Andy Scukanec Lecturer: Steve Marschner

1 Radiometry (cont'd)

Big Fact: Radiance is invariant along a straight line! This assumes no interfering media (ie - a vacuum).

As a proof, look at radiance from a linespace point of view. Assume that S and S' are "small" and "far apart", so that we can use the $\omega = \frac{A}{r^2}$ approximation for solid angle. Light is travelling from left to right, passing through S, then S'. In the image below, S and S' are oriented so that their normals align.

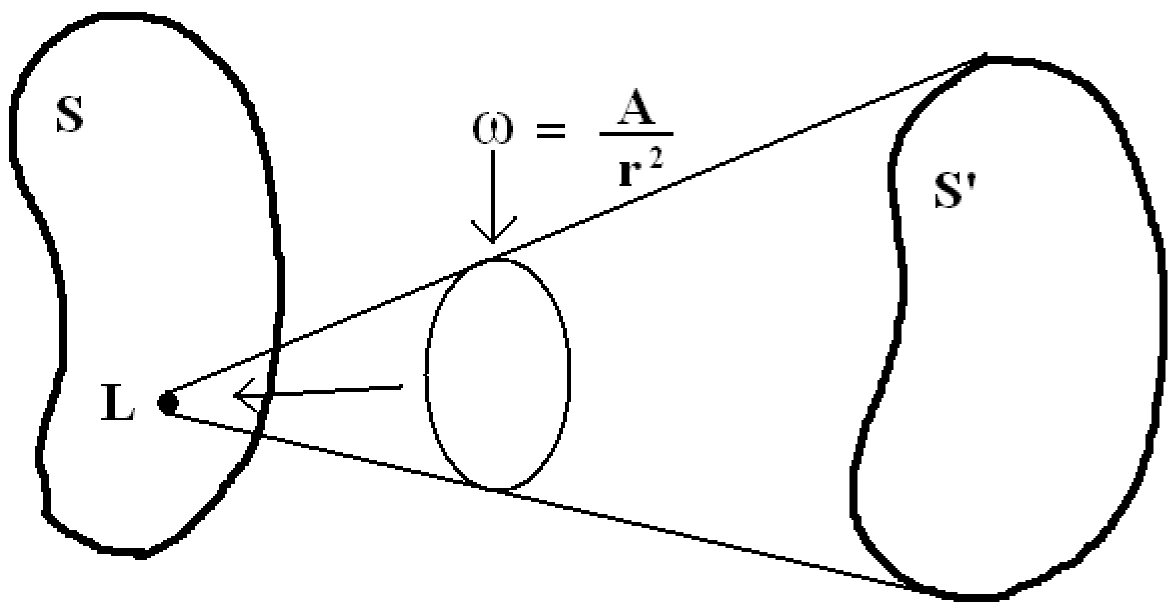

Radiance is defined as the flux divided by the size of the set of lines it is travelling along. In this case, we will look only at the flux that first passes through S, then S'. The set of lines is then all lines that pass through both regions. We can measure the radiance at S, giving us:

$$L = \frac{\Phi^2}{A_S \frac{A_{S'}}{r^2}}$$

or at S', giving us:

$$L' = \frac{\Phi^2}{A_{S'} \frac{A_S}{r^2}}$$

The above set of equations show that L and L' are computed symmetrically with respect to the A_S and $A_{S'}$. The value computed for the radiance measured at S is the same as the amount of radiance measured at S'.

This is an important fact because radiance is essentially what our eyes "see" when we look at a surface. Radiance does not change by bringing the object closer to you (increasing the solid angle), or by making it bigger (increasing the area). This is consistent with the fact that neither of these actions makes a surface brighter. Both of these action increase the amount of power, or flux, of the surface, but not the radiance.

