INTERFERENCE PHENOMENA
AND AN EFFICIENT SOLUTION
TO THE OUT OF GAMUT COLORS PROBLEM

by

BRIAN E. SMITS

A THESIS

Presented to the Department of Computer
and Information Science
and the Honors College of the Universtiy of Oregon
in partial fulfilment of the requirements
for the degree of
Bachelor of Arts

June 4, 1990

APPROVED: [redacted]

Gary W. Meyer

Abstract

An Abstract of the Thesis of

Brian E. Smits for the degree of Bachelor of Arts

in the Department of Computer and Information Science

to be taken June 1990

Title: INTERFERENCE COLORS AND AN EFFICIENT SOLUTION TO THE OUT OF GAMUT COLORS PROBLEM

APPROVED: [redacted]

Gary W. Meyer

The principles of color theory which apply to physically based simulations of natural phenomena are reviewed. The physics of interference phenomena are used to develop a simulation of single and multiple film interference. Three example of interference phenomena are reproduced. To correctly display interference colors, an efficient technique for handling out of gamut colors is presented.

Acknowledgements

I would like to thank my principle advisor Gary Meyer for the time and energy he spent with me throughout the entire thesis process. I would also like to thank my advisors Richard Koch and Joseph Fracchia for their help.

Contents

List of Figures

List of Tables

1 Introduction

Advances in workstation design have made color much more important in computer graphics. The price of color workstations has dropped significantly and their capabilities have increased dramatically in the last ten years. This means that computers with color displays will be used to do more than design, simulate, and display the geometrical properties of objects. As the reproduction of object surface characteristics continues to advance, new image synthesis techniques will allow surface properties such as color and texture to become a part of the creative process. Early design tools and simulations displayed simple line drawings of the objects being modelled. As computer technology improves and becomes more readily available, these design tools will allow a much more accurate simulation of the final product.

A problem that occurs when incorporating color into the simulation or design process is that much of the use of color is not physically correct. Color is often chosen in a haphazard manner because exact specification of color is not readily available. Part of this problem is that color is specified for display devices in a manner different from the way it is represented in nature. Color devices generally produce color as a mixture of three primary colors whereas in nature, color is produced by light interacting on a wavelength by wavelength basis with the surfaces it strikes. When light interacts with a surface, the energy at each wavelength can be reflected, transmitted, scattered, or absorbed. The color of the object is determined by what effect the surface of the object has on the wavelengths of light that strike it. One way of modeling various types of surfaces is to use the physical properties of the object to determine how it actually affects light, convert the light that gets transmitted to the viewer into the form used by the color device, and display it on the device. In this way as accurate a representation of the color can be used as possible.

When color was first used in computer graphics, the color selection process involved choosing red, green, and blue values that produced a color close to the desired color. No care was taken in making sure that this was the actual color of the object in the desired lighting conditions. The technique of using the actual surface properties of the object and the physical properties of the light is fairly new. Robert Cook was one of the first to do this wavelength based approach. He used the surface properties of the object to determine the characteristics of the light being reflected back to the viewer (Cook 1981). The interactions between the light and the surface were modeled and the results were used to determine the color. The color was then converted from a description of the physical makeup of the light into the red, green, and blue values that a color monitor uses (Meyer 1986).

There are many different ways in which the physical properties of surfaces can affect light. The various ways that nature produces color are just beginning to be explored and modeled in computer graphics. Each new technique for modeling natural phenomena expands the set of natural phenomena that can be reproduced in computer graphics. This allows a more complete understanding of how light interacts with the environment. It also allows more accurate simulations of our world. Computer graphic examples of modelling the mechanisms in nature which produce color include the colors of the sky (Klassen 1987), and the colors produced by the refractive properties of gemstones (Thomas 1986), and rainbows (Musgrave 1989).

