Radiative transfer in the cloudy atmosphere

B. Mayer^a

Deutsches Zentrum für Luft- und Raumfahrt (DLR), Oberpfaffenhofen, 82234 Wessling, Germany

Abstract. Radiative transfer in clouds is a challenging task, due to their high spatial and temporal variability which is unrivaled by any other atmospheric species. Clouds are among the main modulators of radiation along its path through the Earth's atmosphere. The cloud feedback is the largest source of uncertainty in current climate model predictions. Cloud observation from satellites, on a global scale, with appropriate temporal and spatial sampling is therefore one of the top aims of current Earth observation missions. In this chapter three-dimensional methods for radiative transfer in cloudy atmospheres are described, which allow to study cloud-radiation interaction at the level needed to better understand the fundamental details driving climate and to better exploit remote sensing algorithms. The Monte Carlo technique is introduced which allows to handle nearly arbitrarily complex atmospheric conditions. The accuracy of the method is discussed by comparison between different models and with observations. Finally, we show some examples and discuss under which conditions three-dimensional methods are actually needed and when commonly-used one-dimensional approximations are applicable. This chapter builds upon the excellent overview of one-dimensional radiative transfer in ERCA Volume 3 [1].

1 Introduction

Radiative transfer is among the most important topics in atmospheric research, if not the most important at all. Solar radiation drives weather and climate of the Earth. Without the sun as an external energy source the atmosphere would be a static and boring place. Atmospheric dynamics is driven by the differential heating by the sun, caused by the latitudinal gradient of the incident solar irradiance and the rotation of the Earth about its axis. On average over the whole globe and over time periods of years, the energy budget of the Earth is balanced. This balance is established by the absorption of solar radiation and the emission of thermal radiation by the Earth's surface and atmosphere: On average, 342 W/m^2 solar irradiance hit the Earth's atmosphere, out of which only 30% (102 W/m^2) are reflected back to space while the remaining 70% (240 W/m^2) are absorbed by the Earth's atmosphere and surface. The same amount is radiated back to space by thermal emission for which reason these 240 W/m^2 determine the average radiating temperature of the Earth. According to Stefan-Boltzmann's law,

$$E = \sigma T^4 \quad (1)$$

($\sigma = 5.67 \cdot 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ is the Stefan-Boltzmann constant; E is the emitted thermal irradiance) the temperature of the Earth is easily calculated to $T = 255 \text{ K}$ ¹. As the spatial and temporal distribution of thermally emitted radiation does not equal the distribution of the

^a e-mail: bernhard.mayer@dlr.de

¹ This is the effective radiating temperature of the Earth-atmosphere system as seen from space. To determine the surface temperature we need a detailed understanding of atmospheric structure, radiative transfer, lapse rates, and the greenhouse effect.

incoming solar irradiance, heat is transported in the atmosphere and oceans, on average from the equator towards the poles. This, together with the Coriolis force which further complicates circulation patterns, is what makes atmospheric science such an interesting subject and what makes numerical weather prediction such a challenging task, thus creating jobs for thousands of scientists. In addition to being the driving force of weather and climate, solar radiation controls atmospheric chemistry. The key reactions in atmospheric chemistry – the photolysis of NO_2 which is the only production mechanism of ozone in the troposphere, and the photolysis of O_3 which causes the formation of the OH radical – are driven by radiation and thus directly proportional to the available radiation in the ultraviolet spectral region. Finally, remote sensing is one of the most important sources of information about the atmosphere: Passive remote sensing methods use reflected solar radiation as well as emitted thermal and microwave radiation to learn something about atmospheric composition. Active remote sensing instruments such as lidar or radar provide their own source and detect the modification of the emitted laser or radar beam to study profiles of atmospheric constituents. To calculate trace gas concentrations or aerosol and cloud properties from the reflected or emitted signals, we need a detailed understanding of atmospheric radiative transfer.

Now, having understood that radiative transfer is an interesting and important topic in atmospheric research, we have a look at what modifies radiation on its way through the atmosphere: Radiation is affected by scattering and absorption by molecules, aerosol particles, and cloud droplets and particles. Among those, the interaction of radiation with clouds is the most complex and for many applications also the most relevant process. The complexity is due to the rapid spatial and temporal variability of clouds which is unrivaled by any other atmospheric species: E.g. air density shows only tiny fluctuations for which reason molecular scattering is relatively constant with time and location. Also, atmospheric trace gases and aerosols usually vary smoothly in the horizontal directions and with time, except close to sources like cities or power plants. In contrast, clouds form and dissipate rapidly. Their impact on the radiation is large and varies over orders of magnitude within a few meters at the cloud edges. Mainly for these reasons, clouds are the largest source of uncertainty in the prediction of future climate: according to the Summary for Policymakers of the IPCC Third Assessment Report [2] “Water vapour changes represent the largest feedback affecting climate sensitivity and are now better understood than in the [third assessment report]. Cloud feedbacks remain the largest source of uncertainty”. Concerning remote sensing, the retrieval of cloud properties is affected by unresolved or neglected inhomogeneities. It is long known that depending on the field-of-view of the satellite, the observed clouds have always more or less unresolved inhomogeneity which causes systematic uncertainties in the retrieved properties.

What is so special about radiative transfer in clouds? As we will briefly outline, radiative transfer is usually done under the assumption that the atmosphere is one-dimensional (1D): that it varies only in the vertical direction while horizontal variations are neglected, at least for radiative transfer applications. The reason was perfectly summarized by Anthony Davis [3]: “Homogeneous plane-parallel clouds may not exist in nature but they are the only ones for which we know how to solve the radiative transfer in a small amount of computer time.” This is certainly a good technical argument for not doing three-dimensional (3D) radiative transfer while scientifically there is little justification, especially in cloudy atmospheres. However, there is another good argument for some applications for not doing 3D radiative transfer, and this is the lack of suitable 3D input data. One example is radiative transfer in climate models: A climate model (general circulation model, GCM) grid box is typically 250 km in diameter, and the typical information available in a GCM is the relative humidity in each model layer which is then converted to cloud cover and liquid/ice water content using a parameterization. This is of course very poor information which does not really warrant an application of 3D radiative transfer. On the other hand, the resolution and detail of cloud-resolving models are increasing continuously, and so is the demand for a realistic treatment of cloud-radiation interaction. The same holds for remote sensing: The spatial resolution of instruments is increasing continuously, and here the only reason for not doing 3D radiative transfer is that 1D approximations are easier to use and cheaper. Here we see the largest potential for 3D radiative transfer: for active and passive remote sensing from space, from the ground, and from aircraft.

Table 1. Definition of the radiation quantities relevant for atmospheric radiative transfer.

Quantity	Symbol	Definition	Unit
Radiant Energy	Q		J
Radiant Power	Φ	$\frac{dQ}{dt}$	W
Irradiance	E	$\frac{dQ}{dA \cdot dt}$	W/m^2
Radiance	L	$\frac{dQ}{dA \cdot \cos \Theta \cdot d\Omega \cdot dt}$	$W/(m^2 \text{ sterad})$

2 Theory

Most of the underlying theory, including definition of quantities, the explanation of the radiative transfer equation, and the most popular 1D solutions, have been introduced in the overview of 1D radiative transfer in ERCA Volume 3 [1]. Here we quickly repeat these quantities which are needed in the following chapters. For most applications, knowledge of two radiation quantities is sufficient: irradiance and radiance, see Table 1. The irradiance E is the energy flux onto a flat surface. If nothing else is said explicitly, we usually mean the energy on a horizontally oriented flat surface, per unit time and unit area. This is the quantity which describes the flux on a specific receiver, either an instrument or e.g. the Earth's surface. Irradiance is therefore the relevant quantity to describe e.g. the input of solar energy into the atmosphere or onto the surface. As is immediately clear, the irradiance is not a function of the radiation alone, but depends also on the orientation of the receiver. If we want to fully quantify the radiation without sensor-specific information, we rather use the radiance L which is the radiant energy per unit time, unit area, and unit solid angle. Please note the difference: The radiance is a function of direction which is usually expressed by the polar angle θ and azimuth angle ϕ , see Fig. 1. Instead of the polar angle θ we often use its cosine $\mu = \cos \theta$. As nearly all equations in radiative transfer depend on $\cos \theta$ and $\sin \theta$ the use of μ facilitates the evaluation of various integrals. The area in the radiance definition is defined as perpendicular to the propagation direction (θ, ϕ) for which reason it is expressed as $dA \cdot \cos \Theta$ – meaning that the area element dA may have any orientation but we consider only its projection perpendicular to the propagation direction: Θ is the angle between propagation direction and area normal. If we know the radiance as a function of location (x, y, z) and direction (θ, ϕ) , the radiation is fully described². The aim of radiative transfer theory is therefore to calculate radiance as a function of location and direction.

