Full-Spectral Rendering of the Earth's Atmosphere using a Physical Model of Rayleigh Scattering

John Irwin

Computer Graphics Unit, Manchester Computing,
University of Manchester, Manchester M13 9PL

WWW: http://info.mcc.ac.uk/CGU/staff/irwin/atmos/

Abstract

The results of a project which developed a rendering system to produce synthetic images of the Earth and its atmosphere using purely physical methods are presented. Two distinct features of this system are emphasised: (1) physically realistic models of the atmosphere's density distribution, the spectral energy distribution of solar radiation impinging on the atmosphere and the molecular scattering of light, which is accurately described by Rayleigh's theory, and (2) the use of wavelength sampling and reconstruction techniques for calculating spectra and colour conversion procedures to produce displayable values. Radiances leaving the atmosphere are calculated using a pseudo-parallel numerical integration technique on a restrictive version of the radiative transfer equation. Results are presented in a number of images of the Earth from various view-points including a sequence depicting the changing colour and brightness of the sky at sunset.

Keywords atmospheric scattering, realistic rendering, spectral sampling.

1 Introduction

Despite the great advances which have been made in the field of realistic image synthesis, there still remains a great reluctance by developers to incorporate a realistic colour model in their rendering systems. The norm these days is to use the Red-Green-Blue (RGB) system for defining colours which, if nothing else, is very easy to use and implement. The motivation for utilising such a model seems to arise simply from the need to display computed images on a colour CRT which invariably use a red/green/blue phosphor set.

However, there are much more sophisticated colour models available[2] which have existed for a long time and which have even been specialised for use in generating realistic computer graphics images[8]. The CIE colour system[12] is based upon the requirement that any colour which the human eye can see will be represented within this scheme. As such, the CIE system should be made the basis for the calculation of colour values in any rendering system which purports to generate realistic images. In practice, the CIE XYZ tristimulus values of a colour act as intermediaries between the colour's computed spectrum and the corresponding RGB values used to drive a graphics monitor. There are well defined procedures available for such colour conversions and these will be covered later.

We see that the major thrust of a realistic image rendering system will be the calculation of the spectrum of light emanating from various directions towards a view-point. The spectrum is represented by samples taken at a number of discrete wavelengths in the visible region, from which the corresponding continuous distribution is reconstructed using an interpolative technique.

To illustrate these colour techniques, a specialised rendering system has been developed which produces synthetic images of a well known natural phenomenon - the scattering of sunlight by the Earth's atmosphere. This effect is very appropriate since we know from experience that the scattering of light by atmospheric gases and aerosols is strongly dependent on the wavelength of light. This dependence gives rise to such visual effects as the blue colour of the clear daytime sky and the reddened appearance of the Sun at sunset. Our aim is to reproduce these colours using a physically-based model.

Following a few observations in the next section concerning the efforts of other researchers in the field of atmospheric modelling for computer graphics, we will begin to see in detail the construction of the model atmosphere and the procedures needed to calculate the colour and intensity of light along any view direction and from any view-point. These procedures, as we have seen, involve the computation of the spectral distribution of scattered light, the conversion to a physical 3-component colour space and the calculation of monitor RGB values.

Following this the results of this project are presented in a number of images of the Earth computed from view-points both inside and outside the atmosphere. For the former, a sequence of images, with a view-point on the Earth's surface, has been computed which capture the changing colour and brightness of the sky at sunset.

Finally, in the last section, suggestions are made for improvements to the atmospheric model which would, in particular, enhance the visual realism of the sunset images.

1.1 Previous work

In the area of atmospheric illumination modelling various attempts have been made to simulate the light from the sky and the Sun considering the scattering and absorption of light in the Earth's atmosphere. Of particular significance here are the works of Klassen[5], Kaneda[4] and Nishita[10].

Klassen's project used a high-resolution spectral sampling technique to calculate the colour of the sky and Sun as seen from an observer fixed to the Earth's surface. Though his model atmosphere was of fixed density, it included a ground based layer of haze to simulate the effects of atmospheric aerosols. Together with a proper consideration for the spectral distribution of the light from the Sun, Klassen's work produced some very realistic results, illustrating the efficacy of using spectrally-based colour models.

Kaneda *et al* developed a rendering system which was used principally to create images of complex outdoor scenes illuminated by direct sunlight and from skylight. For the skylight component a model devised by the CIE was employed, which provided ground-based luminances for various directions in the sky and positions of the Sun. No attempt was made to simulate this directly by accounting for light scattering in the atmosphere. Moreover, their calculations were performed using the RGB colour model and it is believed that the Sun's illumination was modelled merely as a white-light source.

Nishita *et al* produced an atmospheric scattering model very similar to the one presented here, but again they used nothing more sophisticated than the RGB colour model. Although their algorithm incorporates effects due to atmospheric haze and the reflection of skylight from the Earth's surface and ocean water, it must be questioned whether such complexity is warranted given their use of a naive colour model for the calculation of light intensities.

2 Computing single-scatter radiances

The Earth's atmosphere is modelled as a spherically-symmetric envelope of pure air extending from the Earth's surface (at a radius of 6371km) up to a height of 200km. Between these limits the atmosphere's mass density falls exponentially with a *scale height* of 8.4km. Hence, at its outer limit its density is only ~5 x 10^{-11} of that at the surface and is assumed negligible. Note that a fixed height is used principally to contain integrations between finite limits, which aids numerical evaluation. However, the height is set sufficiently high such that the density is all but zero.

