \frac{A_{tri}}{A_{map}} \end{aligned}$$

which is the differential area dx_i associated with the sample point x_i . Note however that every texel may belong to more than one triangle. A practical way to distribute sample points is to rasterize in 2D each triangle by its texture coordinates and interpolate the 3D positions of the triangle. Thus the interpolated position becomes the sample point on the mesh and ambiguity is solved.

3.4 Mesh Unwrap–Based Integration

As opposed to distributing samples across the actual surface of the mesh one might approximate by performing the integration based on distances between

texels and thus completing the integration in the two dimensional unwrap. This assumes connectivity of the unwrap is roughly equivalent to that of the triangular mesh. In other words such that adjacent triangles of the mesh are generally adjacent in the unwrap. The technique was introduced in [BL03] to render faces in the Matrix sequels. The method is done in two primary passes.

1. Render irradiance to a texture map.
2. Convolve against the reflection profile

The first pass is done in a similar way to that of the previous section. We rasterize in 2D each triangle by its texture coordinates and interpolate attributes such as the 3D positions and normals of the triangle. This way we determine the surface position and normal associated with each used texel from the interpolated attributes. Next the inner integral in equation (39) is evaluated and the resulting transmitted irradiance is stored in the texture. Unlike the previous section each texel must map uniquely to one surface point. Thus overlaps in the unwrap, such as mirroring, are prohibited. This is because irradiance is determined from the position (and the normal) and not the location of the texel.

In the second pass, to complete the outer integration, the irradiance texture is convolved with the reflectance profile. The dipole approximation is based on a reciprocal exponential curve and the multipole is based on a summation of these. This typically gives an initial spike followed by a slow fall-off which results in a broad base. As pointed out in [BL03] the broad base provides a wide blur while the spike preserves the original detail. Thus in their paper a conceptual approximation is used based on

$$R_k(r) \simeq \frac{1}{(c+r)^k} \tag{45}$$

which also gives a spike and a broad base. In this approximation c is some small constant used to avoid division by zero at $r = 0$ and k is a tweakable user-defined parameter used to control the behavior of the reflectance profile. A comparison between the physically based dipole approximation given in [JMLH01] and the more adhoc approach in [BL03] is shown in figure 8. No specific choices for k were given in [BL03] for red, green and blue so I have simply chosen values which give profiles reminiscent of those in figure 8(a).

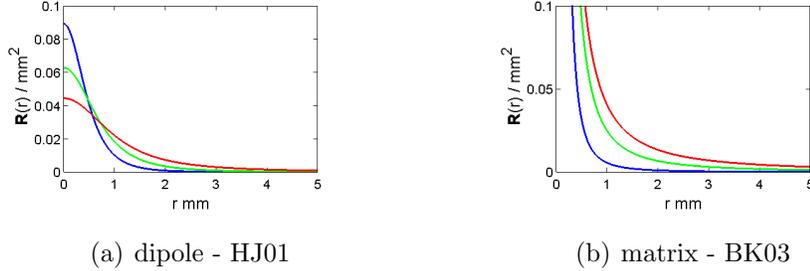


Figure 8: The dipole approximation for skin given in [JMLH01] is shown in figure 8(a). A different empirical approach given in [BL03] is shown in figure 8(b). Though they are numerically different they are conceptually similar. Both have a spike and a broad base.

The resemblance is more striking when compared to the close-up of figure 8(a) shown in figure 7(b).

Given that the method determines reflectance based on distances between texels in the irradiance texture it strictly works for local subsurface scattering. The technique is not able to transmit irradiance through thin layers such as the ears and nostrils. This is because the distance in texels would represent the distance along the surface of the ear/nostril from one side to the other as opposed to the distance through the medium which is much less. This is a limitation also pointed out by the authors themselves. Additional problems which were not addressed in [BL03] are issues such as the fact that approximation (45) is not separable which may lead to a very significant performance hit when convolution is performed. Another problem is what to do about unused texels in the irradiance texture? What to do when locality at the surface does not correspond to locality within the unwrap? And ultimately how to convert from a distance in texels to the corresponding distance across the surface of the mesh.