The interaction of radiation with matter is described by the radiative transfer equation:

$$\frac{dL}{k_{\text{ext}} \cdot ds} = -L + \frac{\omega_0}{4\pi} \int_{4\pi} p(\Omega', \Omega) L(\Omega') d\Omega' + (1 - \omega_0) B_{\text{Planck}}(T) \quad (2)$$

k_{ext} is the volume extinction coefficient of the medium which describes how much of the radiant energy is removed from a collimated beam by scattering and absorption. ω_0 is the so-called single-scattering albedo which is simply the fraction of scattering and extinction coefficient: $\omega_0 = k_{\text{sca}}/k_{\text{ext}}$. In a conservative medium (only scattering, no absorption) the single scattering albedo is 1 while for a non-scattering, absorbing medium the single scattering albedo would be 0. The scattering phase function $p(\Omega', \Omega)$ quantifies the probability that radiation is scattered from

Fig. 1. Definition of *zenith angle* (or *polar angle*) θ and *azimuth angle* ϕ .

² In this text we do not discuss polarization; if polarization is to be included, we would need to provide the four components of the Stokes vector instead of the scalar radiance.

direction $\Omega' = (\theta', \phi')$ into direction $\Omega = (\theta, \phi)$. These three parameters completely describe the optical properties of the atmosphere. $B_{\text{Planck}}(T)$ is the Planck function which provides the radiance emitted by a blackbody of temperature T and thus quantifies the thermal source. If we specify the boundary conditions in addition (e.g. the reflecting properties of the surface and the solar radiation entering the atmosphere from above), we have all the information we need to calculate the radiance anywhere in the atmosphere. With this knowledge we can now understand the radiative transfer equation (2): it describes the change in the radiance passing a small volume element of length ds (left side) by three processes (right side): the removal of radiation from the incident beam by scattering and absorption (first term); the scattering of radiation from different directions Ω' into the considered direction Ω (second term); and the emission of radiation along the transect through the volume element (third term). For more details the reader is referred to [1] or any textbook on radiative transfer.

For most practical purposes in the Earth's atmosphere we don't need to treat thermal and solar radiative transfer at the same time. As Fig. 2 shows the solar irradiance arriving at the surface is relevant between $0.25 \mu\text{m}$ and about $4 \mu\text{m}$ while the thermal irradiance of the Earth's surface prevails between $4 \mu\text{m}$ and $100 \mu\text{m}$. With a few exceptions, like some satellite channels between 3 and $4 \mu\text{m}$ we may do either thermal or solar radiative transfer, but not both at the same time.

Now, (2) doesn't look that complicated, but consider that this is an integro-differential equation where the scattering integral on the right-hand side of the equation couples the radiance from all directions Ω' to the radiance in the desired direction Ω . Without scattering the solution is relatively straightforward, but in a scattering medium we need to solve (2) simultaneously for all directions. This is extremely nasty, and in fact, an analytical solution is not even available for the simplest possible scattering problem: A homogeneous, conservative medium with isotropic scattering. No surprise that people have been trying to simplify (2) before solving it! One of the most common simplifications is the 1D or plane-parallel approximation where one assumes that the optical properties k_{ext} , ω_0 , and $p(\Omega', \Omega)$ vary with height z only, but not in the horizontal directions, see Figure 3. This assumption only slightly modifies the shape of (2): ds is replaced by dz and all quantities depend only on z rather than on (x, y, z) . While this might not seem as a big step for mankind, it actually is one for a mathematician: For the solution of the 1D equation one may apply numerical techniques which are commonly used to solve integro-differential equations. One such solution leads to the so-called "discrete ordinate technique" where a Fourier and Legendre decomposition of the

Fig. 2. Radiative transfer calculation of downward solar and thermal irradiance at the Earth's surface.

Fig. 3. Definition of three-dimensional (3D) and one-dimensional (1D) radiative transfer.

Fig. 4. Typical examples of a 1D and a 3D atmosphere: (left) the ARM (Atmospheric Radiation Measurement) site in Oklahoma; (right) tropical convective clouds, seen from the Space Shuttle.

individual terms and an approximation of the integral by a sum over discrete angles (“discrete ordinates”) transforms the equation into a system of linear differential equations for which the solution is again straightforward (at least for a mathematician). A special case of the discrete-ordinate solution is the two-stream method which considers only radiation into two directions, upward and downward. Two-stream was already used more than 100 years ago [4] to calculate the radiation emitted by stars with a scattering atmosphere. Two-stream methods are very common, as they are employed in most climate models as well as in numerical weather prediction (NWP) models. They may give reasonably accurate results in little computational time.

Two points should be mentioned: (1) The vertical profile of optical properties needs to be considered which is usually done by stratifying the atmosphere in a number of layers; this basically multiplies the number of coupled differential equations to be solved by the number of layers; (2) if radiance is required for a number of wavelengths, or maybe integrated over the solar spectrum, the radiative transfer equation needs to be solved separately for all wavelengths³.

But back to the 1D plane parallel approximation: Is it really justified? Figure 4 shows two examples, of a 1D and a 3D atmosphere, respectively. Far away from sources, the distribution of trace species and aerosol particles in a cloudless atmosphere may be considered as 1D without doubt (left image). Whenever clouds are present, however, the 1D approximation does not look that convincing anymore. While there are parts of the right image where one might expect that a 1D approximation could give reasonable results (e.g. the cloudless fractions at some distance from the cloud, or right in the center of the clouds), the scene is dominated by cloud edges, reflection of radiation at cloud sides, and other effects. And there is no justification why the approximation should be reasonable, except that “we know how to solve the radiative transfer in a small amount of computer time.” Whether the 1D approximation is valid has to be tested in a case-by-case analysis and a lot of work has been done in that respect during recent years. For such tests we need a 3D model to solve the radiative transfer equation (2). As there is no simple closed numerical solution of the 3D radiative transfer equation, iterative solutions are applied, like the spherical harmonics discrete ordinate method, e.g. [5,6]. In this chapter we describe one of the most popular techniques for 3D radiative transfer: The Monte Carlo method. While considered as a brute-force technique by some, we think it is a very elegant method as it allows to consider all relevant atmospheric and surface processes without simplifying assumptions and allows the solution of radiative transfer in nearly arbitrarily complex 3D media. The nice thing about the Monte Carlo method is that you may use it to solve the radiative transfer equation without even knowing the equation! One basically only needs to understand the individual processes.

³ Inelastic (Raman) scattering plays a minor role in the Earth’s atmosphere but is relevant e.g. for the remote sensing of minor trace species like NO₂ – if we need to take inelastic scattering into account the individual wavelengths can no longer be treated separately and the solution is much more complicated.

3 The Monte Carlo technique

Online visualization of a MYSTIC simulation

Monte Carlo radiative transfer is a technique where individual photons are traced on their random paths through the atmosphere. One of the most versatile Monte Carlo codes for atmospheric radiative transfer is MYSTIC, the Monte Carlo code for the physically correct tracing of photons in cloudy atmospheres [7,8]. MYSTIC calculates solar and thermal radiance, irradiance, and actinic flux and has been used for remote sensing as well as for climate and photochemistry applications. It allows the definition of arbitrarily complex 3D clouds, inhomogeneous surface albedo as well as topography. The following description of the Monte Carlo technique follows the implementation in MYSTIC. The path of the photon is divided into 12 steps, marked by roman numbers I – XII. There are certainly other ways to do certain things for which the interested reader is referred to the literature. In particular, [9] gives a broad, 600 page overview of 3D radiative transfer and its application in atmospheric science.

Photon (solar / thermal)

I.

The generation of photons: Monte Carlo radiative transfer is about tracing photons from their source to their end. In the Earth's atmosphere, the source could be either the sun or, in the thermal or "longwave" spectral range, the photon is emitted by the surface or an atmospheric constituent (molecule, cloud droplet, ice crystal, aerosol particle). The solar source is described by the extraterrestrial irradiance of the sun, while thermal emission is defined by the emission coefficient and the temperature. The end of the photon could be either absorption by an atmospheric constituent, absorption by the surface, or when the photon leaves the atmosphere at top which we call here "top of atmosphere" (TOA)⁴.

Photon (solar / thermal)

II.