For pure air a single-scattering model is assumed together with Rayleigh parameters for the scattering coefficient and phase function.

The light from the Sun is modelled with a realistic spectral distribution; values have been taken from Wyszecki and Stiles[12]. We also take the Earth's shadow into account when we determine light paths through the atmosphere and, since we are assuming a parallel light source to represent the radiation from the Sun, the shadow is modelled as a half-cylinder, with a radius equal to the Earth's, extending from the Earth's centre in a direction away from the Sun.

2.1 An integral of the radiative transfer equation

We now use this simplified model of the clear terrestrial atmosphere to calculate the radiance resulting from the single-scattering of sunlight into a particular line of sight. We could start from the radiative transfer equation and find solutions for this particular situation, but it is possible here, because of the single scattering assumption, to approach this problem from first principles. In this way we will appreciate more the processes involved as well as deriving an integral of the transfer equation which can be evaluated immediately .

Our task is, given a finite path through the atmosphere which is illuminated by sunlight (coincident with a particular viewing line), calculate the total spectral radiance which is scattered into the line of sight towards the origin (coincident with the viewing point). The calculations are repeated for each sample wavelength and the results combined to get a representation of the spectral distribution of the light reaching the view-point.

Consider the diagram shown in Figure 1 which depicts a typical situation where the view-point (*E*) is near the ground looking into the sky at an angle along the direction *EO*. Note that this diagram is not to scale: the height of the atmosphere has been greatly exaggerated for clarity of presentation. The Sun is above the local horizon, hence all of the viewing path within the atmosphere is illuminated. Consider a unit volume of air located at point *P* along *EO* whose length parallel to the path is *ds*. Sunlight is incident on this volume from the direction *P'P*. Then the radiance scattered by this volume element towards *E* is given by

$$dI = E_0 \exp(-\tau_1)\, \rho(\mathbf{r}, \alpha)\, \exp(-\tau_2)\, ds \qquad (1)$$

where

- dI is the spectral radiance (units: $Wm^{-2}str^{-1}nm^{-1}$)
- E_0 is the spectral *irradiance* of the Sun's illumination outside the atmosphere (units: $Wm^{-2}nm^{-1}$)
- $\rho(\mathbf{r}, \alpha)$ is the *volume angular scattering coefficient* of the element at position $\mathbf{r}$ and scattering angle α (units: $m^{-1}str^{-1}nm^{-1}$)
- τ_1 and τ_2 are the optical distances along the paths *P'P* and *PE*, respectively (units: nm^{-1})

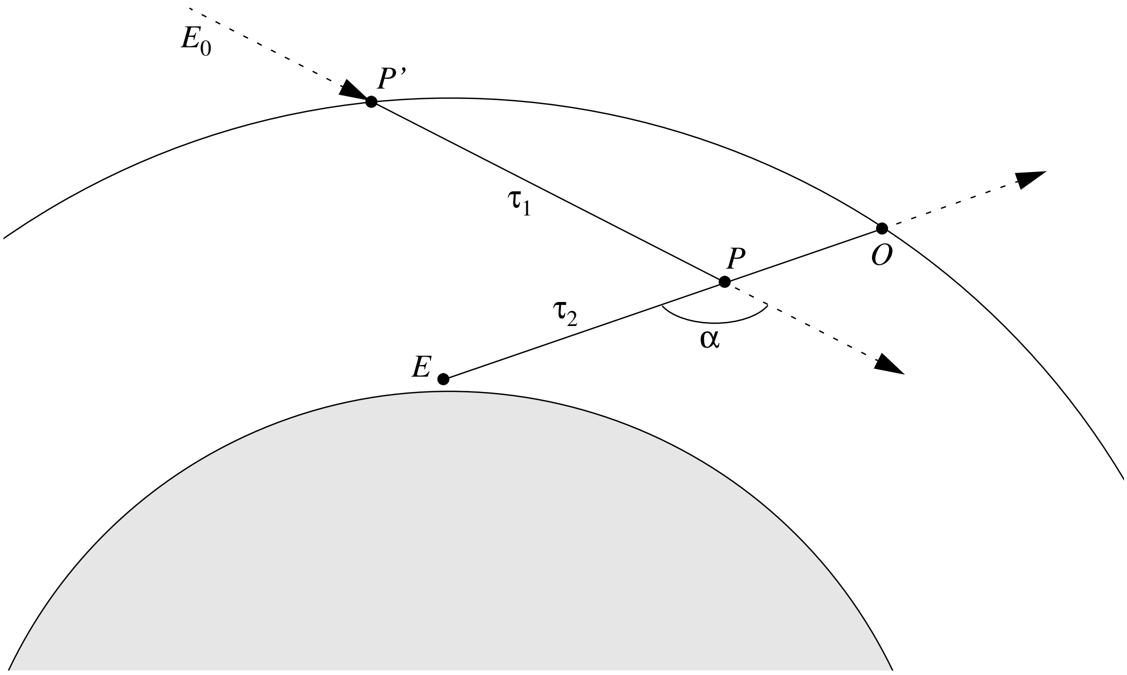

Figure 1 *Calculating the scattered radiance*

We can write the volume angular scattering coefficient as

$$\rho(\mathbf{r}, \alpha) = \frac{\beta(\alpha)\,\sigma(\mathbf{r})}{4\pi} \tag{2}$$

where $\beta(\alpha)$ is the *normalised phase function* and $\sigma(\mathbf{r})$ is the *volume total scattering coefficient*. The total scattering coefficient relates directly with the mass density, hence for our model atmosphere we can write

$$\sigma(\mathbf{r}) = \sigma(r) = \sigma_0 \exp\left(\frac{R-r}{T}\right) \tag{3}$$

This assumes a pure exponential fall of the mass density of the atmosphere with a scale height of Tkm. The quantity σ_0 is the value of the coefficient at the Earth's surface (at a radius R).

