Rendering light propagation in an optically thick medium

Thorsten Renk, 4.10.2011

1 Basic definitions

We see non-luminous objects because photons from a light source reach the object, are scattered there and subsequently reach our eye. If there is no scattering between the object and the eye, the distance we can see is almost unlimited (we easily see objects in space thousands of lightyears distant). What limites the view range are scatterings of light propagating between the seen object and the eye.

Assume that the mean distance for a photon to travel without scattering is λ. This is sometimes also called optical depth or optical thickness. We call a medium *optically thin* if the distance d between object and observer is smaller than λ (i.e. most photons reach us without scattering) and optically thick if $d \gg \lambda$.

Note that any scene to render is almost never purely optically thick or thin — even in dense fog, objects can be sufficiently close so that they are seen clearly. Optically thick and thin are just useful approximations to make.

In an optically thick medium, photons have typically been scattered a number of times before they reach the eye. As a consequence, properties of any individual scattering are largely lost. In the optically thin single scattering approximation, directional (Mie) scattering is a major effect which appears much stronger at small angles with the sun than at large angles. This is not so in optically thick media — consider a rainy sky: It's almost impossible to tell where the sun is above the clouds, because all directional information is lost.

2 Light attenuation

Neglecting any dependence on wavelength, photons propagating directly from an object to the eye are attenuated as

$$dN(x)/dx = -N(x)\sigma\rho(x) \tag{1}$$

where $N(d)$ is the number of photons left after covering distance d, σ is the scattering cross section and $\rho(d)$ is the density of scatterers at position d. For a constant density, this equation is solved by an exponential function with an initial photon number (i.e. light intensity)

$$N(d) = N_0 \exp(-\sigma\rho d) \equiv N_0 \exp(-d/\lambda) \tag{2}$$

where the last expression utilizes the definition of the mean free path $\sigma\rho\lambda = 1$. Thus, if we're two optical depths away from an object, we only see $\exp(-2) \approx 13\%$ of the originally emitted light, the rest is scattered out. If this were all that happens, objects in fog would be much darker. However, in actual reality we see in-scatterd light instead.

3 Light diffusion

What does the rest of the light in the scene do? Assume a scene as depicted in Fig. 1:

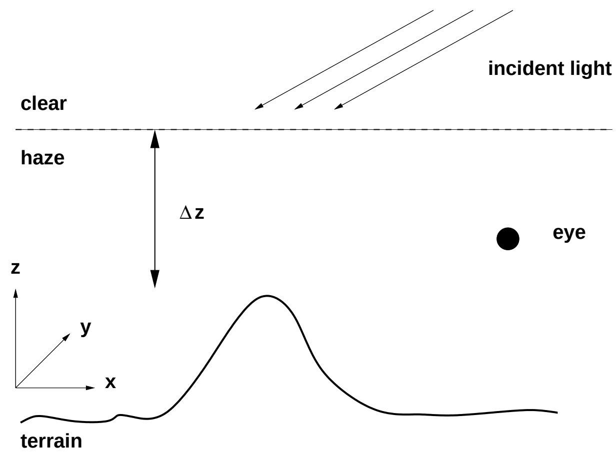

Figure 1: The basic scene of a ground haze layer

Light reaches from the sun to the top of the layer. Since the layer is optically thick, it scatters multiple times in the layer, at which point each photon essentially undergoes a random walk. The mean intensity of light is then described by a diffusion equation

$$\frac{dN(\mathbf{r}, r)}{\partial t} = \nabla \left(D(N, \mathbf{r}), \nabla N(\mathbf{r}, t) \right) \tag{3}$$

where $N(\mathbf{r}, t)$ is the number of photons at spatial position $\mathbf{r}$ and time t and D is a diffusion coefficient which can be computed from the properties of an individual scattering and the local medium density. Assuming that D is not position dependent *inside the haze layer* but vanishingly small outside and some number inside, we can simplify the equation by pulling the constant D out. Furthermore, we are not interested in time dependence — photons still move with the speed of light, so the intensity distribution of light in the layer will adjust practically instantaneously to any changes, thus we're only interested in the static equilibrium limit and can drop the time derivative on the left hand side, hence we get the much simpler

$$D \nabla^2 N(\mathbf{r}) = 0 \tag{4}$$

For an infinite layer symmetry reasons argue that the only gradient can be perpendicular to the interface to clear air, i.e. along the z direction in the figure, i.e. the equation collapses to a 1-dim

$$D\nabla^2 N(z) = 0 \tag{5}$$

This is again solved by an exponential

$$N(z) = N_0 \exp(-\delta \cdot \Delta z) \tag{6}$$

with $D = \delta^2$ and Δz measured from the upper layer edge downwards. For simplicitly, we may assume that $\Delta \sim 1/\lambda$ as both quantities are related to the same scattering process (but in reality, δ knows about directional scattering). Thus, we know the intensity of light at any point inside the layer, and by definition of the optically thick condition they do not have a preferred direction, i.e. light radiates from this point with given intensity into all directions equally.

