

PART III

SOLUTION METHODS

We now have developed a comprehensive theory of radiative transfer within natural water bodies, including the effects of the surface and bottom boundaries. However, our theory is of limited value unless we can readily solve the associated equations. Part III of our book therefore discusses *selected numerical methods for solving radiative transfer equations in conjunction with their boundary conditions*. As in Chapter 4, sections dealing with specific details of numerical algorithms will be set off by bars.

Chapter 6 shows how *Monte Carlo ray tracing* methods can be used to solve the radiance transfer equation within any water body, even those whose boundary conditions and inherent optical properties vary in all three spatial dimensions. Although very general and powerful, Monte Carlo methods are computationally very inefficient for many problems in optical oceanography.

Chapter 7 therefore begins the development of *invariant imbedding* solution methods, which are computationally very efficient for problems of one spatial dimension, e.g. for problems governed by the RTE of Eq. (5.23). *This chapter develops the essential concepts of invariant imbedding theory in the setting of the two-flow irradiance equations*, which provide a mathematically simple set of equations for illustrating the solution method.

Chapter 8 then shows how invariant imbedding methods are applied to the radiance transfer equation itself. This chapter culminates in *a specific numerical algorithm for solving the RTE along with the associated boundary conditions at the air-water surface and at the bottom*.

Chapter 9 begins with an overview of the *discrete ordinates* solution method. This overview sets the stage for a general discussion of *eigenmatrix* solution methods, which are applicable to *homogeneous water bodies*. This chapter finishes with a discussion of the asymptotic radiance distribution, including numerical examples.

Chapter 10 discusses *inverse methods*, whose purpose is to extract information about the inherent optical properties of a water body from light field measurements made either within the water body or remotely. This chapter concludes with a discussion of *ocean-color remote sensing*.

This Page Intentionally Left Blank

Chapter 6

Monte Carlo Methods

In this chapter we present the most general technique for numerical solution of the radiance transfer equation. By "most general" we mean that the technique is applicable to the time-dependent, three-dimensional RTE of Eq. (5.19) in a setting with arbitrary boundary geometry and incident radiance, and with arbitrary inherent optical properties within the water body. This general technique employs Monte Carlo methods, which we already have seen applied in Chapter 4 to the computation of the air-water surface reflectance and transmittance functions. We now extend those ray-tracing techniques to the water body itself.

The idea underlying all Monte Carlo methods is this: if we know the probability of occurrence of each separate event in a sequence of events, then we can determine the probability that the entire sequence of events will occur. In Section 6.1, we show how these probabilities are determined. Our discussion there is straightforward but mathematically naive, and the resulting numerical algorithms are consequently inefficient. In Sections 6.2 and 6.3, we take a more sophisticated look at Monte Carlo methods and learn how to improve their computational efficiency by many orders of magnitude.

6.1 Forward Monte Carlo Methods

For pedagogic purposes, let us consider the measurement of downwelling plane irradiance E_d at some depth z below a random sea surface. Each point on the sea surface is illuminated by an incident sky radiance $L(a; \hat{\xi})$, $\hat{\xi} \in \Xi_d$. Figure 6.1 illustrates a specific example of the fate of three photons. Photon a encounters the sea surface at point 1, where it is refracted into the water in direction $\hat{\xi}_1$. The photon then encounters a phytoplankton cell at point 2, scatters through an angle ψ , and continues on its way in direction $\hat{\xi}_2$ until it encounters the irradiance detector at point 3. Its energy contributes to the irradiance measured at point 3 [see [Supplementary Note 10](#)]. Photon b enters the water, travels a ways, and is absorbed by a water molecule at point 4. Photon c enters the water, undergoes scattering events at points 5 and 6, and continues into the oceanic depths. Clearly, photons b and c do not contribute to the E_d measured at point 3.




Fig. 6.1. Illustration of three photon trajectories and of the computation of E_d .

We can numerically simulate the physical situation just described. A realization of the random air-water surface is generated as described in Sections 4.3 or 4.9. Photon paths representative of the given sky radiance distribution are generated as in Section 4.4, and the air-incident photons are allowed to interact with the surface wave facets. Now, however, rather than just tallying the directions of the transmitted photons, we continue to track the photons as they penetrate into the water body. This is done as follows.

Sampling photon path lengths

Consider photon *a* as shown in Fig. 6.1 and, for simplicity, assume that the water is homogeneous. After the photon enters the water, its direction $\hat{\xi}_1$ is known from its initial direction and subsequent interactions with the water surface. We next must determine how far the photon travels in direction $\hat{\xi}_1$ before it encounters a water molecule or other particle. To see how this can be done, think of photon *a* not as one photon but as a collimated beam or ray of many photons, as was done in Chapter 4. This beam has some radiance $L(\vec{x}; \hat{\xi}_1)$, which decreases with distance according to Eq. (5.15):

$$\frac{dL(\vec{x};\hat{\xi}_1)}{dr} = -cL(\vec{x};\hat{\xi}_1),$$

or

$$L(r;\hat{\xi}_1) = L(0;\hat{\xi}_1)e^{-cr}.$$

Here, as before, r is geometric distance along direction $\hat{\xi}_1$, measured from some reference point \vec{x}_0 at $r = 0$. In terms of the optical path length $l = cr$ defined in Section 5.7, we have

$$L(l;\hat{\xi}_1) = L(0;\hat{\xi}_1)e^{-l}. \quad (6.1)$$

This decrease in radiance can be explained in terms of the fate of individual photons if the probability of any particular photon being absorbed or scattered out of the beam between optical path lengths l and $l+dl$ is

