Efficient Rendering of Glowing Lights

Woodley Packard

June 9, 2004

Introduction

Consider a point light source in an atmospheric medium. The high-level approach used by ray-tracers to render this scene is to cast rays through the volume and calculate the radiance reaching the eye along each ray.

Hence forth let us assume we are considering a single ray. The standard approach for calculating the radiance along the ray is to estimate the value of the volume rendering equation:

$$L = \int_0^\infty e^{-\int_0^t \rho(s)ds} S(t)dt$$

where $S(t)$ represents the sum of the in-scattered light, ρ represents the density at a point, and the emissivity of the point t on the ray. The most common algorithm is numerical integration, i.e. to simply march along the ray, calculate the integrand at each step, and sum the results. The value of $S(t)$ is however a complex quantity which can be difficult to calculate exactly. One approach is to integrate the incoming radiance over all directions around the point t and add the emissivity, but this presents a chicken-and-egg problem: the radiance is exactly what is being computed to begin with. Therefore, after a certain number of recursive steps like this, an approximation is typically made by assuming, for example, that the in-scatter is equal to some constant ambient term. The overall integration process can be quite costly, since it requires time exponential in the number of scattering steps that are required.

Low Density Media

In the case of simple atmospheric scattering, the density of the participating medium is typically extremely low. If the density is constant (i.e. the medium is homogeneous), then the volume rendering equation reduces to:

$$L(x,\omega)=\int_0^\infty e^{-\rho t}S(t)dt$$

Assume further that light is equally likely to be scattered in all directions, i.e. the phase function is constant. Then we can write $S(x)=\int_{S^2}\rho L(x,\omega')d\omega' + \epsilon = \rho\int_{S^2}L(x,\omega')d\omega'+\epsilon$, where ϵ represents the emissivity at x.

Let us now consider the specific case-at-hand: a light-source in pure air, at night time. In this case, it is clear by inspecting any photograph that the radiance in any direction other than that of the light source is at an order of magnitude less bright than the the direct radiance from the light source. Since we are considering point-lights, we need to assume they are "infinitely bright" in some sense, since otherwise they would not produce any contribution to the final scene. Let the light have brightness density A (i.e. a 1×1 unit square area light would appear as bright as our light if it had brightness A). Thus we get $S(x)\approx \rho\frac{A}{r(x)^2}$, where $r(x)$ represents the distance from x to the light source. Plugging this into the volume rendering equation, we get:

$$L(x,\omega)=A\rho\int_0^\infty e^{-\rho t}\frac{dt}{r(t)^2}$$

This is unfortunately not an easy formula to integrate.

Closed Form Approximation

We circumvent this problem by approximating the distance fall-off independently from the scattering. This clearly involves the assumption that our light-source p is close enough that the distance fall-off near p is negligable, since otherwise the two effects are not independent. In our case this was an acceptable assumption. Instead of integrating from 0 to ∞, we can assume the ray collides with some surface at distance T. After our separation assumption, we get:

$$L\approx A\rho\int_0^T\frac{dt}{r(t)^2}+e^{-\rho T}L(x(t))$$

where $L(x(t))$ represents the radiance computed from the distant intersection point. The term $e^{-\rho T}$ is then equivalent to a standard homogeneous medium "fog" effect, so we will drop the distant surface term from now on.

The problem is then to compute $A\rho\int_0^T\frac{dt}{r(t)^2}$. We need to define some coordinates. The relevant points are the location of the light and the line of the ray. These together define a 2-dimensional subspace of $\mathbb{R}^3$ (assuming the light is not on the ray, in which case the brightness should be infinite). It is convenient to constrain all of our work to the plane $\mathbb{R}^2$. Let us pick coordinates such that the segment of the ray from parameter 0 to T coincides with the y axis, from position $(0, a)$ to $(0, b)$. We can pick a translation compatible with this assumption which puts our light source on the x-axis at some point $(p, 0)$. Ignoring the constant terms, our integral takes the form $\int_a^b \frac{dy}{p^2+y^2}$. This integral is easily conquered using trigonometric substitution. We obtain the formula:

$$L \approx A\rho \left[\operatorname{atan}\frac{a}{p} - \operatorname{atan}\frac{b}{p} \right]$$

Calculating a, b, p

There are simple formulas for a, b, and p in terms of standard world-space coordinates. Let P, O, and D be vectors representing the position of the light, the origin of the ray, and the direction of the ray, respectively. Notice that lowercase p is just the distance from P to the ray at its closest point, $-a$ is the distance from O to that point, and b is the distance from the end of the ray segment to that point. Denote that point by Z. We then get:

$$Z = O + D\left(D \cdot (P - O)\right)$$

$$p = ||P - Z||, \qquad a = -||O - Z||, \qquad b = ||O + TD - Z||$$

More general light sources

Since light obeys the superpositionality principle, we can use this approximation to render glows around arbitrarily shaped light sources by simply integrating the single-point form over all the points on a linear or area light source. Monte-Carlo integration or uniform stepping both work. The Monte-Carlo method tends to introduce undesirable noise in the otherwise smooth atmosphere, although it will converge to the correct average value more quickly. We chose to use the uniform stepping approach, since it causes extremely low variance in the final image, even though it introduces some bias since it doesn't sample all points of large light source. This bias is arbitrarily small compared to the error introduced through the many assumptions that were made to derive the approximation, so the smooth appearance outweighed the benifits of Monte Carlo for our purposes.

The resulting images look quite convincing. The size of the glow is controllable either by the brightness of the light or by the density ρ.