Each photon is assigned an initial location and direction. For a solar photon the location will be at top-of-atmosphere, at a random location in the x-y plane, assuming that the model domain is illuminated homogeneously by the sun. The direction is defined by the solar zenith and azimuth angles. Please note that we assume here that the solar zenith and azimuth angles are constant over the model domain which is valid for most practical purposes (except when the domain size is larger than about 100 km and the variation of the angles needs to be taken into account). In the thermal spectral range, the photon is emitted somewhere in the atmosphere. A random location is selected according to the emission coefficient and the temperature. How one determines a random location from a probability density will be explained later, but be assured that we can select photon start locations in a way that after many many photons the distribution of emitted photons will follow the product of emission coefficient and Planck function. As thermal emission has no preferred direction, a random direction is assigned to the photon.

⁴ In reality, the atmosphere has of course no top but for radiative transfer applications top-of-atmosphere is defined at typically 50–120 km because above the atmosphere is so thin that it does not significantly affect radiation anymore.

Photon travel in τ -space: Next, a step width is determined for the photon. The probability for survival p_{sur} of a photon along its straight path is determined by Lambert-Beer's law:

$$p_{\text{sur}}(\tau) = \exp(-\tau) \quad (3)$$

where the optical thickness τ is the integral of the volume extinction coefficient k_{ext} along the photon's path:

$$\tau = \int_0^s k_{\text{ext}}(s') ds'. \quad (4)$$

Let's for a moment stay in optical thickness space which is very convenient in radiative transfer: Knowing (3), we can calculate the probability $p_{\text{ext}} d\tau$ that a photon becomes extinct between τ and $\tau + d\tau$:

$$p_{\text{ext}} d\tau = p_{\text{sur}}(\tau) - p_{\text{sur}}(\tau + d\tau) = -\frac{dp_{\text{sur}}}{d\tau} d\tau \quad (5)$$

and we find more or less surprisingly that the probability that the photon survived to optical thickness τ is equivalent to the probability density that the photon becomes extinct between τ and $\tau + d\tau$:

$$p_{\text{ext}} = p_{\text{sur}} = \exp(-\tau). \quad (6)$$

We now have to draw a random optical thickness so that, when we repeat this step many many times, the histogram of the optical thickness equals the probability density p_{ext} . This can be done applying Von Neumanns Golden Rule of Sampling: First we calculate the cumulative probability density

$$P_{\text{ext}}(\tau) = \int_0^\tau p_{\text{ext}}(\tau') d\tau' = 1 - \exp(-\tau). \quad (7)$$

The advantage of the cumulative probability density is that we can now conclude that it is equally likely that τ falls into the interval $P(\tau) \in [0, 0.1]$ or $P(\tau) \in [0.3, 0.4]$, namely 10%. The respective τ intervals are calculated by inverting $P_{\text{ext}}(\tau)$:

$$\tau(P_{\text{ext}}) = P_{\text{ext}}^{-1}(P) = -\ln(1 - P). \quad (8)$$

Now imagine we have a regular dice with six faces and want to draw a random optical thickness. For that purpose we divide the cumulative probability density into six intervals:

Interval	$P_{\text{ext}}(\tau)$	τ
1	0/6–1/6	0.00000–0.18232
2	1/6–2/6	0.18232–0.40547
3	2/6–3/6	0.40547–0.69315
4	3/6–4/6	0.69315–1.09861
5	4/6–5/6	1.09861–1.79176
6	5/6–6/6	1.79176– ∞

If we would dial a 1, τ would lie between 0 and 0.18232, etc. This is of course a bit rough for our application. One would wish to have a dice with more faces – preferably an infinite

number. The concept is easily generalized to a dice which produces continuous random numbers ρ between 0 and 1. Following the formalism above, for each random number ρ the corresponding τ is easily calculated as

$$\tau = P_{\text{ext}}^{-1}(\rho) = -\ln(1 - \rho). \quad (9)$$

The figure illustrates how that works: The top plot shows 100,000 random numbers ρ equally distributed between 0 and 1, produced by a random number generator. When converted to optical thickness (middle plot) according to (9) the distribution is no longer even, but most photons become extinct at optical thickness between 0 and 2 while only few photons make it to significantly larger optical thicknesses. The histogram of the thus sampled optical thicknesses (lower plot) shows that after 100,000 photons the distribution closely follows Lambert-Beer's law (dashed line). With this methodology, we can sample from arbitrary (normalized) probability densities $p(x)$, by calculating the cumulative probability density $P(x) = \int_0^x p(x')dx'$ and inverting it:

$$x = P^{-1}(\rho) \quad (10)$$

where ρ is a random number between 0 and 1. This procedure was described in detail because it is the fundamental process in any Monte Carlo model: For our application we have a variety of probability density functions including the thermal emission of the photon in the atmosphere, the pathlength of the photon before it becomes extinct, the angular distribution of scattered light, or the directional reflectance of the surface. All of these can be handled by (10), but only very few of them can be solved analytically such as the step width in (9). For those which cannot, we need to setup a table which tabulates x as a function of ρ according to (10).

Photon travel in physical space: Now that we have determined a random optical thickness which the photon travels before it becomes extinct, we need to go back to physical space and determine the actual location of the photon. For that purpose, we need the 3D distribution of the extinction coefficient k_{ext} which is determined by the number and type of molecules, aerosol particles, and water and ice droplets. For Monte Carlo radiative transfer we mostly assume a rectangular grid where the optical properties are constant within each grid cell.

To determine the location, we simply integrate the extinction coefficient along the photon path until we have reached the optical thickness determined by (9):

$$\tau = \sum_{\text{cells}} l_i \cdot k_{\text{ext},i} \quad (11)$$

where l_i is the length which the photon travels in cell i – determined by geometry (the last l_i will of course usually not end at a cell boundary but somewhere in a cell). That way we determine the location (x, y, z) where the photon becomes extinct as well as the cell index.

Scattering and absorption: The cell where the photon has arrived determines the further fate of the photon. That is, the optical properties used for the following considerations are the optical properties of this particular grid cell. There are two ways for a photon to become “extinct”⁵: scattering⁵ and absorption: Which of both happens is determined by the single scattering albedo: With a probability of $\omega_0 = k_{\text{sca}}/k_{\text{ext}}$ the photon is scattered, with a probability of $1 - \omega_0 = k_{\text{abs}}/k_{\text{ext}}$ it is absorbed. Again we draw a random number ρ between 0 and 1: If $\rho \leq \omega_0$ we choose scattering, if $\rho > \omega_0$ we choose absorption. Absorption is the simpler process, as an absorbed photon is simply “gone” and we don’t need to trace it anymore. In our example the photon is scattered. To determine its new direction we need the scattering phase function. In MYSTIC we don’t store the combined scattering phase function but rather the scattering phase functions of the individual scattering species, that is, molecules, aerosol particles, water droplets, and ice particles. If we would store the combined phase function, we would need a different phase function for each grid cell because each grid cell contains an individual mixture of molecules, aerosol particles, water droplets, and ice particles. This is not practicable. Rather, we store one Rayleigh phase function (which is the same for all grid cells), a set of water cloud scattering phase functions for a set of different droplet sizes, etc. which reduces memory consumption and computational cost considerably. For that reason, we need to determine what actually scatters: with a probability of $k_{\text{sca},\text{mol}}/k_{\text{sca}}$ it will be molecular (Rayleigh) scattering, with a probability of $k_{\text{sca},\text{aer}}/k_{\text{sca}}$ it will be aerosol scattering etc. where $k_{\text{sca}} = k_{\text{sca},\text{mol}} + k_{\text{sca},\text{aer}} + k_{\text{sca},\text{wc}} + k_{\text{sca},\text{ic}}$ is the sum of all scattering coefficients. Only some scattering phase functions can be handled analytically, e.g. the Rayleigh phase function p_{Rayleigh} and the Henyey-Greenstein (HG) phase function p_{HG} which is often used as an approximation for aerosol and cloud:

$$p_{\text{Rayleigh}}(\theta) = \frac{3}{4} (1 + \cos^2 \theta) \quad (12)$$

$$p_{\text{HG}}(\theta) = \frac{1 - g^2}{(1 + g^2 - 2g \cos \theta)^{\frac{3}{2}}} \quad (13)$$

where the asymmetry parameter g is typically about 0.85 for water clouds. Integrating these phase functions to obtain the cumulative probability densities, and solving them for the scattering angle θ yields

$$\begin{aligned} \mu_{\text{Rayleigh}} = \cos \theta_{\text{Rayleigh}} &= u - \frac{1}{u} \\ \text{with } u &= \sqrt[3]{-q + \sqrt{1 + q^2}} \\ \text{and } q &= 4\rho - 2 \end{aligned} \quad (14)$$

for the Rayleigh scattering angle θ_{Rayleigh} and

⁵ The term “extinction” is a bit misleading, as the photon still “lives” after scattering; extinction refers only to the process of removing the photon from its initial direction.

$$\mu_{\text{HG}} = \cos \theta_{\text{HG}} = \frac{1}{2g} \left[1 + g^2 - \left(\frac{g^2 - 1}{2g\rho - g - 1} \right)^2 \right] \quad (15)$$

for the HG scattering angle θ_{HG} , where ρ is a random number between 0 and 1. For more complex scattering phase functions like a real Mie calculation for cloud droplets, tables are used which tabulate μ for a set of ρ .