$$p_l(l) dl \equiv e^{-l} dl. \quad (6.2)$$

Note that this $p_l(l)$ satisfies

$$\int_0^{\infty} p_l(l) dl = 1,$$

as is required of any *probability density function* (pdf). The probability that a photon is absorbed or scattered somewhere between $l = 0$ and l is given by the *cumulative distribution function* $P_l(l)$:

$$P_l(l) \equiv \int_0^l p_l(l') dl' = 1 - e^{-l}. \quad (6.3)$$

Now let \mathfrak{R} be a random number drawn from the unit interval between 0 and 1 such that \mathfrak{R} is equally likely to have any value $0 \leq \mathfrak{R} \leq 1$. We denote this distribution of \mathfrak{R} by saying that \mathfrak{R} is uniformly distributed on the interval 0 to 1. The associated pdf of \mathfrak{R} is given by

$$p_{\mathfrak{R}}(\mathfrak{R}) \equiv \begin{cases} 1 & \text{if } 0 \leq \mathfrak{R} \leq 1 \\ 0 & \text{if } \mathfrak{R} < 0 \text{ or if } \mathfrak{R} > 1. \end{cases}$$

We wish to use the randomly drawn \mathfrak{R} , which is a known number, to determine a value for l . The key to this determination is to regard going from \mathfrak{R} to l as a change of variables. Then the probability that \mathfrak{R} is in some interval \mathfrak{R} to $\mathfrak{R}+d\mathfrak{R}$ is $p_{\mathfrak{R}}(\mathfrak{R})d\mathfrak{R}$, and the probability that l is in a corresponding interval l to $l+dl$ is $p_l(l)dl$. These two probabilities must be equal, in which case

$$\int_0^{\Re} p_{\Re}(\Re') d\Re' = \int_0^l p_l(l') dl'.$$

Because $p_{\Re}(\Re)$ is known, the left-hand integral can be evaluated. The preceding equation then becomes

$$\Re = \int_0^l p_l(l') dl' = P_l(l).$$

We now make the following observation, which is sometimes called the *fundamental principle of Monte Carlo simulation* (Cashwell and Everett, 1959): *The equation $\Re = P_l(l)$ uniquely determines l in such a manner that l falls in the interval l to $l + dl$ with frequency $p_l(l)dl$.* This result is completely general.

In the case at hand,

$$\Re = P_l(l) = 1 - e^{-l}$$

gives

$$l = -\ln(1 - \Re). \quad (6.4)$$

We can therefore use Eq. (6.4) to randomly determine the optical path length l traveled by a photon between one scattering or absorption event and another. Distances so chosen, when applied to many photons in a collimated beam, are consistent with the radiance behavior of Eq. (6.1).

Note that since $1 - \Re$ is also uniformly distributed on the interval 0 to 1, we can just as well draw \Re and use

$$l = -\ln \Re, \quad (6.5a)$$

or

$$r = -\frac{1}{c} \ln \Re \quad (6.5b)$$

in terms of the geometric distance $r = l/c$.

We note in passing that the *average geometric distance* $\langle r \rangle = \langle l/c \rangle$ traveled by photons between interactions is

$$\langle r \rangle \equiv \int_0^{\infty} \frac{l}{c} p(l) dl = \int_0^{\infty} \frac{l}{c} e^{-l} dl = \frac{1}{c} \quad (\text{m}). \quad (6.6)$$

The average distance $\langle r \rangle$ is called the *mean free path*, and Eq. (6.6) shows that the beam attenuation coefficient is just the reciprocal of the photon mean free path.

Sampling photon interaction types

The preceding observations on how to randomly determine photon path lengths r give us the conceptual tools necessary to simulate almost any physical process. Suppose we have decided, using the random-number method described above, that a photon should travel a distance r (e.g. from point 1 to point 2 of Fig. 6.1). The photon then interacts with the medium (at point 2). We next need to decide at random if the interaction is to be an absorption or a scattering event. This is easily done by drawing another random number \mathfrak{R} from $U[0,1]$ and comparing it with the albedo of single scattering, $\omega_o \equiv b/c$. If $\mathfrak{R} > \omega_o$, we let the interaction be an absorption event; if $\mathfrak{R} \leq \omega_o$, the interaction results in scattering. Note that this procedure *on average* produces absorption and scattering events in the respective proportions of a/c and b/c . This use of ω_o makes clear why it is sometimes called *the probability of photon survival*.

Sampling scattering directions

If the interaction results in absorption, we terminate the photon, as is the fate of photon b of Fig. 6.1. If the interaction is a scattering event, we can randomly determine the new photon direction by use of the phase function $\tilde{\beta}$. Recall that $\tilde{\beta}(\hat{\xi}' \rightarrow \hat{\xi})$ can be interpreted as a probability density function for scattering from direction $\hat{\xi}'$ into direction $\hat{\xi}$. Thus the probability of scattering into an element of solid angle $d\Omega(\hat{\xi})$ centered on direction $\hat{\xi}$ is

$$\tilde{\beta}(\hat{\xi}' \rightarrow \hat{\xi}) d\Omega(\hat{\xi}) = \tilde{\beta}(\psi, \varphi) \sin\psi d\psi d\varphi.$$

Here ψ and φ are the directions of the scattered photon in a coordinate system centered on the incident direction $\hat{\xi}'$.