If we write $\sigma(r) = \sigma_0\lambda(r)$, where $\lambda(r)$ is now the relative mass density, we can express the optical distances τ_1 and τ_2 as

$$\tau_1 = \int_{P'}^{P} \sigma(\mathbf{r})\, ds' = \sigma_0 \int_{P'}^{P} \lambda(r)\, ds' = \sigma_0 \Lambda(P', P) \tag{4}$$

and

$$\tau_2 = \int_{P}^{E} \sigma(\mathbf{r})\, ds' = \sigma_0 \int_{P}^{E} \lambda(r)\, ds' = \sigma_0 \Lambda(P, E) \tag{5}$$

The functions $\Lambda(P', P)$ and $\Lambda(P, E)$ are the integrals of the relative mass density and are called the *equivalent path lengths* (EPLs) of the intervals concerned. The EPL is the length of atmosphere of constant density (equal to that at the surface) which has the same optical distance as the actual path.

Substituting all this back into Equation 1 gives

$$dI = \frac{E_0 \sigma_0}{4\pi} \beta(\alpha) \lambda(r) \exp[-\sigma_0 \{ \Lambda(P', P) + \Lambda(E, P) \}]\, ds \tag{6}$$

The total radiance scattered back towards E from all points along EO is then given by the integral of Equation 6

$$I = \frac{E_0 \sigma_0}{4\pi} \int_E^O \beta(\alpha) \lambda(r) \exp[-\sigma_0 \{ \Lambda(P', P) + \Lambda(E, P) \}]\, ds \tag{7}$$

This equation can also be derived directly from the radiative transfer equation given an appropriate expression for the source term. Note also that this equation is an integration of a vector quantity (the spectral radiance) along a straight-line path in 3D-space.

2.2 Rayleigh scattering parameters

The scattering of light in a pure air atmosphere is accurately described by Rayleigh's theory which provides precise definitions of the scattering coefficient and phase function. The volume angular scattering coefficient is given by

$$\rho(\mathbf{r}, \alpha) = \frac{\pi^2 [n^2 - 1]^2}{2N\lambda^4} [1 + \cos^2\alpha] \tag{8}$$

where n is the refractive index of the air, N its molecular number density and λ is the wavelength of the light. Note the strong dependence on the wavelength; for example scattering is about 15 times stronger in the blue region of the visible spectrum than in the red region.

From Equation 8 we can derive the volume total scattering coefficient for Rayleigh scattering from

$$\sigma(\mathbf{r}) = \int_0^{4\pi} \rho(\mathbf{r}, \alpha)\, d\Omega = \frac{8\pi^3 [n^2 - 1]^2}{3N\lambda^4} \tag{9}$$

and the normalised Rayleigh phase function is found by substituting Equation 8 and Equation 9 into Equation 2, giving

$$\beta(\alpha) = \frac{3}{4} [1 + \cos^2\alpha] \tag{10}$$

We will also make corrections to these expressions to take into account molecular anisotropy which recognises the fact that air molecules are not simple spherical objects (as assumed by Rayleigh's theory.) The functions that are used are

$$\sigma'(\mathbf{r}) = 1.06\sigma(\mathbf{r}) \tag{11}$$

for the scattering coefficient and

$$\beta'(\alpha) = 0.7629\left[1 + 0.9324\cos^2\alpha\right] \tag{12}$$

for the phase function. For a more comprehensive discussion of Rayleigh scattering see reference [7].

2.3 The radiance reflected from the Earth's surface

The Earth's surface itself is modelled as a texture mapped sphere with a radius of 6371km. The texture map is a composite satellite image of the cloudless Earth. For calculating radiances reflected from the surface, we have assumed a simple diffuse reflection model.

Consider the diagram shown in Figure 2 which depicts a typical situation in which the view-point is high in the atmosphere at E looking at a point P on the Earth's surface. The normal vector at this point is given by **N** and a unit vector in the direction of the Sun is indicated by **L**. Then the spectral radiance of reflected sunlight is expressed by

$$I = E_0 \exp(-\tau_1)\, R_D\, (\mathbf{N} \bullet \mathbf{L}) \exp(-\tau_2) \tag{13}$$

where E_0 is the Sun's irradiance, R_D is the Earth's diffuse reflectivity, and τ_1 and τ_2 are optical distances from P' to P and from P to E, respectively.