In our example, the random process decided for aerosol scattering. The strange thing in step VII symbolizes the angular distribution of the scattering phase function, calculated by Mie theory for a spherical particle. As for all particles larger than the wavelength of the photon, a significant fraction of the scattering goes into the forward direction. When we determine a random direction according to (10) using a precalculated table, the most likely direction is therefore the forward direction. In our case, the photon decided to take a less likely side-ward step. To actually determine a new scattering direction involves a scattering polar angle θ sampled from the scattering phase function, a random scattering azimuth angle, and some elementary geometry to calculate the new direction from those. But now we are done with the first photon interaction: We have a photon with a new initial direction and in principle we now repeat the complete process infinitely, until (1) the photon is absorbed; (2) the photon leaves the model domain at top-of-atmosphere⁶; or (3) the photon hits the ground. In our example, the latter is the case.

Surface interaction: When a photon hits the surface, it may be absorbed or scattered, depending on the surface properties. If we assume a Lambertian surface reflection⁷, we only need to know the albedo which is the probability that the photon is reflected. In reality, most surfaces do not reflect isotropically and we need to apply the BRDF $\rho(\theta_i, \phi_i, \theta_o, \phi_o)$ (bi-directional reflectance distribution function) to determine the new direction (θ_o, ϕ_o) as a function of the direction of incidence (θ_i, ϕ_i) .

Here we introduce the concept of photon weighting which is a common methodology in Monte Carlo radiative transfer: So far we considered only processes where the photon as a whole survived or became extinct. In the case of surface reflection, a photon may either survive (if it is reflected) or die (if it is not). An alternative way is to always reflect the photon and determine a “photon weight” which is the probability that the photon survives. If the surface albedo is e.g. 0.2 we would thus always reflect the photon and multiply the photon weight by 0.2 (at the beginning of the journey, a start weight of 1 is assigned to each photon). One might think that this is a waste of computational time as we trace many photons with only

⁶ If a photon leaves the model domain through the side, we apply periodic boundary conditions: the photon re-enters at the opposite side; this ensures energy conservation and is appropriate for most applications, given that the model domain is large enough so that the process under consideration are not affected by edge effects.

⁷ Lambertian reflection is the most common and simplest assumption about surface reflection: the reflected radiance is equal in all directions; the surface looks equally bright from all directions.

small weights instead of tracing only those which are actually reflected. But the contrary is the case: At the end, when we count the photons, the result might converge faster because more photons per unit time contribute to the result: in our example we actually trace all photons to the very end instead of dropping 80% at the surface. The only exception is a surface albedo of 0 in which case we can safely stop the photon tracing because the photon will not make any contribution to anything anymore. Another advantage of this method is that it allows to consider complex BRDFs in a simple way: Instead of dealing with complex more-dimensional tables of cumulative BRDFs we simply reflect each photon into a random direction and multiply the photon weight with the BRDF $\rho(\theta_i, \phi_i, \theta_o, \phi_o)$ in that direction. After surface reflection, we proceed again from the very beginning (determine a step width in optical thickness space etc.). To determine a random direction seems straightforward: One might believe that we only need to draw a random azimuth angle ϕ between 0° and 360° and a random zenith angle θ between 0° and 90° . For the azimuth angle this is correct, but for the zenith angle we have to consider that a Lambertian surface emits a constant radiance in all directions; as the projected area $dA \cos \theta$ obviously decreases with increasing polar angle, so does the number of photons. This gives a probability density $p(\theta) = 2 \cos \theta$ which we need to integrate to obtain the cumulative probability density (the factor 2 is required for normalization):

$$P(\theta) = \int_0^\theta p(\theta') \sin \theta' d\theta' = 2 \cdot \int_0^\theta \cos \theta' \sin \theta' d\theta'. \quad (16)$$

The $\sin \theta'$ is required as this is actually an integral over solid angle $d\Omega' = \sin \theta' d\theta' d\phi'$ where the ϕ' -integration has been carried out before because the probability density is independent of azimuth. Here we stress once more the advantage of $\mu = \cos \theta$ over θ . With the substitution rule:

$$\int_{\theta_1}^{\theta_2} p(\theta) \sin \theta d\theta = - \int_{\mu_1}^{\mu_2} p^*(\cos \theta) \frac{d \cos \theta}{d\theta} d\theta = \int_{\mu_2}^{\mu_1} p^*(\mu) d\mu \quad (17)$$

with $p^*(\cos \theta) = p(\theta)$. In our case $p^*(\mu) = 2 \cdot \mu$ and

$$P^*(\mu) = 2 \cdot \int_0^\mu \mu' d\mu' = \mu^2 \quad (18)$$

from which we finally obtain the equation to calculate a randomly reflected direction from a random number $\rho \in [0, 1]$:

$$\mu = P^{*-1}(\rho) = \sqrt{\rho}. \quad (19)$$

As mentioned above, one of the big advantages of the Monte Carlo method is that nearly arbitrarily complex processes can be included, which are close to impossible to handle with any other numerical method. One example is a rough surface or topography. The figure shows how topography is treated in MYSTIC: A digital elevation model (DEM) with altitudes on a rectangular grid is used. In between, the surface is interpolated

bi-linearly. For a Lambertian reflection of the surface, we replace the vertical direction with the surface normal at the point where the photon hits the surface and proceed as above.

An alternative treatment of atmospheric absorption:

After having learned about photons weights, we go back to the treatment of absorption earlier in this section: Note that absorption in the atmosphere can also be handled by reducing the photon weight instead of killing the photon, and that is what we actually do in MYSTIC: In that case, the photon path is only affected by scattering while absorption is taken into account by the photon weight. Practically, we only need to replace k_{ext} by k_{sca} in (4)–(9) and to reduce the photon weight. As we know, the probability of photon absorption along the photon path follows again Lambert-Beer's law, and the photon weight between two scatterings will be reduced by

$$\exp(-\tau_{\text{abs}}) = \exp\left(-\sum_{\text{cells}} l_i \cdot k_{\text{abs},i}\right) \quad (20)$$

where the absorption optical thickness is calculated from the absorption coefficient k_{abs} similar to (11).

Radiance and irradiance: With this way of handling absorption and surface reflection, the only possible end of a photon is when the photon leaves the model domain at top-of-atmosphere. Now, the main reason for tracing photons in a Monte Carlo code is that we want to calculate one or more radiation quantities at one or more locations in the model domain. To do that, we have to trace many photons and count the photon whenever it hits the desired location. For example to calculate irradiance at the surface, we simply count the photon weights w_i of all photons hitting the surface:

$$E = E_0 \cdot \cos \theta_0 \cdot \frac{1}{N} \cdot \sum_{i=1}^{N_s} w_i \quad (21)$$

N is the total number of photons traced, N_s is the number of photons hitting the surface, E_0 is the irradiance entering top-of-atmosphere, and $\cos \theta_0$ is the cosine of the solar zenith angle which corrects for the oblique incidence of the sun. Of course this photon counting is not done at the end, but rather every time the photon hits the location where we want to know the irradiance. That way, one may calculate the vertical profile of irradiance by counting photons at each model level. In that case, each photon is “recycled” and may contribute to many irradiance values at different locations. Equally simple is the calculation of the horizontal distribution of irradiance: In that case we may e.g. define a horizontal grid with arbitrary spacing and count photons falling into each individual grid box (k, l) :

$$E = E_0 \cdot \cos \theta_0 \cdot \frac{1}{N} \cdot \frac{A}{A_{kl}} \cdot \sum_{i=1}^{N_{s,kl}} w_i. \quad (22)$$

Here $N_{s,kl}$ is the number of photons hitting grid box (k, l) , A_{kl} is the area of grid box (k, l) , and A is the total area of the model

domain. The area normalization factor considers that only a fraction A_{kl}/A of the photons entering the model domain hits grid box (k, l) . Please note, that the more sampling grid boxes, the more computationally expensive the calculation is: For a sample grid of 100×100 equal boxes, only each 10,000th photon will contribute to the irradiance of a specific grid box and to get only 1,000 photons in each grid box we need to trace at least 10,000,000 photons (remember that not all of them will make it to the location where we sample the photons).