One of the ways nature produces color that has not yet been modelled in computer graphics is through interference phenomena. Interference phenomena cause the colors seen on soap bubbles and oil films. The sequence of colors produced by this type of interference is called Newton's colors as he was the first to study them. Other types of interference produce iridescence which is responsible for the colors found on peacock feathers, opals, and some insect wings. The colors produced by interference are a result of very thin films which selectively reflect or transmit each of the wavelengths of light that strikes it. This is the sort of natural phenomena that can not be modeled accurately without looking at the physical properties of light and surfaces. A visual simulation of interference phenomena leads to a better understanding of the properties of light and what happens to light as it interacts with surfaces.

In the process of modeling the physical interactions between light and surfaces and the colors that are produced, another problem occurs. When some colors are converted into the form used by a color device, they are found to lie outside the range or gamut of colors that the device can produce. Something needs to be done with these colors, called "out of gamut color", in order to represent the image as accurately as possible. This problem of out of gamut colors can occur anytime a color image is transferred from one device or medium to another.

Each color device uses a different representation for color, based on the physical characteristics of the device. Suppose color is used in the design process of an object on a color computer. When it comes time for printing a color picture of the final design, it is important that the color printer produces a picture that is as close as possible to the picture on the color monitor. Because these devices are different, it is very likely that it will be impossible to reproduce the colors exactly with the color printer. It is very important, however, to get a picture that will look as close as possible to the image that was designed on the computer, otherwise the use of color in the design process is pointless. Also, if the exact specification of the color is known on the color workstation or elsewhere, it can be converted into a more standard representation so that when the part is manufactured, the correct color for the manufacturing process will already be specified during the design process.

This thesis will present a physically based computer graphics simulation of interference phenomena and a technique for dealing with the problems that come from representing interference colors on display devices. First there is a section on color theory which gives the background for physically based simulations of color phenomena. The next section will describe two different types of interference phenomena and how they can be reproduced in computer graphics. Then an efficient solution to out of gamut colors will be presented.

2 Color Theory

2.1 Introduction

Color theory is the study of light and how the eye perceives it as color. It is also the study of various representations of color and the benefits and disadvantages of the different representations. By studying what color is and how it is perceived by the observer, much more realistic and accurate simulations of the natural world can be achieved in computer graphics. The principles of color theory are used to take a physical representation of color that comes from modeling the interactions between light and surfaces and convert it into a form that can be used on a color monitor or some other color device. Color theory provides an important background for understanding how to model natural phenomena in computer graphics and for understanding the problems that can arise.

2.2 Light

Light is similar to radar, x-rays, gamma rays, and radio waves in that these are all forms of electromagnetic radiation. Electromagnetic radiation travels at a speed of

3.0 x 10^8 m/s in bundles of energy called quanta. Electromagnetic radiation also behaves like a wave, and the wavelength of the quanta is what we use to distinguish the various forms electromagnetic radiation. Wavelengths between 380 nm and 780 nm stimulate a response in the eye. This range of wavelengths is called the visible spectrum of light. Light of different wavelengths stimulate the eyes in different ways. When light composed mainly of shorter wavelengths enters the eye, it is perceived as a bluish color. The eye perceives light with medium range wavelengths as greenish and that with longer range wavelengths as reddish.

Most of the light that enters the eye usually contains light of all wavelengths in the visible spectrum. If light is all of one wavelength it is called monochromatic. Lasers produce monochromatic light, but most natural sources of light are not monochromatic. The intensity of the light at each wavelength in the visible spectrum determines how the eye perceives the light. Therefore, it is necessary to have a representation of the intensity of light at each wavelength. This representation is called a spectral energy distribution denoted $E(\lambda)$ where $E(\lambda)$ is the intensity of the light at wavelength λ. (Figure 1). The total energy P of the radiation of a spectral energy distribution is defined to be

$$P = \int_{380}^{780} E(\lambda) d\lambda. \tag{1}$$

This quantity is called the radiant flux.