Note: The meaning of $L(x \to \Theta)$ is a limit take as some small area dA approaches a point x, and some small solid angle ω approaches a single direction Θ. More formally:

$$L(x \to \Theta) = \lim_{dA \to x} \lim_{d\omega \to \Theta} \frac{d^2\Phi}{dA\, d\omega\, \cos(\Theta)}$$

Chart from the Preisendorfer paper displaying relationship among radiometric units:

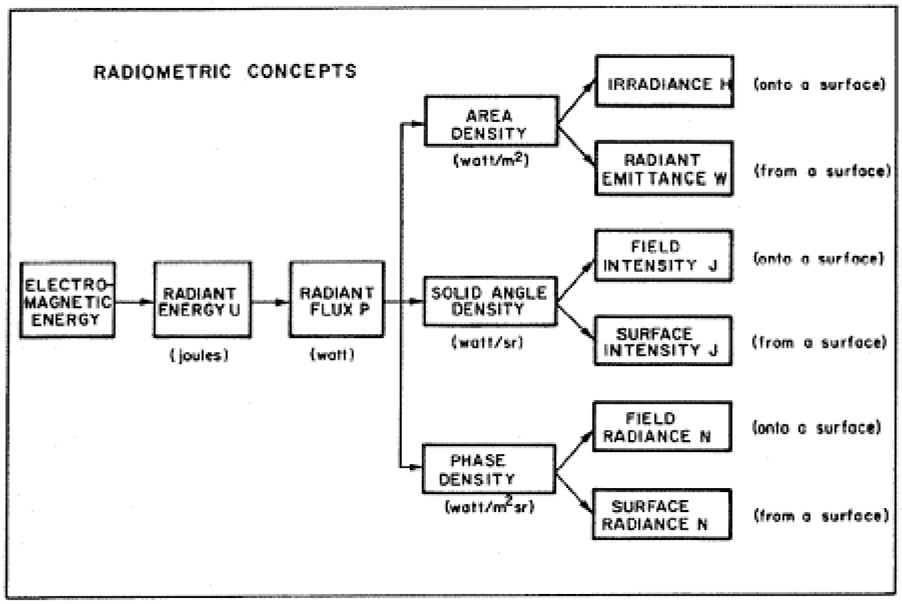

2 Measure Theory view of Integration

Measures are a way of assigning "sizes" to a subset of a domain. The usual way in which this is done is to define the value of some basic unit of the domain (e.g. rectangles in the plane) and approximate any subset of interest with many of these units.

Integration is always done with respect to a measure, μ, and most of the time this is so obvious that we don't even think about it. Primary examples are integration over the real line with respect to the measure of "length", or on the plane with respect to the measure "area". The symbol dx is really a shorthand for the measure theoretically correct way of writing an integral:

$$F(x) = \int f(x)\, dx = \int f(x)\, d\mu(x)$$

$$F(x, y) = \int f(x, y)\, dA = \int f(x, y)\, d\mu(x, y)$$

Where μ is some arbitrary measure

Two commonly used measures in radiometry are: $\sigma(D)$ – the solid angle of D, and $\mu(D) \approx \sigma(D)\cos(\theta)$ – the projected solid angle of D.

Distributions, such as probability distributions, are also defined with respect to some measure.

- distributions vs. densities
- occurrence of Dirac δ functions when distributions don't have densities

Normal PDF's (probability distribution functions) prescribe an infinitely small chance that a variable will take any particular value. For instance, assuming a uniform distribution over some range $[0, L]$, the pdf is:

$$f(x) = \begin{cases} \frac{1}{L} & x \in [0, L] \\ 0 & \text{otherwise} \end{cases}$$

The chance that f takes on any value is not $\frac{1}{L}$, but rather 0, since there are infinitely many values that f may take. Integration will save us here though, as the chance that f will assume some value in the range $[a, b]$ is $\int_a^b f(x)\, dx$. This technique works fine until we encounter a distribution that *does* actually associate a non-zero probability with a single value. For instance, suppose that the probability g would be 2 is $\frac{1}{2}$, and the chance that g would be any other value in the range $[0, 4]$ is evenly distributed amongst the remaining values. Our PDF would look just like f, but with an infinite spike at 2. How do we define density for such a distribution?