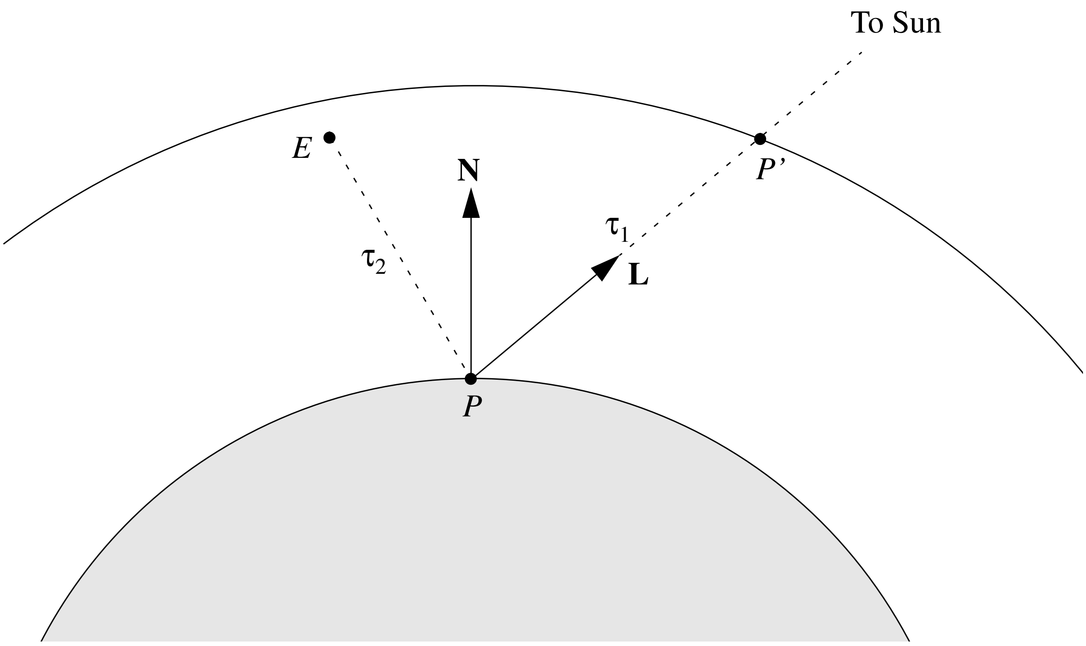

Figure 2 *Calculating the reflected surface radiance*

Using equations similar to Equation 4 and Equation 5 we can write the optical distances as

$$\tau_1 = \sigma_0 \Lambda(P', P) \quad \text{and} \quad \tau_2 = \sigma_0 \Lambda(P, E) \tag{14}$$

hence we can write the reflected radiance as

$$I = E_0 R_D (\mathbf{N} \bullet \mathbf{L}) \exp[-\sigma_0 \{\Lambda(P', P) + \Lambda(P, E)\}] \tag{15}$$

Values of the Earth's diffuse reflectivity are taken directly from the normalised RGB values in the texture map image.

2.4 Computing illuminated paths

Recall that the illuminated path is defined as a finite segment of the viewing-line which is both within the atmosphere and illuminated by direct sunlight. Hence, the illuminated path corresponds to all those points in the line-of-sight which can scatter light towards the view-point. For example, the interval *EO* in Figure 1 is the illuminated path for that particular direction and position of the Sun.

Now, there will be viewing situations whereby the line-of-sight passes through the Earth's shadow. Thus there will points along the viewing direction which do not contribute to the total light reaching the observer. An example is shown in Figure 3 where we have a viewing point situated outside the atmosphere. Indicated are five viewing directions which have different illuminated paths, some of which are of zero length. Viewing direction #3 is the only one which penetrates the atmosphere and is *partially* illuminated.

So we see that the task of determining illuminated paths involves computing these line segments, whose limits are used to evaluate the integral given in Equation 7.

Notice how the computation of the illuminated path is similar to the problem of finding the intersection of a line and a CSG (constructive solid modelling) object. The CSG object in this case is the illuminated part of the atmosphere, composed of the combination of 2 spheres and 1 cylinder. The larger sphere, whose surface corresponds to the atmosphere's upper limit, also acts as a global bounding volume for intersection testing; viewing lines which completely miss this bounding volume require no further processing.