The light that enters the eye after hitting a surface is based on the spectral energy distribution of the light source and the surface properties of the object. The surface of

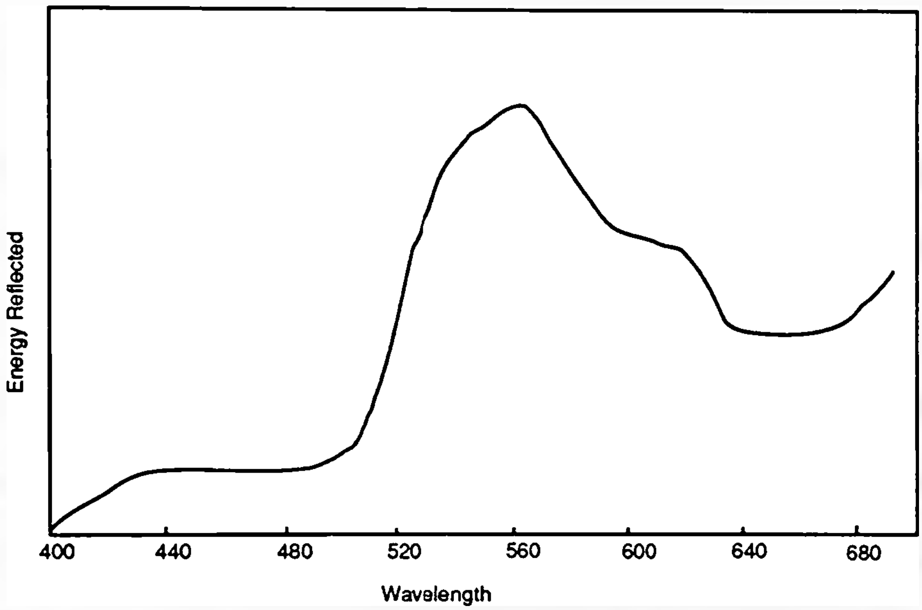

Figure 1: Spectral energy distribution for grass in sunlight (Cornsweet 1970).

an object acts like a filter, screening out certain amounts of light at each wavelength and reflecting the rest away from the object. This property of a surface is called its specular reflectance (Cornsweet 1970) and is denoted $\rho(\lambda)$. A mirror has a reflectance of about 100% at every wavelength ($\rho(\lambda) = 1$) since it filters out almost none of the light. A jet black surface has a reflectance of close to 0% at all wavelengths ($\rho(\lambda) = 0$) since the surface filters out almost all light. A reddish surface would have a higher reflectance at long wavelengths than at short since longer wavelengths will be perceived as red. The intensity or radiant flux reflecting off of a surface is defined to be

$$P_{al} = \int_{380}^{780} E(\lambda)\rho(\lambda)d\lambda. \tag{2}$$

where $E(\lambda)$ is the spectral energy distribution of the light source.

2.3 Color Matching

When light enters the eye, it triggers a response that is based on the wavelengths that make up the light and the intensity of light at each wavelength. In other words, the spectral energy distribution of the light determines what color the eye sees (as long as the surround and other factors are constant). In order to reproduce a color, it is necessary to create a response in the visual system of the viewer that matches the response created by the original light. This does not necessarily mean reproducing the spectral energy distribution of the original sample. Two colors are said to match if an observer looking at both colors on the same backgrounds with the same surrounding light cannot tell the difference between the colors.

A color matching experiment generally takes the following form. A test sample of light of some arbitrary color is projected onto a neutral grey background. Projected next to the test sample is another light or blend of lights over which the subject has control of the intensity. The subject attempts to match as closely as possible the test sample by adjusting the intensity of the other light or lights.

Through various color matching experiments, it has been found that if the subject can control the mixture of three lights, almost every test sample can be matched. Problems occurred when the sample light was monochromatic or nearly so. In this case, the only time a monochromatic light could be matched was if it was the same wavelength as one of the matching lights. If the subject can shine one of the lights onto the the test sample thereby changing the appearance of the test sample, every test sample, including monochromatic lights, can be matched. This was strong evidence for the theory that there are three different types of receptors in the eye, each stimulated differently by light of different wavelengths. (Cornsweet 1970). In order for three lights to be sufficient, it must be impossible for one of the lights that the subject has control over to be matched using the other two. In other words, none of the three lights the subject has control over can be a mixture of the other two. Mathematically, the following holds

$$L_s = L_1 + L_2 + L_3 \quad (3)$$

where $L_i \neq L_j + L_k$ if $i \neq j \neq k$, and L is the intensity of one of the lights. Shining a light on the test sample is represented by using negative values for L_i even though it is physically impossible to have a light with a negative intensity.