The more recent work of Eugene d'Eon et al. [dL07] adapts the work of [BL03] to a more GPU friendly implementation and furthermore addresses all of these problems. As opposed to using approximation (45) a weighted sum of Gaussians is used instead to represent the reflectance profile. The rationale for using multiple Gaussians is that a single Gaussian is not sufficient to

σ^2	Blur Weights		
	Red	Green	Blue
0.0064	0.233	0.455	0.649
0.0484	0.1	0.336	0.344
0.187	0.118	0.198	0
0.567	0.113	0.007	0.007
1.99	0.358	0.004	0
7.41	0.078	0	0

Table 4: These are the weights and the variance σ^2 values, given in [dL07], for the sum of Gaussians used to approximate the reflectance profile of a Caucasian male.

represent the spike and broad base of the reflectance profile.

The two-dimensional Gaussian is given as

$$G(\sigma^2, r) = \frac{1}{2\pi\sigma^2} e^{-\frac{r^2}{2\sigma^2}}$$

and for a fixed number n of Gaussians, and for a known reflectance profile $R(r)$, variance σ_k^2 and weights w_k are chosen such that the squared error is minimized

$$\int_0^\infty r \cdot \left(R(r) - \sum_{k=0}^n w_k G(\sigma_k^2, r) \right)^2 dr$$

In [dL07] a table of values are given such that they represent the reflectance profile of a Caucasian male. According to Eugene d'Eon the values in the table were made using an already made image and a similar image rendered by Eugene himself using a BRDF (no subsurface scattering). From these two images the best match for the weights are reversed. Currently, no public documentation exists on the explicit details of this work. The resulting values are given in [dL07] and also here in table 4. The exact graph when using this table in the sum of six Gaussians is shown in figure 9.

Unlike approximation (45) the sum of Gaussians profile is based on a reciprocal exponential function similar to the multipole (and dipole) technique. The difference is the squared distance r^2 is used directly as opposed to the distance formulas $d_{r,i}(r)$ and $d_{v,i}(r)$ given in section 2.7.1. The significant difference is that Gaussians are separable which allows for a very efficient

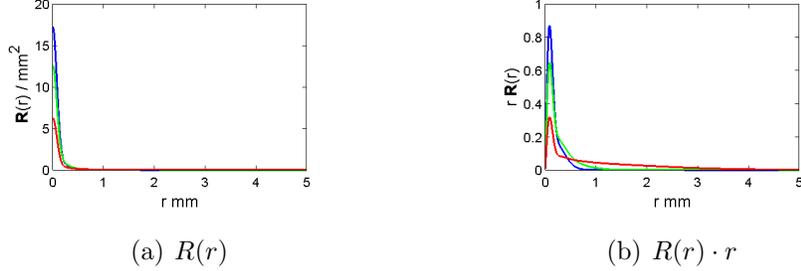


Figure 9: In figure 9(a) we see the reflectance profile used in [dL07] and in figure 9(b) we see the equivalent but multiplied by the distance $rR(r)$.

two-pass implementation. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ then

$$\begin{aligned}
 F(u, v) &= G * f \\
 &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}} f(u-x, v-y) dx dy \\
 &= \frac{1}{2\pi\sigma^2} \int_{-\infty}^{\infty} e^{-\frac{y^2}{2\sigma^2}} \left(\int_{-\infty}^{\infty} e^{-\frac{x^2}{2\sigma^2}} f(u-x, v-y) dx \right) dy
 \end{aligned}$$

First horizontal convolution is performed and then vertical or vice versa. Thus six blurred irradiance textures are created from the original irradiance map using the Gaussians given by table 4. To further accelerate convolution the property that the convolution of two Gaussians results in a wider Gaussian is used

$$G(\sigma_1^2, r) * G(\sigma_2^2, r) = G(\sigma_1^2 + \sigma_2^2, r)$$

Initially, convolution is done by the narrowest Gaussian and then the process iteratively blurs the current already blurred irradiance texture. This is done such that results, at every iteration, are equal to convolution of the original texture by the subsequent wider Gaussian kernel.

The full accelerated algorithm, to achieve subsurface scattering, is outlined by the steps given here

1. Render shadow maps.

2. Render off-screen irradiance texture.
3. For each of the six Gaussian kernels
 - (a) Perform a horizontal blur pass.
 - (b) Perform a vertical blur pass.
4. Render mesh in 3D
 - (a) Sample each of the six convolved irradiance maps and combine linearly using the weights of table 4.
 - (b) Add reflected radiance (specular) for each light source.