In natural waters, $\tilde{\beta}$ depends only on the scattering angle ψ . Therefore $\tilde{\beta}(\psi, \varphi) = \tilde{\beta}(\psi)$, and ψ and φ are independent random variables, which means that two random numbers must be drawn in order to determine both ψ and φ . The independence of ψ and φ also means that the joint pdf can be written as a product of pdf's for ψ and φ :

$$\tilde{\beta}(\psi) \sin\psi d\psi d\varphi \equiv p_\psi(\psi) d\psi p_\varphi(\varphi) d\varphi. \quad (6.7)$$

The azimuthal angle φ (which determines the plane of the scattering event relative to some reference direction) is uniformly distributed on the interval from 0 to 2π . Thus $p_\varphi(\varphi)d\varphi = (1/2\pi)d\varphi$. We can use this result in Eq. (6.7) to identify the pdf $p_\psi(\psi)$ to be used for determining ψ :

$$p_{\psi}(\psi) = 2\pi\tilde{\beta}(\psi) \sin\psi. \quad (6.8)$$

In order to determine the scattering angle ψ , we draw a random number \mathfrak{R} from the uniform distribution on 0 to 1 and set

$$\mathfrak{R} = P_{\psi}(\psi) = \int_0^{\psi} p_{\psi}(\psi') d\psi' = 2\pi \int_0^{\psi} \tilde{\beta}(\psi') \sin\psi' d\psi'. \quad (6.9)$$

Because of the complicated shape of $\tilde{\beta}(\psi)$, Eq. (6.9) must in general be solved numerically in order to determine the value of ψ corresponding to \mathfrak{R} . The same type of analysis applied to φ yields

$$\varphi = 2\pi\mathfrak{R}, \quad (6.10)$$

where this \mathfrak{R} is yet another random number drawn from the uniform distribution on 0 to 1.

We have now seen in some detail how to randomly determine photon path lengths, types of interactions, and scattering directions. The same type of reasoning is easily extended. For example, if we decide that a particular interaction will be a scattering event, then we can randomly decide if the scattering will be by a water molecule or by a large particle, with different phase functions to be used for each case. If the scattering is to be by a water molecule, we can choose whether it will be elastic scattering or Raman scattering to another wavelength λ , where λ will be chosen at random using the Raman wavelength redistribution function of Eq. (5.94) as the relevant pdf.

Solving the RTE

It should be clear that, as we trace many photons, we are building up a solution of the RTE one photon at a time. This is exactly what nature does. The occasional photon that enters the detector shown in Fig. 6.1 contributes to the desired value of E_d . We can compute any other radiometric quantity by a suitable tally of the photons. For example, if we count only the photons traveling within a small element of solid angle at a given location, we are building up the radiance in that direction at that location. Simulated detectors can of course be placed throughout the air and water. Note also that the boundary conditions are automatically "built in" to the solution as the photons interact with the surface or other boundaries. Internal sources can be included in the RTE (5.23) by generating photons at the desired locations within the water. The directional and wavelength distributions of the internal source photons are randomly determined from pdf's chosen to represent the internal source.

For simplicity of notation in Eqs. (6.1)-(6.6), we assumed that the water was homogeneous. This is not necessary. One possible procedure is to divide the water column into many thin layers, each of which is taken to be homogeneous, and then proceed as above within each layer. However, even this complication is unnecessary if we work with optical path lengths l rather than with geometric distances. That is to say, *we can solve the RTE (5.24) in terms of optical depth ζ without regard for spatial inhomogeneities in the actual water body.* If we wish to place our simulated photon detectors at specific geometric depths, then we must convert these depths to the corresponding optical depths at the start of the Monte Carlo simulation, in order to know where to tally the photons. Conversely, we can tally the photons at conveniently chosen optical depths, and then determine the corresponding geometric depths when the final output is produced at the end of the simulation.

The procedures just described are called *forward* or *analog* Monte Carlo methods. "Forward" reminds us that the photons are traced *forward in time*, just as in nature (the significance of this phrase will be seen in the next section). "Analog" means that the mathematical simulation is an exact analogy to the physical processes being studied.

Strengths and weaknesses

We can now make several observations about these forward Monte Carlo methods:

- They are *conceptually simple*. The methods are based on a straightforward mimicry of nature. This in itself endows the methods with a certain elegance.
- They are *instructive*. The methods highlight the fundamental radiative transfer processes of absorption and scattering, and they make clear the connection between the photon-level and the energy-level formulations of radiative transfer theory.
- They are *very general*. The methods are applicable to any geometry, incident lighting, scattering phase function, distribution of inherent optical properties, etc. It is easy to include polarization.
- They are *simple to program* for a digital computer. The methods are also well suited to parallel processing, in which many independent photons are traced at once.

- They provide *no insight* into the underlying mathematical structure of radiative transfer theory.
- They can be *computationally extremely inefficient*. In a situation like the one illustrated in Fig. 6.1, almost all of the photons generated are "wasted" because they never intercept a detector.

A simple calculation allows us to estimate the potential magnitude of this computational inefficiency. As a worst-case scenario, suppose in Fig. 6.1 that the irradiance detector is located at the bottom of the euphotic zone, i.e. at a depth z_{eu} where $E_d(z_{\text{eu}}) \approx 0.01E_d(0)$. Table 3.18 shows that z_{eu} can be greater than 100 m in clear ocean waters. Now there is some region of sea surface centered above the detector, within which the photons incident on the surface have a significant chance of eventually reaching the detector. Suppose this surface region is a circle of radius r_{surf} . In order to properly estimate $E_d(z_{\text{eu}})$, we should let photons fall onto the sea surface throughout the region defined by r_{surf} . But only approximately 1% of the photons incident on the sea surface will reach depth z_{eu} , and of those only a few will actually strike the detector. If the detector is a circle of radius r_{det} , then the ratio of photons actually hitting the detector at depth z_{eu} to the number of photons incident on the surface region is

$$0.01 \frac{\pi r_{\text{det}}^2}{\pi r_{\text{surf}}^2}.$$

For reasonable values of $r_{\text{det}} = 0.01$ m and $r_{\text{surf}} = 100$ m, the proportion of "useful" photons is then 10^{-10} . Thus on average only one photon in 10^{10} is actually counted by the detector in the computation of E_d . Such a situation can arise in problems that are inherently three-dimensional. An example would be a computation of ship-shadow effects on an instrument in the water near the ship.