The trichromatic nature of the visual system is important because it means that a color can be reproduced without reproducing its spectral energy distribution. It makes color reproduction economically feasible because it means only three pieces of information need to be stored, not the entire spectral energy distribution. If the entire spectral energy distribution was needed at each point on a television or monitor, color devices would be infeasible.

Now we look at matching a sample of monochromatic light. By varying the intensity of the monochromatic sample and recording the intensities of the matching lights it was found that the following relationship holds

$$kL_s = kL_1 + kL_2 + kL_3 \tag{4}$$

where k is a constant and L_i is the intensity of the matching light i. This is a very important result. It means there exists a linear dependence between the intensity of the test sample and the intensities of the matching lights.

This linear dependence permits the conversion of a spectral energy distribution to three values (tristimulus values) which correspond to the intensities of the matching lights when light with the given spectral energy distribution is matched. This conversion uses three matching functions which give the intensity of the matching lights at each wavelength. To get the matching functions, we do the matching experiment on monochromatic lights at intervals of one nanometer over the entire visible spectrum. Interpolating between wavelengths gives us a smooth curve for each of the three matching lights called $m_{L_1}(\lambda)$, $m_{L_2}(\lambda)$, and $m_{L_3}(\lambda)$. If three monochromatic lights with wavelength l = 444nm, m = 526nm, and h = 645nm are used, the matching functions look like those in Figure 2. The the matching functions take on negative values because they are unable to match a monochromatic light without adding one or two of the lights to the sample which is represented as negative values for those lights. Now for each wavelength of the sample light we can calculate the intensity of each of the matching lights. At any wavelength, the matching light intensity is the product of the intensity of the light, given by $E(\lambda)$, and the value of the matching function, given by $m(\lambda)$. The intensities of the three matching lights is found by

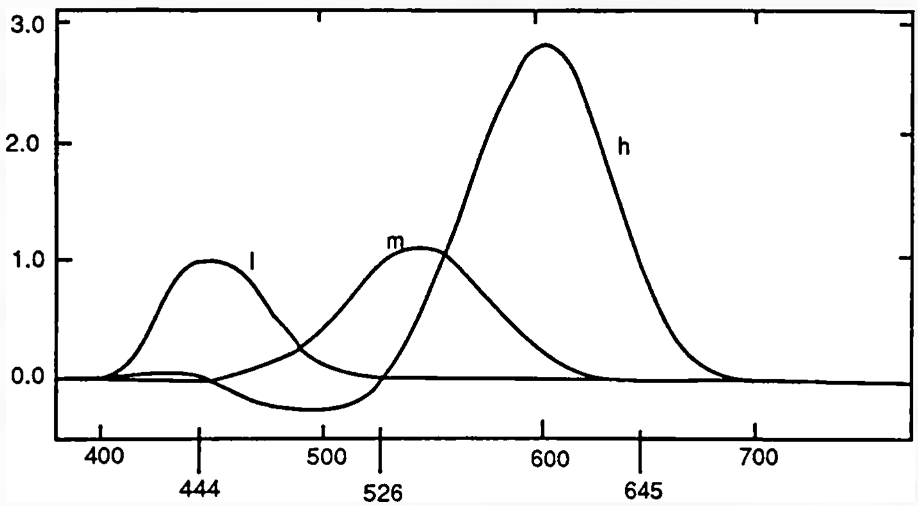

Figure 2: Matching functions for three monochromatic lights with wavelengths 444 nm, 526 nm, and 645 nm.

integrating over the visible spectrum. For light reflecting off of a surface with specular reflectance $\rho(\lambda)$ illuminated by a light source with a spectral energy distribution $E(\lambda)$, the following formulae are used

$$I_1 = \int_{380}^{780} E(\lambda)\rho(\lambda)m_{L_1}(\lambda)d\lambda. \tag{5}$$

$$I_2 = \int_{380}^{780} E(\lambda)\rho(\lambda)m_{L_2}(\lambda)d\lambda. \tag{6}$$

$$I_3 = \int_{380}^{780} E(\lambda)\rho(\lambda)m_{L_3}(\lambda)d\lambda. \tag{7}$$

where m_{L_i} is the matching function for the i^{th} matching light and I_i is the intensity of the i^{th} matching light. Each choice of matching lights gives a different set of matching functions.