Rather than looking at the function as a PDF, we'll look at its integral, since the integral of g (rather than g itself) is what is well defined. This integrated version of g is called its cumulative distribution function, or CDF, and is defined as $G(x) = \int_{-\infty}^{x} g(x)\, dx$. This integral is computed using Dirac δ functions, which act like on/off flags for all "distribution spikes":

$$\int_D f(x)\, d\mu(x) = \delta_a + \int_{D \setminus a} f(x)\, d\mu(x)$$

3 Units from an Integral View

- Radiance $= \frac{W}{m^2\, sr} = L(\vec{x}, \vec{\omega})$.
- Irradiance $= \frac{W}{m^2} = E(\vec{x}) = \int_{H^2} L(\vec{x}, \vec{\omega})\, d\mu = \int_{H^2} L(\vec{x}, \vec{\omega})(\vec{n} \cdot \vec{\omega})\, d\sigma$
- Intensity $= \frac{W}{sr} = I(\vec{\omega}) = \int_A L(\vec{x}, \vec{\omega})\, dA^{\perp} = \int_A L(\vec{x}, \vec{\omega})(\vec{n} \cdot \vec{\omega})\, dA$
- Flux = Power $= W = \Phi = \int_\Omega \int_A L(\vec{x}, \vec{\omega})\, dA\, d\mu = \int_\Omega \int_A L(\vec{x}, \vec{\omega})(\vec{n} \cdot \vec{\omega})\, dA\, d\sigma$
- Note: S^2 = Sphere, H^2 = Hemisphere

4 Examples

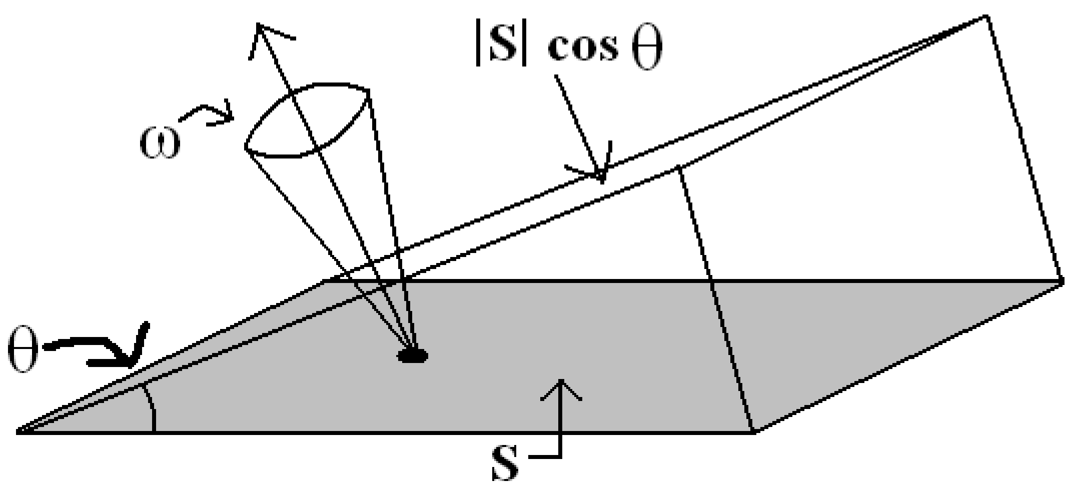

In the image above, we see a lambertian emitter with constant radiance L_0 for all directions and locations. Suppose we'd like to treat this emitter as a point source. What is its intensity distribution? We can get the total intensity in some direction ω for the whole area by integrating the radiance (which is always the same) over the projection of that area normal to the direction ω. The intensity then is given by the simple integral:

$$I(\vec{\omega}) = \int_S L_0 \, dA^{\perp} = \int_S L_0(\vec{n} \cdot \vec{\omega}) \, dA = L_0|S|(\vec{n} \cdot \vec{\omega})$$