The Gaussian kernels are defined relative to the surface of the mesh in millimeters. A simple way to resolve the surface to texel ratio is to use equations (43) and (44) since the magnitudes represent the amount of units at the surface per unit in the texture for the horizontal and vertical directions, respectively, in the texture. This tells us that $\frac{1}{\|\sigma_s\|}$ represents units in the texture per unit in direction σ_s at the surface and similar for σ_t . Another possible way to resolve the same ratios is to use the *ddx*, *ddy* pixel shader instructions. These allow you to compute derivatives, with respect to pixel units, directly in a pixel shader. Either way, the so called *stretch texture* is precomputed and used during the third step to distribute samples in the irradiance texture to perform convolution. In [dL07] the fixed number of seven samples is used for each convolution pass. This is the case for both the horizontal and for the vertical direction.

Additional problems in regards to performing the integration in 2D using the unwrap are: texels which are not covered by the unwrap and also that connected regions at the surface of the mesh can be disconnected in the unwrap. A possible solution given in [dL07] is to use an alpha mask to mark all texels used by the skin. This alpha mask is pre-blurred and used to linearly interpolate between a standard local light evaluation and the subsurface scattered result. However, this creates hard, dry-looking skin, near the boarder of the unwrap. Admittedly, the solution is not a good one and for this reason the results presented in [dL07] were done using a mesh with an unwrap which covers the entire irradiance texture map. This way any sampling done, during convolution, outside the unwrap (and the texture) will be clamped to the nearest valid irradiance sample. However, this can

only be considered an approximation since two triangles which are adjacent at the surface may be disconnected in the unwrap. Thus the correct place to continue sampling may be in an entirely different location in the irradiance texture. Nevertheless, in order to achieve better performance and simplicity, clamping to the boarder is a reasonable approximation for such an unwrap covering the entire texture.

We suggest a different strategy, which is more general, to solve the problem. A better solution is to dilate the initial irradiance texture before convolution. This process copies at every unused pixel the value of the closest valid pixel, i.e., pixels within the unwrap. However, this process is expensive but fortunately, since the unwrap is known prior to rendering, offsets to the closest valid pixel can be determined using a preprocess. With pregenerated offsets runtime dilation of the irradiance texture is trivially done in a single pass. Furthermore, by pregenerating the offsets it is even possible to take into account when triangles which are connected at the surface are disconnected in the unwrap. An unused pixel just outside the boarder of the unwrap can be assigned an offset to a pixel which belongs to the, at the surface, adjacent triangle. However, this will also complicate the preprocess step and does require more analysis which is beyond the scope of this paper. Nevertheless, a preprocess of the offsets, of a standard dilation method, is a simple and sensible solution for unwraps which do not cover the entire texture.

It was mentioned in the beginning of this section that solving the outer integration based on distances between texels only works with local subsurface scattering. This is because distances in the unwrap represent the length of a corresponding path, over the surface, between points. The transmitted distance through the medium, however, may be significantly shorter. In [dL07] this is referred to as global subsurface scattering.

Assuming a point light, let x_o be some occluded surface point and let x_i be the occluding surface point. The point x_i can be resolved by sampling the depth of the shadow map using x_o as input. This way it is possible during shading to resolve the depth of transmittance $d = \|x_o - x_i\|$. Using the convolved irradiance textures local subsurface scattering can be determined at both x_o and x_i . To maintain practicality, of unwrap-based integration, a method is needed to convert the already integrated local subsurface scattering at x_i into transmitted light at x_o . To achieve this the choice is made in [dL07] to approximate by determining the transmitted light for the point x' instead of x_o directly. As shown in figure 10 this is the point, at the bottom, immediately below x_i . Thus the transmitted distance is corrected using scalar

projection onto the surface normal at x_i which gives $d = (\vec{n}_i \bullet \vec{\omega}_i) \cdot \|x_o - x_i\|$. Here we have used that the unit direction from x_o to x_i is the direction towards the light $\vec{\omega}_i$.

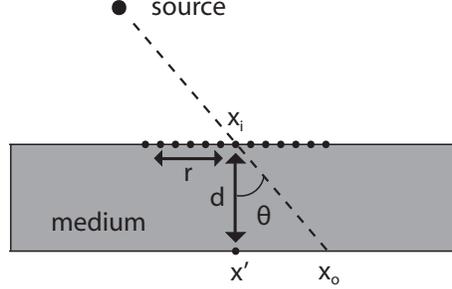


Figure 10: In this constellation x_o is occluded by x_i as seen from the light source. To determine approximately how much light is sent from x_i , through the medium, and to x_o the approximation is made to determine this at x' as opposed to x_o . This will allow us to solve the integration using the already convolved irradiance textures made for local subsurface scattering.