However, the situation just described is unreasonably pessimistic from the viewpoint of most oceanographic problems. We note in particular that *if the water body is horizontally homogeneous, then every photon reaching the depth of the detector can be used to compute E_d at that depth*. In the example above, we are then using 1% of the initial photons, not just 1 in 10^{10} . This is a rather significant improvement, although we are still wasting 99% of our photons.

Another drawback of all Monte Carlo methods is the presence of *statistical sampling noise* in the estimated values. In the above example, each photon collected contributes to the E_d value. But since the photons come from random directions, and since each is weighted by a $\cos\theta$ factor,

the "current estimate" of E_d changes as each new photon is collected. According to standard statistical reasoning, the standard error of the estimate (the sample standard deviation) is proportional to $n^{-1/2}$, where n is the number of photons collected. Thus if we wish to improve the accuracy of a given estimate by a factor of ten, we must collect 100 times as many photons. This slow convergence to the true value of the estimated quantity can be a serious computational disadvantage if very accurate estimates are required.

However, in computations of E_d , the statistical noise is seldom a problem even at great depths. Consider, for example, a Monte Carlo simulation that introduces 10^6 photons at the sea surface. If 1% of these photons reaches the bottom of the euphotic zone, we still have left $n = 10^4$ photons. In this case, $n^{-1/2} = 0.01$, so that we can estimate E_d to 1% accuracy. This is better accuracy than can be achieved with standard oceanographic instruments. Tracing 10^6 photon histories is well within reason for routine computations.

However, if we are computing the zenith (nadir-viewing) radiance, then only the relatively few photons that are backscattered into a small solid angle centered on the zenith direction can be tallied. In such a situation, sampling noise can be a serious problem.

The next two sections present ways to minimize the computational disadvantages of the forward Monte Carlo methods. Backward Monte Carlo methods allow us to make use of almost every photon generated, no matter what radiometric quantity is being estimated. Variance reduction techniques allow us to play our statistical sampling games with "loaded dice," in order to make better use of each generated photon.

6.2 Backward Monte Carlo Methods

As we have just seen, very few of the simulated photons incident on the sea surface are actually collected by a simulated instrument at a particular location. Indeed, if the instrument is approximated as a point, then there is *zero* probability of a photon intercepting the instrument. Moreover, when simulating inherently three-dimensional situations, we cannot know in advance how large an area on the sea surface should receive photons, because this region of influence depends on environmental conditions, water IOP's, and instrument location and orientation.

Reciprocity relations and adjoint problems

We therefore seek to formulate a problem that can be solved efficiently, and whose solution can be transformed into the solution of the

problem at hand. The key to constructing and solving this *adjoint problem* is the following simple but profound observation. Consider for example photon a of Fig. 6.1, which is traveling in direction $\hat{\xi}_2$ when it intercepts the detector. *If a photon were to leave the detector traveling in direction $-\hat{\xi}_2$, then that photon would exactly retrace the path of the original photon, all the way back to the original photon's source.* We already have commented in Section 4.2 on this reciprocity for photons crossing the air-water surface. The same result holds true within the water so long as $\tilde{\beta}(-\hat{\xi}' \rightarrow -\hat{\xi}) = \tilde{\beta}(\hat{\xi}' \rightarrow \hat{\xi})$, which is the case in natural waters. We now ask whether it is in some way possible to generate photons at the detector, trace them on their random journeys away from the detector, tally them as they leave the air-water surface heading upward, and convert this final tally into an estimate for the desired E_d value (or for any other radiometric quantity). If this is possible, then every photon generated will contribute to the estimate of E_d . We now show how this is done.

Figure 6.2(a) illustrates the original problem of Fig. 6.1, and Fig. 6.2(b) shows the corresponding *time-reversed* or *adjoint* problem. Points lying on the boundary surface B of the ocean volume V are denoted by $\vec{x} \in B$. For mathematical preciseness, the boundary B also includes the lateral and bottom boundaries of the volume V , but we show only the sea surface itself (the other boundaries can be taken to be at infinity if desired). Vector \vec{x}_d gives the position of a detector measuring, say, E_d . An arbitrary point within volume V is denoted by \vec{x} . There is an air-incident radiance $L_1(\vec{x}_B; \hat{\xi})$, $\hat{\xi} \in \Xi_d$, falling into the air-water surface. The *outward* normal to the boundary surface is \hat{n} . To avoid notational complications arising with a wave-covered sea surface, we can imagine a horizontal monitoring surface (h.m.s.) just above the mean sea surface, as was done in Section 4.5. The normal to the h.m.s. just above the sea surface is then $\hat{n} = -\hat{i}_3$. The index of refraction of the water is n_w ; that of the air is taken to be $n_a = 1$. Figure 6.2(a) shows a typical photon path going from air to detector.