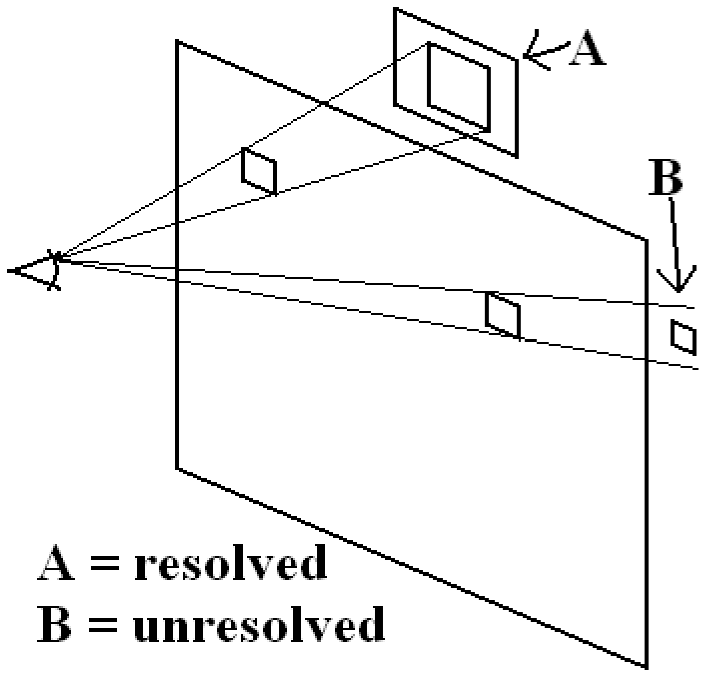

In the above example, we see a problem associated with discretizing a radiometry problem. We have a camera looking through and image plane with boundary lines drawn to indicate the window of view that the camera would see if it were looking through the two indicated pixels on the plane. Behind the image plane, we have two surfaces A and B. A is positioned such that it does not fit entirely into the view of one pixel. B is oriented so that it does. The other thing to keep in mind is that this camera can essentially record only one value for each pixel (lets leave color out of the picture for now), representing the "brightness" of the pixel.

The two situations are marked as "resolved" and "unresolved". The "resolved" situation is the normal one, in that increasing the surface area does not result in a brighter pixel. This correctly simulates the fact that radiance does not increase as a result of increased surface area. The "unresolved" case is one where the solid angle of a pixel in the image fully encloses the surface area of the surface. This effectively makes B a point source. Increasing the size of B does not result in an increase in resulting pixel coverage in the image, but rather results in a brighter pixel at the same spot. This, this pixel measures the intensity.

Side note: it does not make sense to talk about the radiance of the point source, as a point source has no area, making its radiance infinite. Instead, an intensity is usually associated with a point source.

5 Other Related Points

There is a closely related field to Radiometry, which is Photometry. While radiometry is the study of light, photometry is the study of light detection by people.

Up until now, a dimension of information has been completely lacking from our discussion of light - wavelength, or color. All of the quantities (radiance, irradiance, etc) can also be divided up depending on the color of incident light. This only really becomes an issue when dealing with real measurement tools in a lab, since most have associated sensitivity curves.

The number given back by a receptor is really the integral of that receptors sensitivity curve against the incoming spectral distribution. The lecture slides give the sensitivity curve for the human eye, and we can use this curve as an example. Pretend you were completely color blind and could only see shades of light and dark. This sensitivity curve then would dictate the brightness of any particular direction. This curve is given by the equation $\bar{Y}(\lambda)$, where λ is the wavelength.

Given some spectral distribution g (in other words the amount of light for any particular wavelength), the brightness humans see is given by the integral:

$$\int \bar{Y}(\lambda) g(\lambda) \, d\lambda$$

The domain of the integral is technically all valid wavelengths, but that is an infinite domain, and the computation is usually done just over the support of the response function.

Although human eyes have their own response function, so does every other photometric device. This response function needs to be taken into account during calibration etc.

CS667 **Lecture 3: The BRDF, Light Transport in a Vacuum** 3 February 2004

Jon Moon Lecturer: Steve Marschner

1 Radiometry examples

To conclude discussion of radiometry, we present three example exercises.