Let r represent the parameter for the local distance between x_i and surrounding sample points used during convolution. Conceptually, the distance from x' to the sample points, which surround x_i , is given by $r' = \sqrt{r^2 + d^2}$. By exploiting the squaring of r , in the Gaussian, the dependency on the transmitted distance d is factored out

$$\begin{aligned} R(r') &= \sum_{k=1}^6 w_k G(\sigma_k^2, \sqrt{r^2 + d^2}) \\ &= \sum_{k=1}^6 \left(e^{\frac{-d^2}{2\sigma_k^2}} w_k \right) G(\sigma_k^2, r) \end{aligned}$$

This allows us to reuse the six irradiance textures by sampling these at the texel location which corresponds to x_i . As we see the weights w_k used to sum up these contributions are scaled by $e^{\frac{-d^2}{2\sigma_k^2}}$ which will penalize, the contribution, as d increases. Furthermore, narrow Gaussians are penalized harder

than the subsequent wider Gaussians. Because of this the choice is made, in [dL07], to only sample the last three convolved irradiance textures to determine global subsurface scattering.

A significant disadvantage to this solution, to account for global subsurface scattering, is that until now the convolution only had to be performed once independently of the number of point lights in the scene. Using the suggested approach, for global subsurface scattering, one must compute convolved irradiance textures for each light separately. By default it is necessary to perform the convolution per unwrap. The added penalty of performing the convolution per light makes the algorithm inherently more expensive. An additional issue with this solution for global subsurface scattering is that, in situations where the texel locations corresponding to x_o and x_i are relatively close, global and local subsurface scattering will overlap. This issue will give an unwanted halo at the silhouette as seen from the light source since the overlap will cause contributions to be added twice. An approximate solution, used in [dL07], is to fade off the contribution received by global subsurface scattering when the texel locations corresponding to x_i and x_o are close to each other.

3.5 Diffuse Color Maps

The reflectance profile implicitly determines the color of a material. Consider once again the laser beam incident on a homogeneous semi-infinite plane-parallel medium. The percentage of transmitted light which will scatter around internally and finally make it back up to the surface is given by the total diffuse reflectance (37). The remainder $1 - R_{tot}$ is absorbed inside the medium. The value R_{tot} is wavelength specific and represents the color of the material. The single reflectance profile was derived based on the assumption that the medium is homogeneous and thus gives only one color per material. This presents a problem since a complex material such as skin presents both low and high frequency variation in diffuse color across the surface. This is where the assumption of homogeneity breaks down. To model such variation accurately we require a full participating media simulation which is significantly more expensive.

Another option is to use, instead of one reflectance profile per material, an entire texture map of generated reflectance profiles. This is impractical and furthermore the entire derivation is based on the assumption that σ_s and σ_a are constant per slab. In fact the reflectance profile is radially symmetric

due to this assumption (and others given in section 2.6). Since the reflectance profile is given as an analytical function in [JMLH01] the profile is evaluated using the local material parameters at x_i . Though it is not said so explicitly this implies a texture, mapped to the surface, of the parameters $\sigma_s(x_i)$ and $\sigma_t(x_i)$.

It would be impractical to apply this method to the multipole method, by [DJ05], since reevaluation of the reflectance profile per sample x_i is slow. For this reason only a single reflectance profile is used per material but with color variation stored in a texture mapped to the surface. This texture of colors represents the varying total diffuse reflectance which presents a problem since the reflectance profile also determines the material color. To solve this the color must be neutralized in one or the other. Given the comprehensive effort made in [DJ05] to derive a physically based reflectance profile it is decided to normalize the texture map instead. This is done in [DJ05] by scaling the channels such that the average intensity is the same across all three $R_{avg} = G_{avg} = B_{avg}$.

In the paper by [dL07] the alternative is chosen where the reflectance profile is normalized by dividing it by the total diffuse reflectance. Thus the values of the texture map are preserved and used instead to define local surface color. It is noted in this paper that modulation of the irradiance by this diffuse color texture can be done either before or after convolution by the reflectance profile. Modulating after convolution will preserve the high frequency detail of the texture but will also neglect to color bleed the skin tones. However, if the diffuse color texture came from a scan/photographs of real skin then natural color bleeding has already occurred and thus modulating after the convolution would be more appropriate. A compromise is suggested in [dL07] where a user-defined value $\alpha \in [0; 1]$ is used to control how much of the color is applied before convolution $col[x_i]^{1-\alpha}$ and how much is applied after $col[x_i]^\alpha$. Though there is no real physically based justification for this approach one might interpret, using $\alpha = \frac{1}{2}$, the process as an infinitesimal absorption layer at the surface. Thus absorption by this layer occurs once when light enters the medium and once again, after multiple scattering, as it exits the medium.