4 Bringing it together

Thus, what do we see? The basic problem is depicted in Fig. 2:

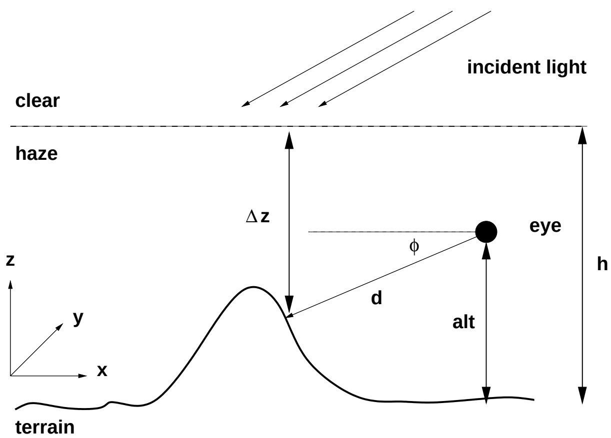

Figure 2: Viewing an object inside the haze

We see the light directly transmitted from an object across the distance d in addition to all light that is scattered into the path and hasn't been scattered out. Thus, a fraction $T = \exp(-d/\lambda)$ is direct light from the object, the rest is background light.

Consider one particular point d_0 on the line between eye and observed object. Using the labelling in the figure, the Δz at this point is $(h - alt) + d_0 \sin(\phi)$ (if ϕ is taken to be positive looking up and negative looking down) and the corresponding light intensity is $\exp\left[-((h - alt) + d_0 \sin(\phi))/\lambda\right]$ with the fraction of this light actually reaching the eye being $\exp(-d_0/\lambda)$. The total amount of stray light being inscattered and seen by the eye (i.e. not outscattered again) is then the integral along the path

$$N_{stray} = \int_0^d dd' \exp\left[-\frac{(h-alt)+d'\sin(\phi)}{\lambda}\right] \exp\left[-\frac{d'}{\lambda}\right]. \qquad (7)$$

This looks complicated, but has a very easy approximate solution which is appropriate for most cases relevant in practice. This solution is: Take the stray light at the value of one optical thickness λ along the path (or the layer edge or the terrain, whatever comes first) as the solution to the integral. Thus, the part of light that is not transmitted should be taken with the intensity

$$N_{stray} = N_0 \exp\left[-\frac{(h-alt)+\lambda\sin\phi)}{\lambda}\right] \qquad (8)$$

and the light reacing the eye is then $T \cdot \text{direct light} + (1-T) \cdot N_{stray}$.

In the general case, the path must be divided. The following situations can occur and must be dealt with separately:

- eye inside haze, object inside haze: this has been discussed above
- eye inside haze, object (or skydome) above haze: in this case, only a portion d_{haze} of the total path contributes to attenuation
- eye outside haze, object below haze: in this case, only a portion d_{haze} of the total path contributes to attenuation
- eye outside hze, object outside haze: there is no stray light from the optically thick limit and the whole computation must be done in the optically thin limit

In general, attenuation is multiplicative, i.e. the total transmission through two layers A and B is $T_A \cdot T_B$, i.e. as long as at least one layer is optically thick, the results will be determined by that layer.

5 Intensity perception

Using an exponential attenuation of intensity gives the result of a photon flux measurement correctly, but doesn't reproduce the perception of the human eye. Here the Weber-Fechner law states that the sensitivity to differences in intensity of a stimulus is approximately proportional to the magnitude of the stimulus. Thus, if light at the top of the haze layer has the $(rgb)_{light}$ values of the primary light source, the color of a pixel in the layer should *not* be $(rgb)_{light} \cdot N_{stray}$ but instead

$$(rgb)_{haze} = (rgb)_{light}(1 + a\ln(N_{stray}/N_0)) \qquad (9)$$

where a is a parameter to be tuned by comparison with real situations. The net result of taking a logarithm here is a cancellation of the exponential, i.e.

$$(rgb)_{haze} = (rgb)_{light} \cdot (1 - a \cdot \Delta z/\lambda) \quad (10)$$

For a more detailed modelling, it is of course also possible to compute the attenuation explicitly for each color channel with slightly different attenuation constants.

6 Some Flightgear specifics

Cloud layers do attenuate light, but they are not done using volumetric haze, so the information must be inserted manually. Currently we do this via */rendering/scene/scattering* which is supposed to stand for a Weber-Fechner corrected reduction of the light intensity below a cloud layer.

In principle, */rendering/scene/saturation* does the same thing, but that's a more dangerous beast as it directly modifies the intensity of the primary light source - it will in essence dim the sun itself, and that is only a good idea if the visibility is generally poor and the sky can't be seen at all.