Case (1957) shows that the following *reciprocity relation* holds true:

$$\begin{aligned} & \int_{\hat{\xi} \cdot \hat{n} < 0} d\Omega(\hat{\xi}) \int_B dB |\hat{\xi} \cdot \hat{n}| \left[L_1(\vec{x}_B; \hat{\xi}) L_2(\vec{x}_B; -\hat{\xi}) - L_2(\vec{x}_B; \hat{\xi}) L_1(\vec{x}_B; -\hat{\xi}) \right] \\ &= \int_{\Xi} d\Omega(\hat{\xi}) \int_V dV \left[\frac{1}{n_w^2} L_1(\vec{x}; -\hat{\xi}) S_2(\vec{x}; \hat{\xi}) - \frac{1}{n_w^2} L_2(\vec{x}; \hat{\xi}) S_1(\vec{x}; -\hat{\xi}) \right]. \end{aligned} \quad (6.11)$$

On the left hand side of Eq. (6.11), the integral over B is an integral over the entire boundary surface (dB is an element of area), and the integral over Ω is an integral over all directions $\hat{\xi}$ toward the surface. The integrals on the right hand side are over the entire volume (dV is an element of volume)




Fig. 6.2 Illustration of the original (forward) and adjoint (time-reversed) problems used to develop backward Monte Carlo methods.

and over all directions. $L_1(\vec{x}_B; \hat{\xi})$ is the incident radiance falling onto the boundary surface (since $\hat{\xi} \cdot \hat{n} < 0$) in the original problem, and $L_1(\vec{x}_B; -\hat{\xi})$ is the radiance leaving the surface (since $-\hat{\xi} \cdot \hat{n} > 0$). Likewise, $L_2(\vec{x}_B; \hat{\xi})$ is the radiance falling onto the surface in the adjoint problem, and $L_2(\vec{x}_B; -\hat{\xi})$ is the radiance leaving the surface. We shall consider only that part of the total boundary B that corresponds to the air-water surface. Thus $L_1(\vec{x}_B; \hat{\xi})$ and $L_2(\vec{x}_B; \hat{\xi})$ correspond to the air-incident radiances for the original and adjoint problems, respectively, and $L_1(\vec{x}_B; -\hat{\xi})$ and $L_2(\vec{x}_B; -\hat{\xi})$ are the water-leaving radiances for the original and adjoint problems. S_1 and S_2 are respectively the internal sources for the original and adjoint problems. The origin of the n_w^{-2} factors on the right hand side of Eq. (6.11) can be traced to the L/n_w^2 terms in Eq. (5.19). In Eq. (6.11) note the juxtaposition of original-problem quantities whose photons are traveling in a given direction, with adjoint-problem quantities whose photons are traveling in the opposite direction. This very general result is the quantitative foundation of backward Monte Carlo methods.

Example

Gordon (1985) nicely exploited Eq. (6.11) in the oceanic setting. We shall illustrate his development with a specific example. Let us assume that

we wish to compute $E_d(\vec{x}_D)$, that the incident sky radiance is a cardioidal distribution, and that there are no internal sources in the original problem. We therefore take

$$L_1(\vec{x}_B; \hat{\xi}) = L_o (1 + C |\hat{\xi} \cdot \hat{i}_3|) \quad \text{and} \quad S_1(\vec{x}; -\hat{\xi}) = 0, \quad (6.12)$$

where L_o and C are positive constants. By definition the desired irradiance $E_d(\vec{x}_D)$ in the original problem is

$$E_{d1}(\vec{x}_D) = \int_{\hat{\xi} \cdot \hat{i}_3 > 0} L_1(\vec{x}_D; \hat{\xi}) \hat{\xi} \cdot \hat{i}_3 d\Omega(\hat{\xi}), \quad (6.13)$$

where we have recalled our choice of coordinate systems with \hat{i}_3 positive downward.

Now *choose* the internal source for the adjoint problem to be

$$S_2(\vec{x}; \hat{\xi}) \equiv \begin{cases} -J_o \hat{\xi} \cdot \hat{i}_3 \delta(\vec{x} - \vec{x}_D) & \text{if } \hat{\xi} \cdot \hat{i}_3 < 0 \\ 0 & \text{if } \hat{\xi} \cdot \hat{i}_3 \geq 0, \end{cases} \quad (6.14)$$

where J_o is a positive constant. Then the right hand side of Eq. (6.11) reduces to

$$\int_{\hat{\xi} \cdot \hat{i}_3 < 0} d\Omega \int_V dV \frac{1}{n_w^2} L_1(\vec{x}; -\hat{\xi}) (-J_o \hat{\xi} \cdot \hat{i}_3 \delta(\vec{x} - \vec{x}_D)) \quad (6.15a)$$

$$= -\frac{J_o}{n_w^2} \int_{\hat{\xi} \cdot \hat{i}_3 < 0} d\Omega L_1(\vec{x}_D; -\hat{\xi}) \hat{\xi} \cdot \hat{i}_3 \quad (6.15b)$$

$$= \frac{J_o}{n_w^2} \int_{\hat{\xi} \cdot \hat{i}_3 > 0} L_1(\vec{x}_D; \hat{\xi}) \hat{\xi} \cdot \hat{i}_3 d\Omega = \frac{J_o}{n_w^2} E_{d1}(\vec{x}_D). \quad (6.15c)$$

In going from Eq. (6.15a) to (6.15b) we have used the properties of the Dirac delta function, and we have let $\hat{\xi} \rightarrow -\hat{\xi}$ and recalled Eq. (6.13) in going from (6.15b) to (6.15c).