4 Conclusion

In this paper we have presented, in detail, the analysis of Henrik Wann Jensen et al. [JMLH01] and Craig Donner et al. [DJ05]. Furthermore, we provide a thorough walkthrough, and derivation, of the mathematical process used to derive their separate, but based on similar principles, BSSRDFs for sub-surface multi-scattered light. In the paper [DJ05] it was argued that the reflectance profile in [JMLH01], for human skin, is inaccurate because this is really composed of individual layers such as epidermis and dermis which have different scattering and absorption values. It is reported in [DJ05] that visual results become overly blurred and waxy as a consequence of this. For this reason a multi-layer model of skin is used to produce a more accurate reflectance profile based on a table of absorption and scattering values referenced to [Tuc00]. However, it is unclear to us whether or not the reflectance profile is wrong due to the single layer as claimed in [DJ05] or because the chosen absorption and scattering values reversed and given in [JMLH01] are wrong. We pointed out at the end of section 2.6 that the reflectance profile given by equation (4) in [JMLH01] has errors in it which were not remedied until later on. It seems plausible that the table of reversed parameters listed in [JMLH01] could be wrong depending on whether or not the errors in this equation were only in the paper or in the calculations as well. It would be interesting to pursue these questions in our future work.

We have implemented the multi dipole model and generated the reflectance profile using the table of absorption and scattering values given in [DJ05]. The result is shown in figure 6(b) and appears as a very thin spike followed by a base at zero which is the profile of the impulse function which is not desirable. To verify our implementation, to generate the reflectance profile, we turned to a freely available program by the name (MCML) Monte Carlo for Multi-Layered media by the authors Lihong Wang and Steven L. Jacques [WJ]. We gave MCML the table values listed in [DJ05] and though their implementation uses Monte Carlo simulation the resulting profiles are strikingly similar to the result produced using our implementation of the multi dipole method (see figures 6(b), 7(c) and 7(d)). Since the multi dipole model is well established in the optics and medical physics community it seems unlikely that the problem is rooted in this method. Furthermore, given the well matched results between MCML and our implementation of the multi dipole method (see section 2.8) we conclude that the problem is more likely due to errors in the table given in [DJ05]. The table is referenced

to [Tuc00] but according to the author Craig Donner the table will not be found there in this exact form.

An observation, which supports our suspicion, is that the values for σ_a and σ'_t in [DJ05] are significantly larger compared to those given for skin in [JMLH01] which could indicate a problem since both are defined relative to mm . The multipole and the dipole generated reflectance profile is fundamentally a sum of $e^{-\sigma_{tr} \cdot d}$ where d is a distance and $\sigma_{tr} = \sqrt{3\sigma_a\sigma'_t}$. Such a function will have a rapid fall-off when σ_{tr} is large which would indicate very little light transmits through the first and second layer in [DJ05]. Ultimately, we have decided to dismiss the validity of the table.

A more recent paper by Craig Donner et al. [DJ06b] proposes a two layer model where absorption and scattering values are produced by choosing inputs to control oil, melanin and hemoglobin. These will allow you to match different skin types such as Caucasian, African and Asian and various combinations of these. It seems unlikely that a designer would want to interface with such parameters directly but assuming the method works it might be a good way to create presets. For future work it would be interesting to investigate this approach further, and verify the procedure, by using the resulting scattering and absorption values in our multi dipole implementation.

In regards to designers it seems more realistic that their ideal interface would be to edit a spline curve in a way that is similar to the concept of tweaking a ramp for a post filter. Possibly aided by a list of presets for default reflectance profiles, to choose from, in the form of tweakable spline curves. Considering that the reflectance profile is known to consist of a spike and a broad base this might be easier than tweaking, for instance, scattering and absorption directly. As an example it was mentioned in section 3.4 that the reflectance profile used in the Matrix sequels [BL03] was made using a simple function (45) which is known to appear as a spike and a broad base.

At this point the better reflectance profile appears to be the one presented by Eugene d'Eon et al. [dL07] as a sum of Gaussians. This profile was reversed, by the author, using analysis of two similar images of a human head. One is an already existing image with subsurface scattering and the other is a BRDF evaluation. The exact details on how they extract such a reflectance profile are not public.