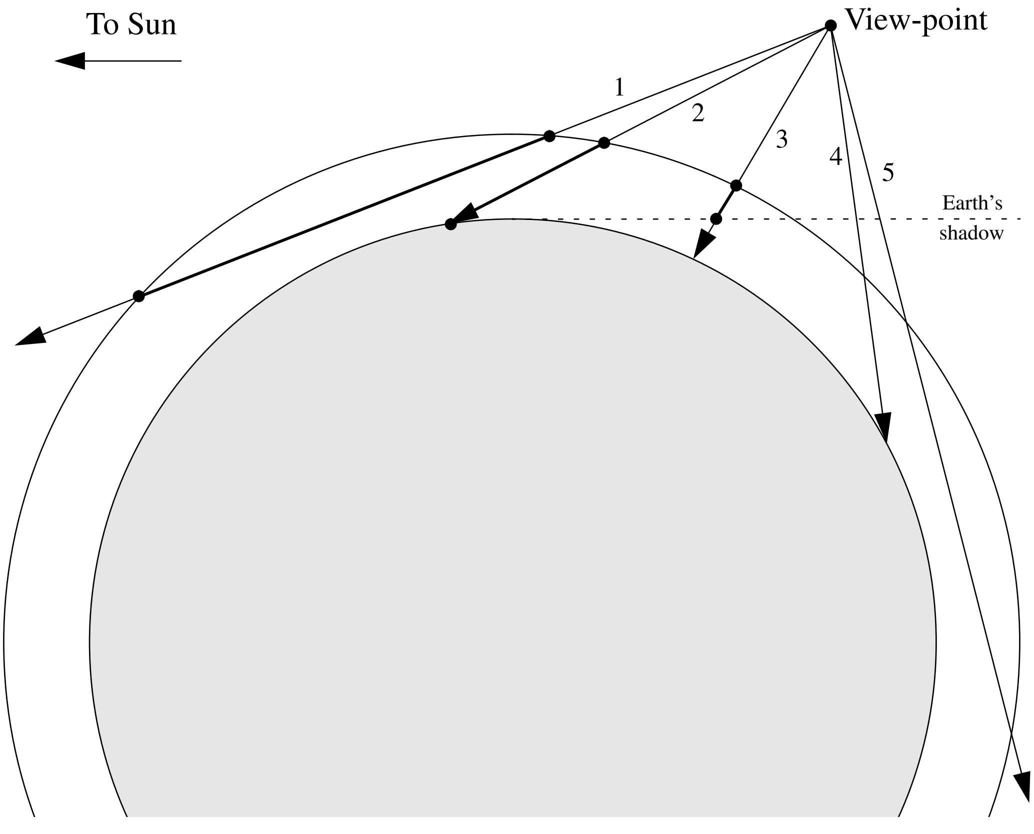

Figure 3 *Determining the illuminated path*

2.5 Numerical integration

Given the parameters describing a particular illuminated path, we can begin to evaluate Equation 7 for the scattered spectral radiance. For this purpose we have used the adaptive Simpson quadrature method[1] which estimates the integral in a recursive manner until convergence is reached.

Notice how Equation 7 has two levels of integration, an inner one for calculating the optical distances τ_1 and τ_2, and an outer one for calculating the radiance. Both the inner and outer integrations use the same adaptive technique for their evaluation, which can potentially cause problems with accuracy, should any inaccuracies in the calculation of the τs be transmitted up to the outer integration. It is important to avoid this happening, otherwise visual artifacts will appear in the final image.

Remember that Equation 7 has to be evaluated for each sample wavelength. For efficiency, a pseudo-parallel integration scheme has been devised which evaluates each sample concurrently. This avoids the repetition of effort on the wavelength independent parts of the equation. Essentially, what this means is that at each step of the integration, new estimates of the radiance at each wavelength are calculated and independently tested for convergence. Once a sample has converged it takes no further part in the integration, though in practice it was found that all spectral samples converged within 1 or 2 steps of each other.

3 From spectrum to RGB

Given the spectrum representing the radiance of scattered light, our task is to determine an RGB triplet which, when displayed on a graphics monitor, produces the same visual response in the eye as that from the original spectrum. Considering that the typical RGB monitor has a limited dynamic range and colour gamut does not make this calculation easy. However, we can still make good progress in the reproduction of a colour if we can limit the dynamic range in a consistent manner. This would not seem too unreasonable remembering that no one really objects that a photograph, for example, looks un-natural even though the range of brightnesses on the picture is far smaller than in the original scene. It seems that if we can reproduce the *quality* of the colour with good fidelity on an RGB monitor then we will have largely succeeded in our original aim.

3.1 Reconstructing the spectrum

Before transforming our spectral samples to an RGB triplet, we must estimate the equivalent continuous distribution in a process called *reconstruction.* Our problem is that we have, for computational reasons, undersampled the calculation of the spectral radiance suitable for further processing. So we need to make a reasonable estimate of the spectrum between the sample points. We do this simply by fitting a *natural cubic spline*[1] to the sample data, which approximates the spectrum with the minimum of calculation.

For our calculations we used 14 sample wavelengths spaced uniformly from 380nm to 770 nm. This gives a sample interval of 30nm. After reconstruction the spectrum is resampled at a higher resolution, at 5nm intervals, and kept in a lookup table ready for subsequent calculation.

3.2 CIE XYZ tristimulus values

A stepping-stone towards finding the RGB triplet is the conversion of the spectrum to XYZ tristimulus values[12]. This is a physical 3-component colour space for representing colours visible to the human eye and is based upon a set of three *colour matching curves* defined across the visible spectrum. These curves were derived from the results of extensive experiments on human colour vision by the CIE in 1931 and were constructed such that any colour can be matched by a linear combination of the three curves.

We can use these curves to calculate the XYZ values for a given spectral distribution, $P(\lambda)$, using the defining equations

$$X = k\int_{\lambda} P(\lambda)\,\bar{x}(\lambda)\,d\lambda \qquad Y = k\int_{\lambda} P(\lambda)\,\bar{y}(\lambda)\,d\lambda \qquad Z = k\int_{\lambda} P(\lambda)\,\bar{z}(\lambda)\,d\lambda \tag{16}$$

where

- $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$ are the colour matching curves (see Figure 4)
- k is a normalising factor (if $k = 1$ then the units of X, Y and Z will be $Wm^{-2}str^{-1}nm^{-1}$)
- and the integration is performed over the whole of the visible spectrum, typically from 380nm to 770nm.