The calculation of radiances is similar: In addition to sampling only those photons which hit the target level at the correct location, we sample only those photons which hit the target from the specified direction – of course we have to allow some angular interval, e.g. a cone with specified opening angle. One may easily guess that this is computationally very inefficient: Imagine e.g. that an opening half angle of 5° (which is large considering e.g. that the sun has an opening half angle of only 0.25° and we might be interested in the circum-solar radiation) corresponds to only 0.2% of the 4π solid angle of the sky and typically only a similar percentage of photon directions will fall into this cone while the remaining 99.8% will not contribute. Luckily, there are better ways to calculate radiances, see next section. When counting the photon, one must not forget to weight with $1/\cos\theta$, where θ is the polar angle under which the photon hits the target, in order to account for the slant incidence on the target area (see the definition of radiance in Table 1).

Photon statistics: Monte Carlo radiative transfer is a method where a number of photons is randomly traced through the atmosphere. As such, the result is inherently noisy. As one can easily imagine, the noise decreases with the number of photons. The calculation of the standard deviation of the result is straightforward: The calculation of any radiation quantity by a Monte Carlo model can be considered as a series of yes/no experiments where the photon either makes it into the result with a probability p or not with a probability $(1 - p)$ (forget photon weights for a moment). This results in a binomial distribution for which we can easily calculate average $\mu = N \cdot p$ and standard deviation $\sigma = \sqrt{p \cdot N \cdot (1 - p)}$ where N is the number of tries, see any textbook on statistics. With $N_s \approx p \cdot N$ (the number of photons sampled into the result) we obtain

$$\frac{\sigma}{\mu} \approx \sqrt{\frac{N - N_s}{NN_s}}. \quad (23)$$

This is only an approximation because we replaced the (unknown) probability p with the estimate N_s/N , but this approximation will be the better, the larger N is. If $p \ll 1$, (23) turns into the result for the Poisson distribution,

$$\frac{\sigma}{\mu} \approx \frac{1}{\sqrt{N_s}} \quad \text{if } N_s \ll N. \quad (24)$$

With that knowledge we can quickly estimate how many photons are required to get a desired accuracy. Let's go

Monte Carlo result as a function of sample size

back to the above example where we wanted to know the irradiance in a grid of 100×100 ground pixels. Assuming an atmospheric transmission of 0.5 the probability that a photon will be counted in a certain pixel is $p = 0.5 \cdot 1/10,000 = 5 \cdot 10^{-5}$. If we require a standard deviation of 1%, we need $N_s = 1/0.01^2 = 10,000$ photons according to (24). As $N_s = p \cdot N$, we need to trace a total of $N = 10,000/(5 \cdot 10^{-5}) = 200,000,000$ photons. This sounds like a lot and it actually is. Currently, a state-of-the-art code like MYSTIC can run about 100,000 photons per second on a standard INTEL processor for cloudless conditions (clouds increase computational time, as the number of scattering events increases). Such a calculation takes about half an hour. While 200,000,000 photons sounds like “a lot”, please consider what the sun does: For high sun, the irradiance at the surface reaches 1000 W/m^2 under cloudless conditions. If we roughly translate that into number of photons assuming an effective wavelength of $\lambda_{\text{eff}} = 500 \text{ nm}$ (the peak wavelength of the solar spectrum) and a photon energy of $hc/\lambda_{\text{eff}} = 3.9757 \cdot 10^{-19} \text{ J} = 2.5 \text{ eV}$, we find the incredible number of $4 \cdot 10^{21} \text{ photons} \cdot \text{m}^{-2} \cdot \text{s}^{-1}$ which the sun provides every second for every square meter. With a typical detector size of 1 cm^2 we would receive $4 \cdot 10^{17} \text{ photons} \cdot \text{s}^{-1}$ from which one can quickly derive the relative standard deviation if we averaged over a second: according to (24), $\sigma/\mu = 1/\sqrt{4 \cdot 10^{17}} = 1.6 \cdot 10^{-9}$. What a waste! Who wants to know the solar irradiance with an accuracy of 10^{-9} anyway? But seriously, while the main purpose of sunlight is not to be measured by scientists but to drive photosynthesis, the photon noise causes problems with spectral measurements e.g. in the ultraviolet spectral range where a detection limit of $10^{-6} \text{ W}/(\text{m}^2 \text{ nm})$ is desired. Considering the very limited area of a monochromator and the transmission of the necessary entrance optics we are in the range where the Monte Carlo code actually provides more photons per second for the result than the sun and the observation is actually limited by photon noise.

But back to model applications: If one is happy with an uncertainty of 10% instead of 1%, 1/100 of the above mentioned 200,000,000 photons would be sufficient as the uncertainty decreases with the square root of the number of photons according to (24). If, on the other hand, we don't need the spatial resolution but are only interested in the area-averaged irradiance with a standard deviation of 1% we would need only 20,000 photons which, again depending on the atmospheric conditions, is a matter of tenths of seconds to a few seconds on a modern processor. Hence, the common knowledge that Monte Carlo radiative transfer is computationally very expensive is not generally true: As long as only area-averaged quantities are concerned, Monte Carlo is not necessarily much slower than standard 1D algorithms like the discrete ordinate method. However, if more information is requested, e.g. horizontal distributions, then Monte Carlo is much more expensive. In that case, however, the higher computational cost comes with a gain in information.

Model accuracy: The first step to determine the accuracy of a model result is the comparison between different

radiative transfer codes. MYSTIC is operated in the framework of the freely available radiative transfer package libRadtran [10] which allows us to conveniently compare different solvers of the radiative transfer equation for identical input conditions. The figure shows such a comparison between MYSTIC and the discrete ordinate solver disort2 [11] for a 1D atmosphere. For cloudless as well as for overcast conditions both solvers agree within the numerical noise of the Monte Carlo simulations which in this case was smaller than 0.1%. As two completely different methods agree perfectly, we have some confidence that both methods as well as the code are correct, and we may speak of “exact” solutions of the radiative transfer equation. In the Intercomparison of 3D Radiation Codes (I3RC), a detailed comparison for a variety of 3D atmospheres showed an agreement between a group of models including MYSTIC to about 1% or better [12]. The validation of a 3D radiative transfer model with observations, however, is not at all easy: We would require radiation observations as well as measurements of the atmospheric conditions with high detail and high accuracy. Well-characterized 3D atmospheric conditions are scarce, and in particular clouds change so rapidly in space and time that the observation of the 3D structure of a cloud at a given moment is close to impossible. While for clouds a real validation of 3D models is still open, we used other 3D cases for the validation of MYSTIC, in particular the simulation of a total solar eclipse where measurements and simulations agreed perfectly well [8,13]; or the effect of topography on the radiation in the METCRAX experiment [14].

3.1 Some tricks

In the previous section a basic Monte Carlo code to calculate radiance and irradiance in the atmosphere has been described. Although this allows to calculate correct results, the described method is not necessarily efficient. Here we briefly describe some methods which speed up the calculation by improving the photon statistics. All of those used are in MYSTIC.

3.1.1 Local or directional estimate techniques

As we already mentioned above, the calculation of radiance by sampling photons into a cone centered around the desired direction is rather time-consuming because we trace many photons which do not contribute to the result at all since they do not fall into the cone. In the example only photon 1 makes it into the cone while photon 2 fails. In reality, the probability that the photons hit the cone is very small. A second disadvantage is that we do not really calculate the desired radiance for a given direction, but rather the radiance averaged over the solid angle interval formed by the cone. We can avoid that by making the cone smaller, but reducing the opening angle by a factor of two means that we sample four times less photons which also increases the computational time by a factor of four.

An alternative method is the so-called local or directional estimate technique⁸ which calculates at each scattering point the probability that the photon is scattered into the direction of the sensor, rather than into the actual scattering direction (the actual photon path is not affected by this technique!). The probability for scattering into the direction of the sensor is given by the phase function, and of course one needs to take into account the extinction between scattering and detector. At each scattering point add a weight w_v to the radiance:

$$w_v = w_0 \cdot p(\Theta_p) \cdot \exp(-\tau_{\text{ext}}) / \cos \Theta_d \quad (25)$$

where w_0 is the photon weight. Θ_p is the angle between photon direction (before scattering) and the radiance direction. In consequence, the phase function $p(\Theta_p)$ gives the probability that the photon is actually scattered into the direction of the detector. We can think of a virtual photon created at each scattering point. Of course, this virtual photon suffers extinction along its path from the scattering point to the detector which is considered by the Lambert-Beer term $\exp(-\tau_{\text{ext}})$ where τ_{ext} is the extinction, integrated over all cells on the way from the scattering to the detector. Finally, Θ_d is the angle between the direction of the virtual photon and the detector, to account for the slant area in the definition of the radiance in Table 1. The same needs to be done when a photon hits the surface, using the albedo or BRDF for the photon weight instead of the scattering phase function. After tracing many photons, the radiance is then the sum of all virtual photon weights w_v . That way, each photon contributes not only once but several times.