In the adjoint problem there is a source S_2 within the water, and photons are leaving the surface. However, there is no air-incident radiance, so $L_2(\vec{x}_B; \hat{\xi}) = 0$. The left hand side of Eq. (6.11) therefore reduces to

$$\int_{\hat{\xi} \cdot \hat{n} < 0} d\Omega \int_B dB |\hat{\xi} \cdot \hat{n}| L_1(\vec{x}_B; \hat{\xi}) L_2(\vec{x}_B; -\hat{\xi}) \quad (6.16a)$$

$$= L_o \int_B dB \int_{\hat{\xi} \cdot \hat{n} < 0} d\Omega |\hat{\xi} \cdot \hat{n}| \left(1 + C |\hat{\xi} \cdot \hat{i}_3| \right) L_2(\vec{x}_B; -\hat{\xi}) \quad (6.16b)$$

$$\equiv L_o \int_B dB \check{E}_{u2}(\vec{x}_B). \quad (6.16c)$$

Here we have substituted for $L_1(\vec{x}_B; \hat{\xi})$ from Eq. (6.12) and defined the quantity

$$\check{E}_{u2}(\vec{x}_B) \equiv \int_{\hat{\xi} \cdot \hat{n} < 0} d\Omega(\hat{\xi}) |\hat{\xi} \cdot \hat{n}| \left(1 + C |\hat{\xi} \cdot \hat{i}_3| \right) L_2(\vec{x}_B; -\hat{\xi}). \quad (6.17)$$

As already noted, $L_2(\vec{x}_B; -\hat{\xi})$ is the adjoint radiance *leaving* the surface ($\hat{\xi} \cdot \hat{n} < 0 \Rightarrow \hat{\xi} \in \Xi_d$, so $-\hat{\xi} \in \Xi_u$). Thus $\check{E}_{u2}(\vec{x}_B)$ is similar in form to the adjoint upward plane irradiance just above the surface, $E_{u2}(\vec{x}_B)$. The difference is that in computing $\check{E}_{u2}(\vec{x}_B)$, *each photon leaving the surface is weighted by a geometric factor proportional to the angular distribution of the incident radiance $L_1(\vec{x}_B; \hat{\xi})$ in the original problem.* The quantity $\check{E}_{u2}(\vec{x}_B)$ is easily computed by appropriately tallying the photons that leave the air-water surface.

We have now reduced Eq. (6.11) to

$$L_o \int_B dB \check{E}_{u2}(\vec{x}_B) = \frac{J_o}{n_w^2} E_{d1}(\vec{x}_D). \quad (6.18)$$

This equation relates the desired quantity $E_{d1}(\vec{x}_D)$ in the original problem to known or computable quantities in the adjoint problem. We can rewrite Eq. (6.18) by recalling from problem 1.5 that the original downward plane irradiance onto the sea surface for the cardioid distribution of Eq. (6.12) is

$$E_{d1}(\vec{x}_B) = 2\pi L_o \left(\frac{1}{2} + \frac{C}{3} \right).$$

Moreover, the total power P_2 emitted by the adjoint source of Eq. (6.14) is

$$P_2 = \int_{\Xi} d\Omega(\hat{\xi}) \int_V dV S_2(\vec{x}; \hat{\xi}) = \pi J_o.$$

Using these results in Eq. (6.18) gives

$$\frac{E_{d1}(\vec{x}_D)}{E_{d1}(\vec{x}_B)} = \frac{n_w^2}{P_2} \frac{1}{\left(1 + \frac{2}{3}C\right)} \int_B dB \check{E}_{u2}(\vec{x}_B). \quad (6.19)$$

Thus the ratio of the irradiance at the detector to the irradiance incident on the sea surface in the original problem is proportional to the ratio of a computable quantity (with units of watts) to the power emitted by an imbedded source in the adjoint problem. This statement holds in general, even though we have deduced it by reference to a specific problem.

Strengths and weaknesses

We now have seen the essence of backward Monte Carlo methods. Given an original problem, we must first conjure up an adjoint source S_2 that will allow us to extract the original quantity of interest from the right-hand side of Eq. (6.11). Using the adjoint source, we then generate and trace photons. Each photon leaving the boundary gets weighted by a factor proportional to the incident radiance distribution of the original problem. A tally of the weighted photons leaving the boundary then yields the original quantity of interest. Table 6.1 shows the source functions S_2 that yield commonly used radiometric quantities.

There is one common situation for which the above scheme requires modification: the case of a collimated incident radiance distribution. In this case the weighting function for the photons exiting the sea surface is a Dirac delta function specifying the direction of the collimated incident radiance. Thus almost all photons exiting the surface in the adjoint problem would receive a weight of zero, since they would not exit parallel to the incident photons of the original problem. Gordon (1985) presents a suitable simulation strategy for this special case.

In addition to their obvious computational efficiency (almost all generated photons are eventually tallied), backward Monte Carlo methods avoid several subtle problems inherent in forward methods. For example, we do not need to decide in advance how large a region of the water surface should receive incident photons. Also, the size of the detector does not influence the number of photons recorded. Backward Monte Carlo methods can even simulate point detectors, which is not possible in forward Monte Carlo methods.