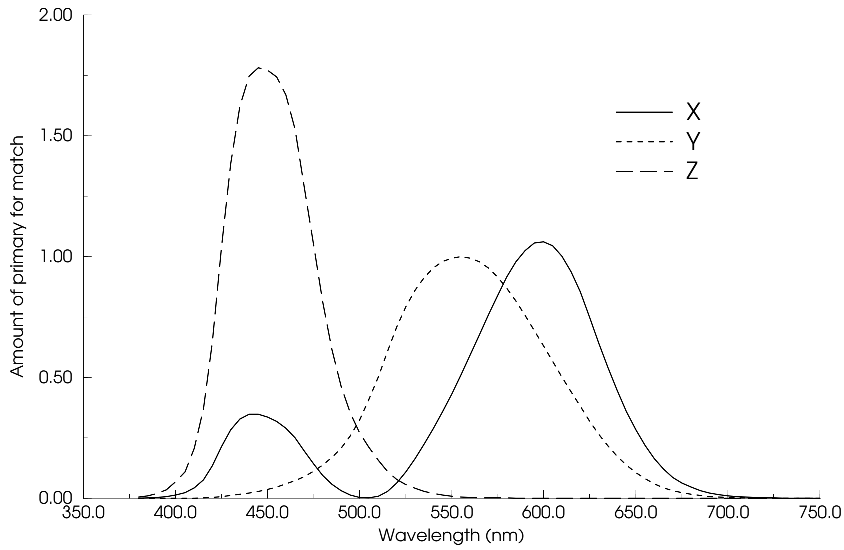

Figure 4 *The CIE colour matching curves*

In practice, both the spectral power distribution and the colour matching curves are held with sufficient resolution in data tables such that we can numerically evaluate these integrals to a good approximation as weighted Riemann sums:

$$X = k\sum_{i} P_i \bar{x}_i \Delta\lambda \qquad Y = k\sum_{i} P_i \bar{y}_i \Delta\lambda \qquad Z = k\sum_{i} P_i \bar{z}_i \Delta\lambda \tag{17}$$

Sufficient resolution here means ~5nm or less, which is why earlier we resampled the spectrum at these intervals. This should not cause any appreciable accuracy problems for most natural sources of light. Tables of $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$ can be found in reference [12].

The $\bar{y}(\lambda)$ colour matching curve was actually defined to be identical to the *luminous efficiency function* of the human eye, which has a maximum of $683lmW^{-1}$ at $\lambda = 555$nm. Hence, we can calculate the *luminance*, L, of the spectrum by setting $k = 683lmW^{-1}$ in the equation for the Y-value (units of $P(\lambda)$ are $Wm^{-2}str^{-1}nm^{-1}$, hence the units of L are $lmm^{-2}str^{-1}$ or cdm^{-2}.)

The RGB and XYZ colour systems are simply related by a linear transformation which we can write in matrix form as

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{18}$$

where X_r, Y_r and Z_r are the tristimulus values of the red phosphor at full output, and similarly for the green and blue. This equation assumes that the phosphors respond linearly to the voltages applied to each of the monitor's electron guns. This is not generally true, of course, but once we have solved Equation 18 for R, G and B, gamma correction is applied before displaying the colour on a graphics screen. The matrix of tristimulus values can be derived from knowledge of the x, y, z chromaticities of the phosphors and white point using a very simple procedure which is, for example, covered adquately by Watt[11].

3.3 Dynamic range compression

It is very common to find that RGB values obtained from Equation 18 are not in the range [0, 1] necessary for display on a monitor screen. This is because we still haven't dealt with the problem of the limited dynamic range of the monitor phosphors. Generally speaking, values greater than 1 are a dynamic range problem but values less than 0 mean that the colour cannot be displayed at all, at any intensity, i.e. it lies outside the *colour gamut* of the monitor. In tests of the software, we found that very few values were out of gamut and of those which were, their magnitudes were <0.05. These instances were treated by setting their value to zero.

There are a number of ways to deal with the dynamic range problem[2]. We can adjust the RGB values *locally* by, for example, clamping the offending value to the required range, or we can adjust the values *globally* by considering the full range within the entire image and scaling everything at the same time so that no values are out of range. The second method is generally preferred since it avoids *limit errors*, leading to Mach banding in the image, though in practice what has been done is to introduce a scaling factor which is applied to every calculated spectrum, such that no RGB values greater than 1 are actually encountered.

4 Results

The author has developed a ray-tracing rendering system which is built around the scattering algorithms described above. The software was used to generate a number of images of the Earth and atmosphere, some of which will be shown here[13]. We will start with a view of the Earth from outside the atmosphere - see Figure 5.

Each image is a 4° view of the Earth from 240,000km above the equator at longitude 0° and the Sun is overhead at latitude 23°N and longitude 110°E. This corresponds to late afternoon (BST) around midsummer in the Northern Hemisphere. The image on the left shows only the light which has been scattered by the atmosphere. We can clearly see the effect of the Earth's shadow - light is reaching parts of the atmosphere above points on the surface which are in darkness. The image on the right shows a complete view of the Earth and atmosphere.

The next set of images, shown in Figure 6, show a view from within the atmosphere. Each image is a 90° (horizontally) view of the horizon from 2m above the ground. The Sun's position is indicated by the small disk (for display purposes only, though its angular size is correct) at approximately 10°, 5°, 0°, -5° and -10° from the horizon. When the Sun is below the horizon the view-point is actually inside the Earth's shadow, so we see less illuminated atmosphere and thus less intensity from this part of the sky, as expected.