In section 3 we covered existing methods on accelerating the numerically solved integration. Jensen himself [JB02] uses an octree to adaptively reduce precision, for faster processing, at increasing distance which agrees with the profile of the reflectance function. The method consists of two steps where

the first pass evaluates the lighting, without subsurface scattering, for some dense distribution of sample points. The second pass computes the subsurface scattering, using the recorded irradiance, during the integration process. A very practical aspect of this algorithm is that the subsurface scattering part is done once independently of the applied light model and the amount of sources. This is a property that is worth preserving in an adaptation of the integration step to the GPU. For local subsurface scattering this is achieved using the method in [dL07].

Acknowledgments. This author would like to thank Swaminathan Narayanan and Manchor Ko for helpful discussions, Kenny Erleben and Trine Mikkelsen for constructive comments and proof reading.

References

- [BL03] George Borshukov and J. P. Lewis. Realistic human face rendering for "the matrix reloaded". In *SIGGRAPH '03: ACM SIGGRAPH 2003 Sketches & Applications*, pages 1–1, New York, NY, USA, 2003. ACM.
- [Cha50] Subrahmanyan Chandrasekhar. *Radiative Transfer*. Oxford Univ. Press, 1950.
- [CMZ97] Daniele Contini, Fabrizio Martelli, and Giovanni Zaccanti. Photon migration through a turbid slab described by a model based on diffusion approximation. i. theory. *Appl. Opt.*, 36(19):4587–4599, 1997.
- [DJ05] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Trans. Graph.*, 24(3):1032–1039, 2005.
- [DJ06a] Craig Donner and Henrik Wann Jensen. Rapid simulation of steady-state spatially resolved reflectance and transmittance profiles of multilayered turbid materials. *J. Opt. Soc. Am. A*, 23(6):1382–1390, jun 2006.

- [DJ06b] Craig Donner and Henrik Wann Jensen. A spectral bssrdf for shading human skin. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pages 409–418, June 2006.
- [dL07] Eugene d’Eon and David Luebke. *GPU Gems 3*, chapter 14, pages 293–347. Addison-Wesley Publishing, 2007.
- [DS03] Carsten Dachsbacher and Marc Stamminger. Translucent shadow maps. In *EGRW ’03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 197–201, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.
- [EVNT78] G. Eason, A. Veitch, R. Nisbet, and F. Turnbull. The theory of the backscattering of light by blood. *J. Physics*, 11:1463–1479, 1978.
- [FPW92] Thomas J. Farrell, Michael S. Patterson, and Brian Wilson. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Med. Phys.*, 19(4):879–888, 1992.
- [Ish78] Akira Ishimaru. *Wave Propagation and Scattering in Random Media*. Oxford University Press, 1978.
- [JB02] Henrik Wann Jensen and Juan Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graph.*, 21(3):576–581, 2002.
- [Jen01] Henrik Wann Jensen. *Realistic image synthesis using photon mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001.
- [JMLH01] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *SIGGRAPH ’01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 511–518, New York, NY, USA, 2001. ACM.
- [Kaj86] James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, 1986.

- [Mac48] T.M. MacRobert. *Spherical Harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications., 1948.
- [Mik08] Morten S. Mikkelsen. Simulation of wrinkled surfaces revisited. Master’s thesis, Department of Computer Science at the University of Copenhagen, 2008.
- [NRH⁺77] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. *National Bureau of Standards (US)*, 1977.
- [PCW89] M. Patterson, B. Chance, and B. C. Wilson. Time resolved reflectance and transmittance for the non-invasive measurement of tissue optical properties. *Appl. Opt.*, 28:2331–2336, 1989.
- [SH81] Robert Siegel and John R. Howell. *Thermal radiation heat transfer*. McGraw-Hill Education, 1981.
- [SKP09] Musawir A. Shah, Jaakko Konttinen, and Sumanta Pattanaik. Image-space subsurface scattering for interactive rendering of deformable translucent objects. *IEEE Comput. Graph. Appl.*, 29(1):66–78, 2009.
- [TS67] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57(9), 1967.
- [Tuc00] Valery Tuchin. *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*. SPIE Press, 2000.
- [Tur92] Greg Turk. Re-tiling polygonal surfaces. In *SIGGRAPH ’92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 55–64, New York, NY, USA, 1992. ACM.
- [Tur93] Greg Turk. *Graphics Gems I*, chapter 1, pages 24–28. Morgan Kaufmann Publishers, 1993.
- [WJ] Lihong Wang and Steven L. Jacques. (mcml) monte carlo for multi-layered media. Available at <http://omlc.ogi.edu/software/mc/>.