The local or directional estimate technique can be proven by formally integrating the radiative transfer equation into a Von-Neumann-Series, see e.g. [9]. However, we can also easily understand the technique: We actually sample the radiance at each point in the model domain where there is a potential that the photon might change its direction into the sensor direction either by scattering or surface reflection and multiply the photon weight by the probability that this actually happens.

If we look carefully at (25) we find that the photon weights w_v may now span a considerable range of values, essentially from the maximum of the scattering phase function $p(\Theta_p)$ – usually the forward peak – to more or less 0. In a “normal” (non-directional estimate) calculation, the weights are usually between 0 and 1 (except for the BRDF weighting) while we now may get contributions much larger than 1. A large contribution sounds nice at the first glance but unfortunately these large contributions occur only rarely which implies that, when they occur, we get a spike in the result which no longer smoothly converges towards the average but in a series of steps. The figure shows an extreme but realistic example where we try to calculate TOA nadir radiance for an optically thick water cloud at a wavelength of 350 nm. The top plot shows the scattering phase function $p(\theta)$ which spans six orders of magnitude: The

⁸ In the literature the term “local estimate” is often used while the technique should better be called “directional estimate” because we calculate the contribution into a certain direction.

convergence of the result is rather slow because some photons have a weight which is several orders of magnitude larger than that of most others. The lower plot shows the result as a function of the number of photons. Completely different from what we saw above, the result (grey line) does not converge with $1/\sqrt{N}$ but it looks like we are not close to convergence even after 10^7 photons where we would have expected the relative uncertainty of the result to be in the order of $1/\sqrt{10^7} \approx 3 \cdot 10^{-4}$. Rather we find noise and systematic jumps in the order of several %. This is due to the fact that (23) has been derived under the assumption of equal weights for all photons, while we now have weights spanning many orders of magnitude.

Different techniques have been applied to reduce this noise or so-called “hot spots”: A common technique is to simplify the phase function by removing the strong forward peak and assuming that the removed fraction is not scattered but simply goes on un-scattered – this technique is known as peak-truncation or delta-scaling. While these are approximations rather than exact solutions, there are so-called “variance reduction” techniques which may improve convergence considerably without introducing a bias. A description of these techniques would go far beyond the scope and size of this chapter. The main idea is that whenever we have contributions to the result which have large weights but occur only rarely, we force more of these contributions and reduce their weight accordingly. Remember e.g. the treatment of the BRDF in section 3: With the described method, photons are reflected in random directions and their weight is scaled with the BRDF. For a nearly flat water surface, most photons will be assigned a very small weight and contribute little to the result; only those which are accidentally reflected close to the specular reflection direction will contribute considerably. This results in bad statistics because every 1000th photon or so causes a spike in the result and the convergence is slow. There is a solution to that: Instead of waiting for every 1000th photon, we force more photons into the specular reflection direction and reduce the photon weight accordingly. That way we force more spikes while at the same time reducing their magnitude which may improve statistics considerably. Similar techniques can be applied to atmospheric scattering. The black line shows the result when all variance reduction techniques in MYSTIC are switched on. The convergence is much better and we get a reliable result with a reasonable number of photons.

3.1.2 Backward Monte Carlo

According to the reciprocity principle first formulated by Hermann von Helmholtz, the path of light from A to B is reversible. An exhaustive discussion of this principle is presented e.g. by [15] from which we copy the following translation of Helmholtz’ principle, as formulated in his Treatise on Physiological Optics, Volume 1, published in 1856: *Suppose light proceeds by any path whatever from a point A to another point B, undergoing any number of reflections or*

"plain" MYSTIC

delta-scaling

variance reduction

Radiance simulations, all with identical computational time; hot spots appear in black.

Hermann v. Helmholtz
(1821-1894)

refractions en route. Consider a pair of rectangular planes a_1 and a_2 whose line of intersection is along the initial path of the ray at A ; and another pair of rectangular planes b_1 and b_2 intersecting along the path of the ray when it comes to B . The components of the vibrations of the aether particles in these two pairs of planes may be imagined. Now suppose that a certain amount of light J leaving the point A in the given direction is polarised in the plane a_1 , and that of this light the amount K arrives at the point B polarised in the plane b_1 ; then it can be proved that, when the light returns over the same path, and the quantity of light J polarised in the plane b_1 proceeds from the point B , the amount of this light that arrives at the point A polarised in the plane a_1 will be equal to K . Apparently the above proposition is true no matter what happens to the light in the way of single or double refraction, reflection, absorption, ordinary dispersion, and diffraction, provided that there is no change of its refrangibility, and provided it does not traverse any magnetic medium that affects the position of the plane of polarisation, as Faraday found to be the case.

What it basically means for our purpose is that it doesn't matter if we trace the photons from the source to the detector or the other way round. Tracing photons backward can be an enormous advantage compared to forward tracing whenever the irradiance or radiance is only needed at a certain location and not everywhere in the model domain. Think as an example of a detector in the center of our model domain: In standard Monte Carlo (as described above) we would illuminate the whole model domain at top of atmosphere and sample photons in a small area centered around the detector. Most photons would not contribute to the result as they would not hit our sample area. Going backward, however, all photons would be started at the detector and hence contribute to the result. The remaining problem is that the backward photons only contribute to the result if they hit the source, e.g. the sun. While we could in principle again sample photons going into a small cone centered around the sun, the directional estimate technique helps again to speed up the calculation considerably. If we want to calculate radiance at a point location, we would start the backward photons at the detector into the desired radiance direction. If we are interested in irradiance, we would start the photons in all directions, equivalent to the Lambertian reflection of photons described above⁹. In the latter case, the problem of calculating the irradiance at a point location caused by the parallel illumination of the top of the model domain is replaced by the equivalent problem of calculating the radiance into the direction of the sun, caused by a Lambertian point source.

⁹ Please note that this way we only calculate the diffuse component; the direct irradiance is easily calculated by sampling the extinction along the direct path from the detector to the source and applying Lambert-Beer's law which is basically a directional estimate before we even start the photon.

Summary of MYSTIC specifics

At the end of this section, we summarize the described method by showing how it is implemented in the MYSTIC Monte Carlo code [7,8]. MYSTIC is a forward/backward Monte Carlo code which allows the realistic treatment of inhomogeneous clouds, surface albedo and BRDF, and topography. The model atmosphere consists of a 1D background of molecules and aerosol particles and a 3D grid of cloudy cells. A schematic diagram of the algorithm is shown in Figure 5. The involved physics is simulated as closely as possible on the basis of the input atmosphere, without any further simplifying assumptions. All processes involving random numbers are marked with a double-frame. As the random generator is the heart of a Monte Carlo model, special care should be used in the choice of the random number generator. Remember that nowadays we run up to 10^9 or more photons which means that we need in the order of 10^{10} – 10^{11} random numbers per Monte Carlo calculation. Requirements are therefore that the random number generator has a long period, produces real random numbers without correlations or periodicity, and is fast. In MYSTIC, the random number generator was already replaced twice, as the used one was found not to be “random enough”. Currently we use the MT19937 random number generator [16]. According to the documentation, MT19937 “is a variant of the twisted generalized feedback shift-register algorithm, and is known as the Mersenne Twister generator. It has a Mersenne prime period of $2^{19937} - 1$ (about 10^{6000}) and is equi-distributed in 623 dimensions. It has passed the diehard statistical tests.” This sounds not only impressive but also gives us the good feeling that we should have no problems even if we run as many photons per second as the sun does.

Photons are traced from scattering to scattering. Absorption is considered implicitly, by reducing the photon weight by the Lambert-Beer factor $\exp(-\tau_{\text{abs}})$. Interaction with the surface is either treated by actually absorbing the photon if it is not reflected, or by always reflecting the photon in a random direction and reducing the photon weight by multiplying it with the Lambertian albedo or BRDF. A number of optimizations has been introduced to speed up computational time and to reduce noise: Radiances may be calculated either by cone-sampling or by the local estimate technique. Variance reduction techniques are used to reduce the inherent intermittency of directional estimate radiances. A backward Monte Carlo mode has been introduced which has proven extremely useful whenever results are not required everywhere in the model domain but only at few selected locations. Backward Monte Carlo is also the method of choice for calculations in the thermal spectral range where the computational time may be reduced by many orders of magnitude especially in spectral regions where the atmosphere is optically very dense.