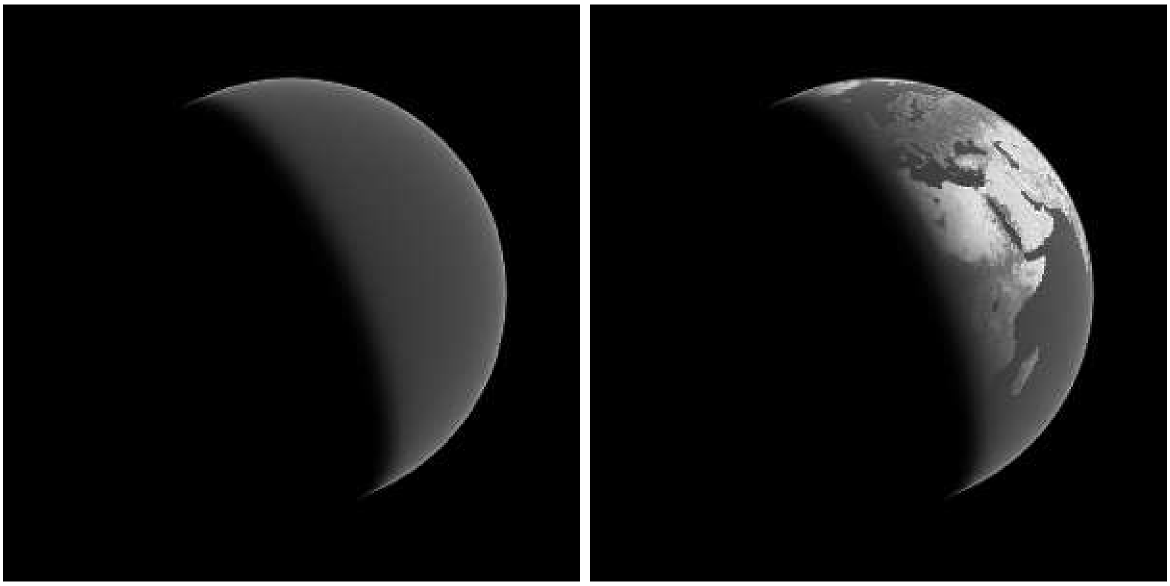

Figure 5 *External view of the atmosphere*

5 Ideas for further work

There are at least two significant problems remaining with the model thus far described. The first concerns *skylight* reflected from the Earth's surface. As the model stands at the moment the surface is only illuminated directly from the Sun. In reality, the surface is also illuminated from a secondary source of light which comes from all of the atmosphere in the hemisphere above each point on the surface: this is skylight. Skylight is a diffuse form of lighting, dependent on the position of the Sun in the sky as seen from the surface, which would not be easy to calculate using the ray-tracing rendering technique developed here. The ideas from distributed ray-tracing would be of value here.

The other problem is evident in the sunset images shown previously, especially those with the Sun on or below the horizon. We can see that there is a distinct lack of blue colour in the sky when in reality the sky remains quite blue (though of a different quality) long after sunset. There are two effects at work here which the model does not address: *multiple-scattering* and *ozone absorption*, of which the latter has its most noticable effect on the sky colour after sunset[3]. Incorporating multiple-scattering is fraught with computational problems, which in all honesty should be dealt with using other rendering techniques. However, because we are at the moment only concerned about scattering in pure air (an optically thin medium), the development of a *double-scattering* light model may be a good compromise[6]. The incorporation of ozone absorption would be fairly easy, as a spherical shell of constant density with an appropriate height and thickness in the atmosphere.

Finally, it would be interesting to include the effects of *refraction* in our model. Then we may be able to simulate such effects as the flattening of the Sun's disc at sunset and the *green-flash phenomenon*[9], which is thought to be caused by dispersion in the Earth's atmosphere.

And, of course, it would be nice to add some clouds, but that's another story!