Backward Monte Carlo methods hold the greatest computational advantage over forward methods when the incident light source is spatially distributed and the detector acceptance angle is small. This is precisely the situation in hydrologic optics when sun and sky light falls in the water

Table 6.1. Source functions $S_2(\vec{x}; \hat{\xi})$ for use in backward Monte Carlo simulations leading to various radiometric quantities^a.

desired radiometric quantity	source function $S_2(\vec{x}; \hat{\xi})$
$E_d(\vec{x}_D)$	$\begin{cases} -J_o \hat{\xi} \cdot \hat{i}_3 \delta(\vec{x} - \mathbf{x}_D) & \text{if } \hat{\xi} \cdot \hat{i}_3 < 0 \\ 0 & \text{if } \hat{\xi} \cdot \hat{i}_3 \geq 0 \end{cases}$
$E_u(\vec{x}_D)$	$\begin{cases} 0 & \text{if } \hat{\xi} \cdot \hat{i}_3 \leq 0 \\ J_o \hat{\xi} \cdot \hat{i}_3 \delta(\vec{x} - \mathbf{x}_D) & \text{if } \hat{\xi} \cdot \hat{i}_3 > 0 \end{cases}$
$E_o(\vec{x}_D)$	$J_o \delta(\vec{x} - \mathbf{x}_D)$ for all $\hat{\xi} \cdot \hat{i}_3$
$L(\vec{x}_D; \hat{\xi}')$	$J_o \delta(\hat{\xi}' - \hat{\xi}) \delta(\vec{x} - \mathbf{x}_D)$

a. Based on Gordon (1985)

surface, and when we wish to compute radiometric quantities at fixed points within the water. If the incident radiance distribution in the original problem is generated by a point source, then backward Monte Carlo methods hold no advantage over forward methods. An example of a forward Monte Carlo simulation of bioluminescence modeled as an embedded point source is seen in Gordon (1987). However, the reciprocity relation (6.11) can still be of utility in point-source problems. Gordon (1984) used it to great advantage in an analytic study of the effect of the depth of a bioluminescent point source on the remotely sensed signature of the point source.

A path different than Eq. (6.11) but still leading to backward Monte Carlo methods was taken in pioneering papers by Collins, *et al.* (1972) and Adams and Kattawar (1978). They viewed the reformulation from forward to backward Monte Carlo as a change of variables from "source" to "detector" coordinates. The end result is the same as our development based on the reciprocity relation (6.11).

We end this section by pointing out that the reciprocity relation (6.11) rests on the validity of the time-reversal symmetry of photon paths, i.e. on the fact that two photons with directions $\hat{\xi}$ and $-\hat{\xi}$ trace out the same paths in space. This symmetry itself can be traced back to the time-

reversal invariance of Maxwell's equations. Note that time reversal is *not* the same thing as simply interchanging the source and receiver locations.

6.3 Variance Reduction Techniques

Regardless of whether we use forward or backward Monte Carlo methods, there are random errors in the estimated quantities because of the inherently statistical nature of Monte Carlo methods. However, it is often possible to reduce these random errors (for a given amount of computational expense) through the use of clever sampling strategies.

The simplest such strategy is this: *never waste a photon*. For simplicity of the preceding discussions and to highlight the relevant physics, we have talked about simulating photons one at a time. This is not good practice because an absorbed photon never reaches a detector. The computational expense of generating and tracking the photon before its absorption is therefore wasted. This inefficiency is easily avoided by considering each "photon" in the above discussions to be a packet or ray of many identical photons. The "weight" w of this photon packet is usually set to $w = 1$ when it is generated. At each subsequent interaction, the weight is multiplied by $\omega_o = b/c$, and the packet is scattered into a new direction for further tracing. Multiplying the current weight by ω_o is equivalent to saying that a fraction $1 - \omega_o = a/c$ of the photons in the packet is absorbed, and a fraction b/c continues on its way in the scattered direction. This is analogous to what was done in Section 4.4 when rays interacting with the air-water surface had their weights (radiant powers) multiplied by Fresnel reflectances and transmittances. A ray is traced until it is tallied by a detector or until its weight drops below some prechosen value, say 10^{-6} .

Biased Sampling

There are many schemes for reducing the variance (i.e. the statistical error) in estimates by the use of *biased sampling*. The general theory of these techniques is simple. Suppose we wish to compute the expected value $\langle f \rangle$ of some function $f(x)$ of a random variable x , which is drawn from the (correct) probability density function $p(x)$. Then by definition

$$\langle f \rangle \equiv \int f(x) p(x) dx ,$$

and the variance in f is by definition

$$\sigma^2(f) \equiv \int [f(x) - \langle f \rangle]^2 p(x) dx . \quad (6.20)$$

Now suppose that we draw the random variable x from some other density $p_b(x)$, the *biased* pdf. Then

$$\langle f \rangle = \int f(x) \frac{p(x)}{p_b(x)} p_b(x) dx \equiv \int f(x) w(x) p_b(x) dx \equiv \langle fw \rangle_b. \quad (6.21)$$

The quantity $w(x)$ is called the *statistical weight*. Note that the expected value of f computed when using the correct pdf, $\langle f \rangle$, equals the expected value of fw computed when using the biased pdf, $\langle fw \rangle_b$. The variance of fw computed with the biased distribution is

$$\begin{aligned} \sigma_b^2(fw) &\equiv \int [f(x) w(x) - \langle fw \rangle_b]^2 p_b(x) dx \\ &= \int [f(x) w(x) - \langle f \rangle]^2 p_b(x) dx. \end{aligned} \quad (6.22)$$

Now if $w(x)$ is carefully chosen, the variance computed in Eq. (6.22) will be less than the variance computed by Eq. (6.20). But the converse is also true: a poor choice of $w(x)$, i.e. a poor choice of $p_b(x)$, can increase the variance of the computed expectation.