MYSTIC is operated as one solver of the freely available libRadtran radiative transfer package [10], see <http://www.libradtran.org> (MYSTIC is currently not part of the free package). libRadtran translates atmospheric properties like pressure, temperature, ozone concentration,

Fig. 5. Schematic overview of the basic MYSTIC model for surface irradiance calculation, without directional estimate or variance reduction techniques. The double-framed boxes include a random number.

Fig. 6. Cubic clouds (left) and the transmittance of the clouds at the surface (right), calculated with 3D radiative transfer and approximated with the independent-column (or independent-pixel) and the plane-parallel approximations.

or liquid water content to optical properties, and processes the results to obtain calibrated spectra, weighted doses, or photolysis frequencies.

4 Applications

4.1 Examples from the cubistic period

Monte Carlo radiative transfer is not at all a new subject. Various studies have been done in the 1970s, e.g. [17], but in the good old days computers were slow, and studies were restricted to a few simple cases, like e.g. cubic clouds, see Figure 6. While these might seem unrealistic and over-simplified, they are very well suited to demonstrate the main 3D radiative transfer effects. Therefore we show two examples for cubic clouds, in particular the transmitted irradiance below a cubic cloud field, and the reflected radiance above a cubic cloud field.

Figure 6 illustrates the 3D calculation of the transmittance (defined as E/E_0 where E_0 is the extraterrestrial irradiance) as well as two widely used approximations: In the *plane-parallel approximation*, we simply assume that the cloud is plane-parallel (homogeneous in the horizontal directions); for that purpose, the cloud optical or microphysical properties are averaged horizontally; in our application that means that we average 50% clouds with $\tau = 50$ and 50% cloudless sky and obtain $\tau_{\text{avg}} = 25$. The solid line shows the 1D plane-parallel calculation of the transmittance as a function of optical thickness from which we can directly obtain the transmittance of the averaged cloud $T_{\text{PPA}} = 0.24$. As the diagram already illustrates, the relationship between optical thickness and transmittance is non-linear for which reason we certainly introduce an error if we calculate the radiation after averaging the cloud field. From the shape of the curve we can infer that we will usually underestimate the actual transmittance for which reason the uncertainty is called the *plane-parallel bias*: In the extreme case of a very large optical thickness, the cloudy part would have a transmittance of about 0 while the cloudless part has a transmittance of 1 which would give an averaged transmittance of 0.5. In plane-parallel approximation we would, however, get a transmittance of about 0 because the averaged optical thickness $\tau_{\text{avg}} = \tau/2$ is still very large.

In contrast, the so-called independent pixel (or independent column) approximation tries to avoid the plane-parallel bias, by not averaging the cloud properties but instead doing separate calculations for the cloudy and cloudless fraction (assuming plane-parallel homogeneous clouds for each part), and averaging both:

$$T_{\text{IPA}} = c \cdot T_{\text{cloudy}} + (1 - c) \cdot T_{\text{cloudless}} \quad (26)$$

where c is the cloud fraction (0.5 in our example). The dashed line in Figure 6 illustrates how the independent column approximation calculates the weighted average. For many cases, the

independent column approximation provides reasonably accurate results, but in particular for broken clouds it may fail as our example illustrates: in our example, the solar zenith angle is 50° which implies that the shadow of the cloud falls onto the surface between the clouds – in other words, the cloudless fraction is influenced by the clouds and vice versa and the columns are not really independent. In this particular case, no direct radiation at all hits the surface and one might guess that this has strong consequences for the radiation. Figure 6 also shows the result of the 3D Monte Carlo calculation – there are two data points because the result slightly depends on the incidence azimuth of the sun, and these two points were calculated for incidence parallel and diagonal to the cube faces. In this particular example the true result lies between the plane-parallel and the independent column approximation, but this is not a general rule: in principle it may lie anywhere above, below, or between the PPA and IPA results. The PPA and IPA are the most common methods for radiative transfer in cloudy atmospheres. They are used in all climate models where the PPA has meanwhile been replaced by the IPA to avoid the plane-parallel bias, by specifying a cloud fraction in addition to the optical thickness in each model layer; but there is a long way to go before there will be a real application of 3D radiative transfer in a climate model: First, we don't have the computational power to afford 3D radiative transfer in a climate model; second, even if we did, there is no easy way to translate the climate model output to a 3D distribution of water and ice, unless the climate model has a high enough resolution to resolve clouds: First steps into that direction have been made with the so-called super-parameterization.

For remote sensing we face similar problems: Generally we don't have any idea about the variability of the radiation within the sensor's field-of-view which varies between several 100 meters for satellites aimed at cloud detection (e.g. TERRA/MODIS, NOAA/AVHRR) up to several 10s or 100s of km for satellites aimed at trace gas observations (e.g. AURA/OMI, ENVISAT/SCIAMACHY). Within the field-of-view (or "satellite pixel") we usually assume that the cloud is constant (plane-parallel assumption) as we have little idea about the unresolved sub-pixel variability¹⁰. For the retrieval of atmospheric properties we also assume that we can treat each satellite pixel separately which simplifies retrieval techniques tremendously (independent pixel assumption). For the plane-parallel assumption we already know that the result will be biased, but let's use our cubic clouds to study the effect of the independent pixel assumption: Figure 7 shows the setup: We use our cloud field from above, but this time with an optical thickness of 20. Assume our cloud field is located at the equator and we do a calculation for spring equinox, where the sun travels along the horizontal line and reaches the overhead position at local noon. In the following we show the result for two satellite viewing geometries, one for the sensor looking straight down (nadir) and the second for an inclined view along the arrow, with a viewing zenith angle of 30° . Figure 8 shows a series of calculations for the second viewing geometry for the sun at various positions. We clearly see that the reflected radiance is affected by cloud top as well as cloud sides and the independent column approximation would give wrong results here. A quantitative analysis of the results is presented in Figure 9 – each of the stars corresponds to an individual radiance calculation as shown in Figure 8. The independent column approximation gives the highest radiance at noon, which is what one would expect. The actual maxima, however, occur at a solar zenith angle of 50° which was already obvious in Figure 8. This behaviour is caused by two effects acting in the same direction: First, the radiation reflected by the cubes is reduced as part of the radiation entering the top of the cube leaks out from the sides and is not reflected. For this reason, the radiance at local noon is reduced compared to the independent pixel calculation. Second, a much larger fraction of the radiation is intercepted by the cloud for lower incidence angles because the cloud sides intercept additional radiation; this is the reason for the enhanced reflectivity for low incidence

Fig. 7. Geometry of the problem.

¹⁰ Some information about sub-pixel variability can be obtained from spectral information or from simpler sensors with higher-resolution on the same satellite platform.

Fig. 8. A series of 3D calculations for different positions of the sun on its daily course.

Fig. 9. Reflected radiance as a function of the solar zenith angle for nadir viewing geometry (left) and an 30° inclined viewing geometry (right).

angles. In combination, we find deviations up to a factor of 2 between the true radiance and the independent pixel approximation. Most realistic clouds are less pathological than these cubic ones, but the cubes are well suited to show the general behaviour.

4.2 Application to remote sensing

Remote sensing of clouds with passive and active remote sensing instruments is currently of highest scientific interest, as only satellites allow to observe clouds over larger spatial and temporal scales, required for better understanding their role in climate, see introduction of this chapter. The validation of remotely sensed cloud properties is extremely difficult, as it is close to impossible to provide independent in-situ observations of cloud properties at the scale of the satellite instrument. Imagine that e.g. the validation of an optical thickness derived from an instrument such as SEVIRI on Meteosat Second Generation would mean that we need to measure the average optical thickness of the cloud over a $3 \times 3 \text{ km}^2$ area (the pixel size). From a ground station this is difficult to achieve because all remote sensing instruments like cloud radar, microwave radiometers, lidars etc. either provide only indirect information about the cloud or on a limited spatial scale (e.g. directly overhead the instrument). From aircraft observations of liquid water content and droplet radii the optical thickness can in principle be calculated, but that would imply that the aircraft samples the complete volume ($3 \text{ km} \times 3 \text{ km} \times \text{cloud height}$) preferably in a few minutes or less, as the cloud changes not only spatially but also with time. For that reason, we suggest artificial satellite images as an alternative approach to test remote

Fig. 10. Simulated satellite observation at 600 nm, based on the output of a numerical weather prediction model, COSMO-DE; high-resolution calculation with a pixel size of 0.56 km; (left) 3D calculation; (right) independent pixel approximation.