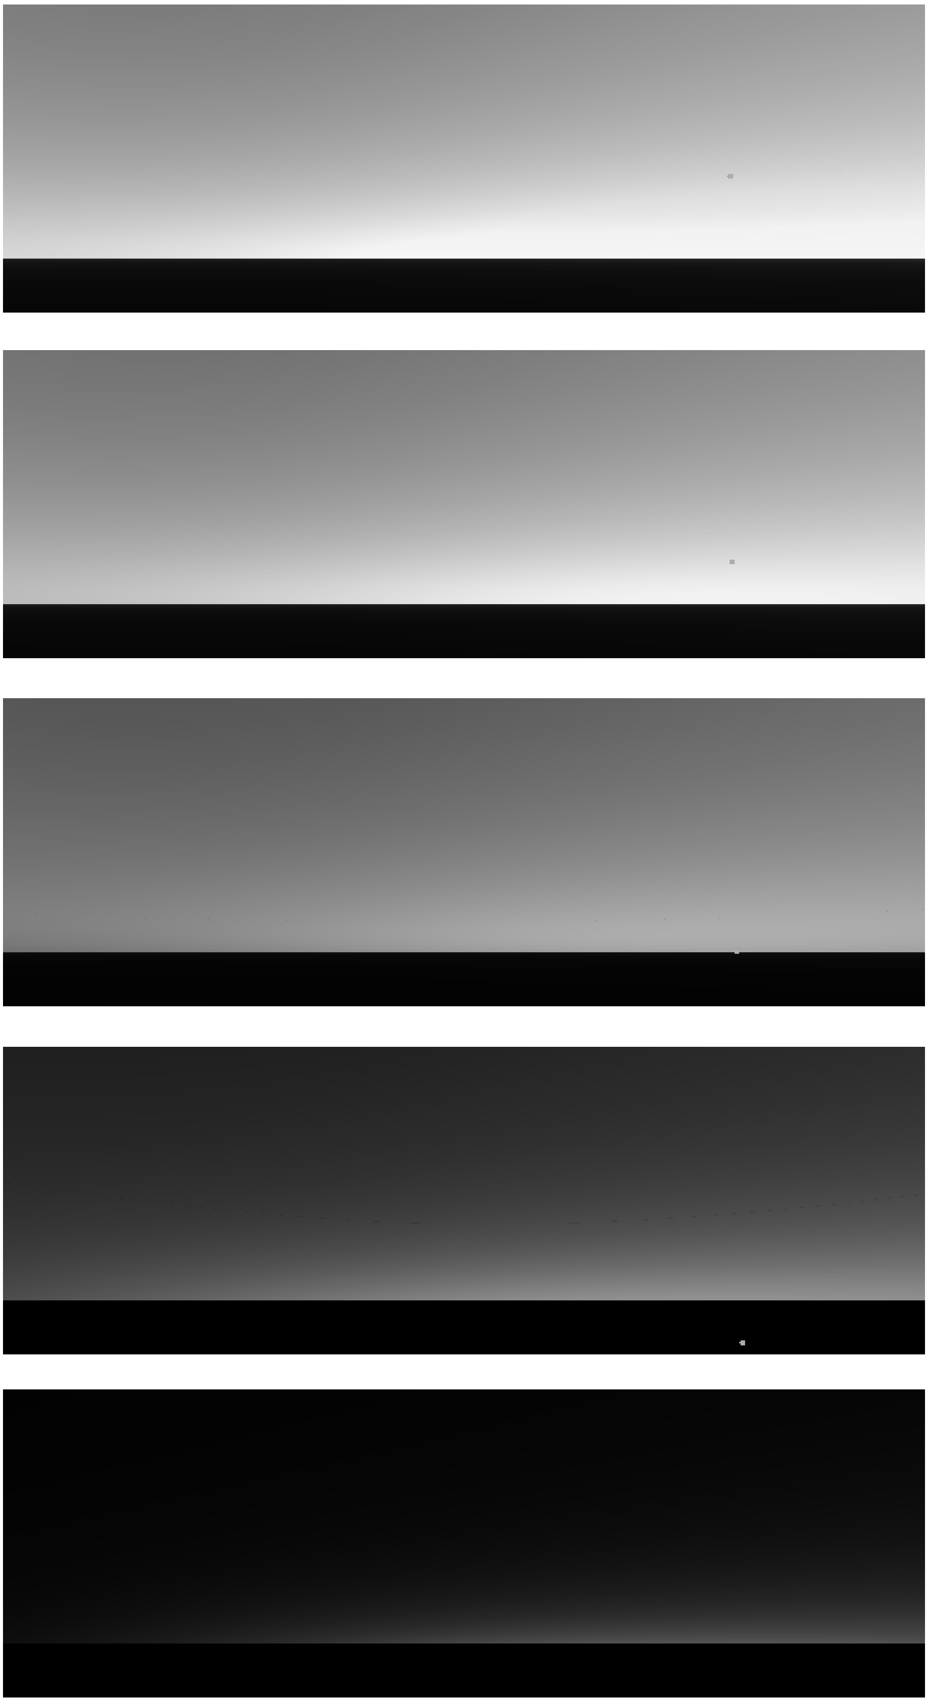

Figure 6 *Internal view of the atmosphere: a sunset*

Acknowledgements

The author would like to thank Terry Hewitt, Director of the Computer Graphics Unit, for providing the resources which enabled this work to be undertaken, and Chris Lilley of the CGU for valuable discussions on the theory of colour and colour spaces. Lastly, thanks go to Martin Preston, Paul Lever and Chris Lilley (again) of the CGU for encouraging me to write this paper for the EG-UK conference.

References

[1] Burden R L and Faires J D, *Numerical Analysis*, Prindle Weber and Schmidt, Boston, 1993.

[2] Hall R, *Illumination and Color in Computer Generated Imagery*, Springer-Verlag, New York, 1987.

[3] Hulburt E O, *Explanation of the brightness and color of the sky, particularly the twlight sky*, J. Opt. Soc. Am. **43**(2), February 1953, 113-118.

[4] Kaneda K T, Okamoto T, Nakamae E and Nishita T, *Photorealistic image synthesis for outdoor scenery under various atmospheric conditions*, The Visual Computer **7**(1991), 247-258.

[5] Klassen R V, *Modeling the effect of the atmosphere on light*, ACM Trans. Graph. **6**(3), July 1987, 215-237.

[6] Lee R L, *Horizon brightness revisited: measurements and a model of clear-sky radiances*, App. Opt. **33**(21), July 1994, 4620-4628.

[7] McCartney E J, *Optics of the Atmosphere*, John Wiley, New York, 1976.

[8] Meyer G W, *Wavelength selection for synthetic image generation*, Comp. Vis. Graph. Image Proc., **41**(1988), 57-79.

[9] Minnaert M, *The Nature of Light and Color in the Open Air*, Dover Publications, Mineola NY, 1954.

[10] Nishita T, Sirai T, Tadamura K and Nakamae E, *Display of the Earth taking into account atmospheric scattering*, ACM Computer Graphics Proceedings, Annual Conference Series, 1993, 175-182.

[11] Watt A, *Fundamentals of Three-Dimensional Computer Graphics*, Addison-Wesley 1989, p348.

[12] Wyszecki G and Stiles W S, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, John Wiley, New York, 2nd Ed, 1982.

[13] Colour versions of these and other images can be found via the URL given at the beginning of this report.