1.1 Solid angle of a disc or sphere

Given a disc or sphere of radius R with center at distance r from the point of interest, what is the solid angle it subtends?

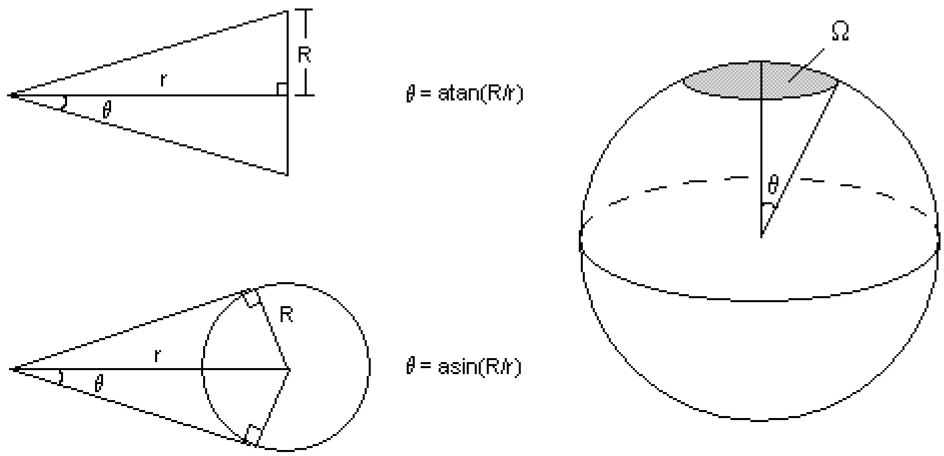

The solid angle of this region is equal to the area of the cap of the unit sphere, Ω. If we define the measure σ as

$$\sigma(\Omega) = |\Omega|$$

then the area is

$$\int_\Omega d\sigma = \int_0^{2\pi} \int_0^{\theta} sin\ \theta\ d\theta\ d\phi$$

Note that θ is the angle down from the top of the sphere, and ϕ is the angle in the equatorial plane, measured counterclockwise. This will be the convention used for the duration of this course. We evaluate the integral as follows:

$$= 2\pi \int_0^{\theta} sin\ \theta\ d\theta = 2\pi[-cos\ \theta - (-cos\ 0)] = 2\pi(1 - cos\ \theta)$$

Note that the quantity $(1 - cos\ \theta)$ is the height of the spherical cap, and that this result generalizes for arbitrary bands around the sphere:

$$\text{Area of band} = 2\pi(cos\ \theta_1 - cos\ \theta_2)$$

1.2 Radiance of the sun, approximately

This example gives an idea of the magnitude of radiance values in the real world. Consider the irradiance of the sun on a flat surface at noon.

- known irradiance of the sun: 500 W/m^2
- known angular subtense: $1/2^o$ or 1/100 radian

We estimate the solid angle by treating the spherical cap as a disc, with diameter 1/100 on the unit sphere and thus has area $\pi/40000$ steradian. Since the illumination is perpendicular, this is also the projected solid angle. Since radiance is irradiance per unit projected solid angle, we find

$$L = \frac{500 \text{ W/m}^2}{\pi/40000 \text{ sr}} \approx 6 \times 10^6 \text{ W/m}^2\text{sr}$$

1.3 Reflection from a Lambertian reflector

A Lambertian reflector reflects a fraction R of its incident flux, emitting it uniformly in all directions. That is,

$$\text{(radiant exitance) } M = R * E \text{ (irradiance)}$$

Recall also that

$$M(\mathbf{x}) = \int_{\mathbb{H}^2} L(\mathbf{x}, \omega)\, d\mu(\omega) = \int_{\mathbb{H}^2} L\, d\mu = L \int_{\mathbb{H}^2} d\mu = \pi L$$

where μ is the projected solid angle measure. Combining these two results yields

$$\pi L = R * E \quad \text{and so} \quad \mathrm{L} = \frac{\mathrm{R}}{\pi}\mathrm{E}$$

2 The Bidirectional Reflectance Distribution Function (BRDF)

2.1 Definition

Surface reflection is an operator, taking as input an incident radiance distribution L_i and producing a reflected radiance distribution L_e as output. That is, $L_e = \mathcal{R}(L_i)$.