sensing algorithms. The idea is to use realistic and complete cloud information, obtained, e.g., from synthesized observations or from cloud-resolving models, use these as input to radiative transfer simulations to generate artificial satellite images, apply the remote sensing algorithm, and to compare the retrieved cloud properties with the ones used as input. That way, we have a consistent model-generated approach for the validation of remote sensing products. In order to obtain meaningful results, we need highly realistic cloud properties, preferably covering large areas as remote sensing algorithms also make use of cloud structure and behave differently over different surfaces. Figure 10 shows an example of such a simulation, based on the output of the numerical weather prediction (NWP) model COSMO-DE of the German Weather Service. The pixel size is 0.56 km in this case (created by statistical downscaling of the NWP model) and the area covered is about $100 \times 100 \text{ km}^2$. For the simulation we assumed a solar zenith angle of 65° with the sun in the West-South-West (azimuth 66°). The sensor zenith angle was $\text{SZA} = 59^\circ$ with the sensor approximately in the South. This is a typical afternoon viewing geometry for Meteosat Second Generation over central Europe. For comparison, the left plot shows the 3D calculation while the right plot is the respective independent pixel approximation. The cloud top structure becomes pretty obvious in the 3D simulation, as higher clouds cast shadows on lower clouds, and the sunlit faces of the clouds appear very bright. In contrast, the independent pixel approximation appears rather homogeneous and the only signs of structure are the gaps in the cloud where the ground is visible. The sunlit portions in the 3D simulation are brighter than any 1D approximation could explain. That implies that if we apply a standard remote sensing algorithm to that image, the sunlit portions would lie outside the range of pixels covered by the 1D lookup-table used for the retrieval and cause a large overestimation of the optical thickness. The opposite is true for the cloud shadows where the optical thickness would be underestimated. Due to the non-linearity of the relationship of optical thickness and radiance, both effects do not cancel, though. In consequence we need to consider 3D effects when we simulate artificial satellite images for cloud retrieval validation purposes. While this is clear for such high-resolution images, the question is if this still plays a role for lower-resolution instruments like SEVIRI on Meteosat Second Generation. Figure 11 answers this question: Here the radiances from Figure 10 were averaged over 5×5 high-resolution pixels in order to obtain the resolution of SEVIRI, and even here we see strong 3D effects.

Fig. 11. Same as Figure 10 but resolution degraded to 2.8 km by averaging 5×5 pixels.

5 Summary and further reading

In this chapter we tried to explain atmospheric radiative transfer processes by illustrating the Monte Carlo method – the most flexible way to solve the radiative transfer equation, which allows to consider nearly arbitrarily complex 3D atmospheres and also complex surfaces with inhomogeneous BRDF and even topography. The Monte Carlo method is computationally more expensive than a numerical 1D solution such as the discrete ordinate method. Very expensive simulations, however, are usually connected with a gain in information like the spatially resolved radiance reflected by a 3D cloud field which is not available at all from a 1D simulation. While the Monte Carlo method is extremely versatile, it is not necessarily the method of choice for all applications. For 95% of the radiative transfer applications 1D methods are to be preferred. 3D radiative transfer is required for applications like the high-resolution remote sensing of clouds, the effect of topography on the radiation budget, the cloud-radiation interaction in cloud-resolving models, or the passive remote sensing of inhomogeneous surfaces, to name but a few. While 3D radiative transfer has been mainly used to quantify the uncertainty of 1D approximations historically, current applications also concentrate on directly applying 3D radiative transfer, e.g., in the high-resolution remote sensing of clouds [18,19].

Playing around with a Monte Carlo model helps students get real insight into the processes which affect solar and thermal radiation on their complex path through the atmosphere. At our institute we offer a workshop each year where students develop their own Monte Carlo code over the course of a week. Within that time frame it is perfectly possible to develop a simple 1D Monte Carlo code which correctly calculates transmission in an atmosphere containing molecules and clouds. Every student involved in radiative transfer is therefore encouraged to sacrifice some time to develop a simple Monte Carlo code to get acquainted with the subject. After that, however, we also encourage the student to check if a ready-made tool is already available for the specific purpose because models like MYSTIC and SHDOM [5] have been continuously developed over many years to detect and remove errors, to add many features, and to optimize the code for speed and memory. A good overview of what's available to date is e.g. given by the two large intercomparisons of 3D radiative transfer codes, I3RC (<http://i3rc.gsfc.nasa.gov/>) and RAMI (<http://rami-benchmark.jrc.it/>). While the former focusses on 3D atmospheres, the latter addresses mainly the interaction with complex surfaces.

For those who want to learn more about the subject we recommend the textbook by Marshak and Davis [9] which gives a broad overview over 3D models and approximations as well as applications. Here one may also find the exact theory which was only briefly touched in this introductory chapter. Those who need the Monte Carlo method for new and challenging subjects will also find a wealth of information in the pioneering book by Marchuk [20].

The careful reader will have noticed that the author used the term “we” repeatedly in the text. This work is based on more than 10 years of Monte Carlo experience of the author and his group. In particular, Dr. Tobias Zinner, Dr. Claudia Emde, and Dr. Robert Buras heavily contributed to the development of MYSTIC and the knowledge briefly summarized in this chapter. The chapter is also based on several years experience with a workshop organized by the author, together first with Dr. Ronald Scheirer and then with Dr. Claudia Emde, where students develop a fully-working Monte Carlo code within one week. Finally, I want to thank Dr. Tobias Zinner, Tanja Reize, Dr. Robert Buras, and Dr. Claudia Emde for reading the manuscript and for many helpful comments and improvements.

References

1. B. Pinty, M.M. Verstraete, ERCA **3**, 67 (1998)
2. IPCC, Climate change 2007: The scientific basis, Technical report, Intergovernmental Panel on Climate Change (IPCC), IPCC Secretariat, c/o World Meteorological Organization, Geneva, Switzerland, 2007
3. A.B. Davis, A. Marshak, E. Kassianov, G.M. Stokes, EOS Trans. **80**, 622 (1999)
4. A. Schuster, Astrophys. J. **21**, 1 (1905)
5. K.F. Evans, J. Atmos. Sci. **55**, 429 (1998)
6. C. Emde, S.A. Buehler, C. Davis, P. Eriksson, T.R. Sreerekha, C. Teichmann, J. Geophys. Res., 109 (2004) doi:10.1029/2004JD005140
7. B. Mayer, In *Intercomparison of three-dimensional radiation codes: Abstracts of the first and second international workshops*, Institute of Atmospheric Physics, University of Arizona, 2000, ISBN: 0-9709609-0-5
8. C. Emde, B. Mayer, Atmos. Chem. Phys. **7**, 2259 (2007)
9. A. Marshak, A.B. Davis (eds.), *3D Radiative Transfer in Cloudy Atmospheres* (Springer, Berlin, Heidelberg, New York, 2005)
10. B. Mayer, A. Kylling, Atmos. Chem. Phys. **5**, 1855 (2005)
11. K. Stamnes, S.C. Tsay, W. Wiscombe, K. Jayaweera, Appl. Opt. **27**, 2502 (1988)
12. R.F. Cahalan, D. Silberstein, J.B. Snider, J. Atmos. Sci. **52**, 3002 (1995)
13. A. Kazantzidis, A.F. Bais, C. Emde, S. Kazadzis, C.S. Zerefos, Atmos. Chem. Phys. **7**, 5959 (2007)
14. S.W. Hoch, B. Mayer, C.D. Whiteman, In *Proceedings of the 13th Conference on Mountain Meteorology*, August 11–15, Whistler, BC, Canada, 2008
15. E. Veach, *Robust Monte Carlo Methods for Light Transport Simulation*, Ph.D. thesis, Stanford University, 1997, available at http://www-graphics.stanford.edu/papers/veach_thesis/
16. M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, F. Rossi, *GNU Scientific Library Reference Manual - Revised Second Edition*, Free Software Foundation, Boston, MA, USA, 2008, ISBN 0954161734; available from <http://www.gnu.org/software/gsl/>
17. P. Wendling, J. Atmos. Sci. **34**, 642 (1977)
18. T. Zinner, B. Mayer, M. Schröder, J. Geophys. Res. **111** (2006) doi:10.1029/2005JD006062
19. T. Zinner, A. Marshak, S. Lang, J.V. Martins, B. Mayer, Atmos. Chem. Phys. **8**, 4741 (2008)
20. G. Marchuk, G. Mikhailov, M. Nazaraiev, R. Darbinjan, B. Kargin, B. Elepov, *The Monte Carlo Methods in Atmospheric Optics* (Springer-Verlag, New York, 1980) ISBN-13: 978-0387094021