Examples

We can illustrate these ideas with two specific examples. Gordon (1987) studied the spatial distribution of irradiance at the sea surface resulting from an isotropically emitting point source of bioluminescence embedded at a fixed depth within the ocean. The physically correct pdf that describes an isotropic point source is $p(\theta) = \frac{1}{2}\sin\theta$, where θ is the usual polar angle (see problem 6.1). However, Gordon wisely observed that photons emitted into downward directions, $\hat{\xi} \in \Xi_d$ or $0 \leq \theta < 90^\circ$, are usually "wasted" because of the small chance that they will be scattered into upward directions and eventually reach the surface. He therefore chose to generate the photon polar angles by using a biased distribution

$$p_b(\theta) \equiv \frac{\sqrt{1 - \epsilon^2}}{\pi (1 + \epsilon \cos\theta)},$$

where $0 \leq \epsilon < 1$. This distribution generates a disproportionate number of photons in upward directions, $90^\circ \leq \theta \leq 180^\circ$. It is these upward photons that are most likely to reach the surface and contribute to the estimate of the irradiance. Each photon was given an initial weight

$$w(\theta) = \frac{p(\theta)}{p_b(\theta)} = \frac{\pi}{2\sqrt{1 - \epsilon^2}} (1 + \epsilon \cos\theta) \sin\theta,$$

so that the surface irradiance estimated using the biased distribution would equal the irradiance estimated if the correct, unbiased distribution had been used; recall Eq. (6.21). After generation, each photon was tracked using forward Monte Carlo methods, which was appropriate since this problem involved a point source and a spatially extensive detector (at the sea surface).

The directional biasing just described is an example of what is known as *importance sampling*. Gordon biased the photon generation in favor of those directions that make the most important contribution to the quantity of interest. Since more photons were created heading upward, and thus more reached the surface, the variance of the estimated irradiance is smaller (for a given number of photons generated and traced).

Another example can be based on the same problem. Suppose that now we are interested in the underwater light field set up by the embedded point source. Any photon (i.e. packet of photons) that is emitted into an upward direction $\hat{\xi} \in \Xi_u$, and which makes its way from the source to the surface and into the air without any collisions in the water body, is wasted as far as the underwater light field simulation is concerned. If the source is near the surface in clear water, randomly drawing the photon optical path lengths l from $p(l) = e^{-l}$, as in Eqs. (6.2) and (6.5a), will result in many photons that escape the water body without any interactions.

Let l_s be the optical path length along direction $\hat{\xi} \in \Xi_u$ from the point of generation (or scattering) to the water surface. Now let us sample the path length l from the biased pdf

$$p_b(l) = \frac{e^{-l}}{1 - e^{-l_s}}, \quad 0 \leq l \leq l_s.$$

Use of this distribution guarantees that there will be an interaction somewhere along the path to the surface. Use of this sampling scheme is therefore called the method of *forced collisions*. It is a good example of the principle of never wasting a photon. The corresponding weight is

$$w(l_s) = \frac{p(l)}{p_b(l)} = 1 - e^{-l_s}.$$

The photon weight at the time of generation or scattering must be multiplied by $w(l_s)$ in order to remove the effects of the biased sampling on the estimated quantity of interest. Because this weight is always less than one, its use will always reduce the variance.

Each of the above examples illustrates a general tenet stated by Hammersley and Handscomb [(1964), quoted by permission of Chapman and Hall, Ltd.]: "never sample from a distribution merely because it arises

in the physical context of a problem, for we may be able to use a better distribution in the computations and still get the right answer."

6.4 Summary

We have now seen that Monte Carlo methods provide powerful ways to solve general radiative transfer problems. However, Monte Carlo methods are often computationally very expensive. This is especially true if an accurate estimate of the entire radiance distribution is desired, or if radiometric quantities are desired at large optical depths. In computing upwelling radiances via backward Monte Carlo methods, photon packets must still be initiated in downward directions. Relatively few photons are scattered upward and eventually reach the surface where they are tallied. Thus many more downward photons must be initiated than would be the case if only downwelling irradiances, say, were desired. Moreover, because the number of surviving photons decreases exponentially with optical path length, the number of initial photons required to yield a given number of photons at a detector increases exponentially with the optical path length or depth. Thus computing a quantity at $\zeta = 10$ optical depths, versus the same computation at $\zeta = 5$, requires roughly $e^{10}/e^5 \approx 150$ times as many photons to be traced. For these computational reasons, Monte Carlo methods are best suited to computations of irradiances at shallow depths ($\zeta \leq 5$). However, for such problems, Monte Carlo methods can be very efficient, especially when one accounts for the programmer's time required in writing and debugging the associated computer code. It is also easy to incorporate polarization into Monte Carlo methods; see, for example, Kattawar and Adams (1989). For problems with irregular geometry, there may be no other applicable solution method.

The literature abounds with Monte Carlo methods, variance reduction techniques, and tricks-of-the-trade for speeding up Monte Carlo calculations. However, any further discussion of these matters would take us too far from the subject of hydrologic optics. The books by Hammersley and Handscomb (1964), Spanier and Gelbard (1969), and Cashwell and Everett (1959) each provide good introductions to Monte Carlo methods. A mathematically advanced treatment of the subject is given in Marchuk, *et al.* (1980).

6.5 Problems

- 6.1. Develop a random sampling scheme that could be used in a Monte Carlo simulation to generate photons emitted from an isotropically emitting point source.
- 6.2. Develop an algorithm for randomly choosing scattering angles ψ in a manner that is consistent with the phase function for pure water given by Eq. (3.20).
- 6.3. Develop a sampling scheme to incorporate a Lambertian surface into a Monte Carlo simulation, i.e. show how to determine what happens to a photon packet that intersects such a surface.
- 6.4. We saw in Eq. (6.6) that the average distance traveled by photons is $1/c$. What is the standard deviation of the distance traveled?