2.2 Linearity of the BRDF

A key property of $\mathcal{R}$ is linearity: $\mathcal{R}(A + B) = \mathcal{R}(A) + \mathcal{R}(B)$ This linearity allows us to treat a radiance distribution A as a sum of small light sources A_j, each contributing radiance L_j from solid angle Ω_j around ω_j, and have $\mathcal{R}(A) = \sum_j \mathcal{R}(A_j)$

This means that to predict the reflection of any radiance distribution, we only need to know the reflection for small sources. This is exactly what the BRDF tells us: the reflected distribution from a small source. We can define the BRDF f_r as the exitant radiance in a direction per incident radiance from a direction per unit projected solid angle. That is,

$$f_r(\omega_i, \omega_r) = \frac{L_r}{L_i} / \mu(\Omega_i)$$

Equivalently,

$$L_r = f_r(\omega_i, \omega_r)\, L_i\, \mu(\Omega_i)$$

For our sum of small light sources A_j, we have

$$\mathcal{R}(A)(\omega_r) = \sum_j f_r(\omega_j, \omega_r)\, L_j\, \mu(\Omega_j)$$

Or as the limit as Ω_j gets small:

$$L_r(\omega_r) = \int_{\mathbb{H}^2} f_r(\omega_i, \omega_r)\ L_i(\omega_i)\ d\mu(\omega_i)$$

Two other ways to think about the BRDF are:

- $f_r(\cdot, \omega_r)$ represents the "sensitivity" to radiance per unit projected solid angle
- $f_r(\omega_i, \cdot)$ represents the reflected radiance for a collimated incident beam.

2.3 Properties of the BRDF

It should be obvious that a BRDF needs to conserve energy: the flux leaving a surface (radiant exitance) must be $\leq$ the flux incident on the surface (irradiance) for all incident distributions:

$$\int_{\mathbb{H}^2} L_r\ d\mu \leq \int_{\mathbb{H}^2} L_i\ d\mu$$

This is true if and only if it holds for collimated illumination:

$$\int_{\mathbb{H}^2} f_r(\omega_i, \omega_r)\ d\mu(\omega_r) \leq 1 \tag{1}$$

The forward implication is obvious, and the reverse implication is shown via integration:

$$M = \int_{\mathbb{H}^2} L_r\ d\mu = \int_{\mathbb{H}^2} \int_{\mathbb{H}^2} f_r(\omega_i, \omega_r)\ L_i(\omega_i)\ d\mu(\omega_i)\ d\mu(\omega_r)$$

Swapping the order of integration, we have

$$\int_{\mathbb{H}^2} L_i(\omega_i) \underline{\int_{\mathbb{H}^2} f_r(\omega_i, \omega_r)\ d\mu(\omega_r)}\ d\mu(\omega_i)$$

By (1), the underlined integral must evaluate to ≤ 1. Thus

$$M \leq \int_{\mathbb{H}^2} L_i(\omega_i)\ d\mu(\omega_i) = E$$

as claimed.

A less obvious property is Helmholtz reciprocity, which states that the BRDF has a symmetry with respect to swapping its arguments:

$$f_r(\omega_i, \omega_r) = f_r(\omega_r, \omega_i)$$

The physical interpretation for this reciprocity is that the sensitivity distribution looks like the radiance distribution:

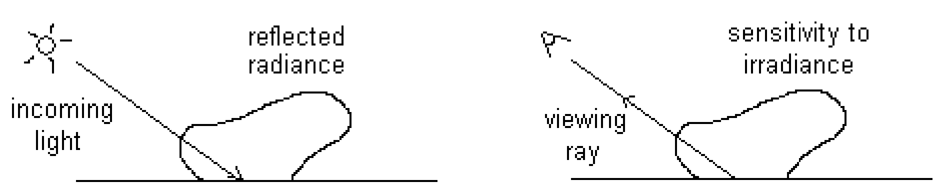

This is a very important property, and is fundamental to many rendering algorithms.

3 Light Transport in a vacuum

Consider the transport of light through a vacuum, by which we mean there is no participating medium. Take the following as ground rules:

- The scene is composed of surfaces floating in a vacuum. Let all the surfaces considered together be a piecewise smooth surface (a 2-manifold) $\mathcal{M}$.
- Reflection occurs pointwise, as all surfaces are opaque and obey valid BRDFs.
- The output we are interested in - the camera image - is just a set of averages over the light reflected from the scene surfaces, with one measurement made per pixel.
- There is an enclosure surrounding all of $\mathcal{M}$, to avoid special cases for the background.
- All light in the scene is initially emitted from the surfaces

Also define:

- $L_e(\mathbf{x}, \omega_e)$ is the exitant radiance from point $\mathbf{x} \in \mathcal{M}$ to direction ω_e.
- $L_e : \mathcal{M} \times \mathbb{H}^2 \rightarrow \mathbb{R}$
- $L_i(\mathbf{x}, \omega_i)$ is the incident radiance on point $\mathbf{x} \in \mathcal{M}$ from direction ω_i
- $L_i : \mathcal{M} \times \mathbb{H}^2 \rightarrow \mathbb{R}$
  note that ω always faces away from the surface!
- $f_r(\mathbf{x}, \omega_i, \omega_e)$ is the BRDF at point $\mathbf{x}$
- $f_r : \mathcal{M} \times \mathbb{H}^2 \times \mathbb{H}^2 \rightarrow \mathbb{R}$

From all this, the BRDF definition gives:

$$L_e(\mathbf{x}, \omega_e) = \int_{\mathbb{H}^2} f_r(\mathbf{x}, \omega_i, \omega_e)\ L_i(\mathbf{x}, \omega_i)\ d\mu(\omega_i)$$

or

$$L_e = \mathbf{K} L_i \quad \text{where } \mathbf{K} \text{ is the reflection operator}$$

We can think of $\mathbf{K}$ as the whole surface reflectance for all points everywhere rolled into a single linear operator. We also include emittance, which adds to the reflection:

$$L_e = \mathbf{K} L_i + L_e^0$$

Where $L_e^0(\mathbf{x}, \omega_e)$ is the radiance emitted from point $\mathbf{x}$ in direction ω_e.

At this point, this is just a restatement of surface reflection. To make a solvable equation we need to relate L_i to L_e. Fortunately, because we are considering light transport in a vacuum, they are the same function - only with permuted domains. That is, $L_i(\mathbf{x}, \omega) = L_e(\mathbf{y}, -\omega)$ for the point $\mathbf{y}$ that is visible from $\mathbf{x}$ when looking in the direction ω. This is ray casting, essentially.

We can then define a transport operator $\mathbf{G}$ such that $L_i = \mathbf{G} L_e$:

$$(\mathbf{G} L_e)(\mathbf{x}, \omega) = L_e(\psi(\mathbf{x}, \omega), -\omega)$$

Where ψ is the ray casting function, with $\psi(\mathbf{x}, \omega) = \mathbf{y}$, and $\psi : \mathcal{M} \times \mathbb{H}^2 \rightarrow \mathcal{M}$

Finally, we can substitute this into our surface reflection equation, resulting in

$$L_e = \mathbf{K}\mathbf{G} L_e + L_e^0$$

This is a very compact way to write down the rendering problem and to expose the algebraic structure. As a final note, let **1** be the identity operator. Then we have

$$\mathbf{1}L_e - \mathbf{KG}L_e = L_e^0$$
$$L_e = (\mathbf{1} - \mathbf{KG})^{-1}L_e^0$$
$$L_e = L_e^0 + \mathbf{KG}(L_e^0 + \mathbf{KG}(L_e^0 + ...))$$

Which is an intuitive representation for recursive ray tracing.

Next lecture we will examine Kajiya's formulation of the rendering equation using areas.