2.4 Color Spaces

The tristimulus values given by the matching functions form a space with three dimensions called a color space. Because of the linear dependence of the matching process, this space is a vector space. This is important because it means that the color space can be transformed to new color spaces with new matching functions using linear transforms. The new matching functions are determined by applying the linear transform to the three values of the set of matching functions at each wavelength as follows

$$\begin{bmatrix} m'_{L_1}(\lambda) \\ m'_{L_2}(\lambda) \\ m'_{L_3}(\lambda) \end{bmatrix} = T \begin{bmatrix} m_{L_1}(\lambda) \\ m_{L_2}(\lambda) \\ m_{L_3}(\lambda) \end{bmatrix} \tag{8}$$

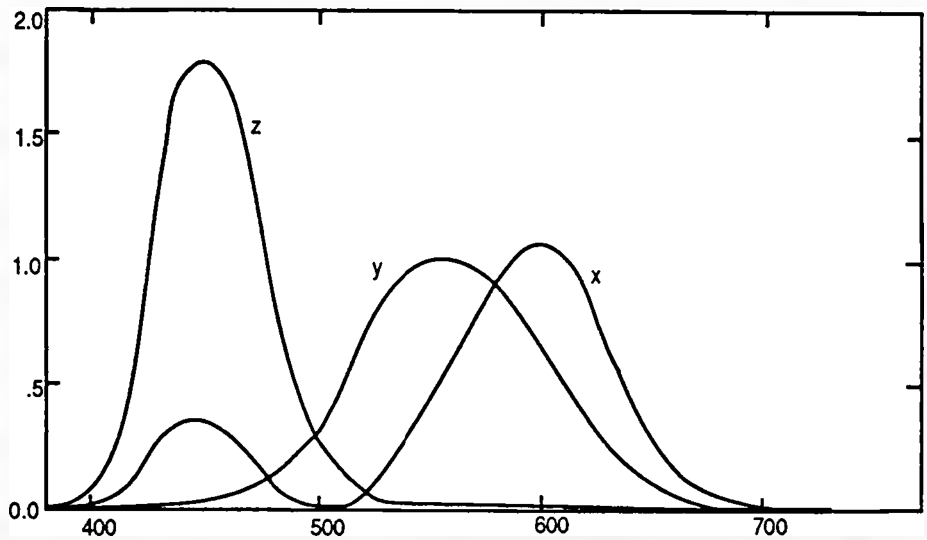

Figure 3: CIE XYZ matching functions.

where $m_{Li}(\lambda)$ is the i^{th} matching function, $m'_{Li}(\lambda)$ is the i^{th} new matching function, and T is a linear transform. By choosing the correct linear transform, it is possible to get matching functions that are all positive. The Commission Internationale d'Eclairage (CIE) in 1931 adopted a set of matching functions that is positive for all wavelengths (Figure 3). These new matching functions define the XYZ color space. Also, these matching functions have the property that white or achromatic light (light having the same intensity at all wavelengths) gives equal values for X, Y, and Z. The X, Y, and Z primaries are defined mathematically and cannot be represented by physical lights. If they could, they would be able to match any monochromatic light with only positive intensities and this is impossible. The XYZ color space has the additional property that the Y value is a measure of the luminance of the color.

The result of mixing colors together becomes easy to calculate in these linear color spaces. Given these color spaces are vector spaces, mixing two colors in a color space is equivalent to vector addition. Because of the properties of vector spaces, two lights with variable intensities can produce a range of colors that lie in a plane formed by the coordinates of the two colors being mixed and the origin. The range of colors is limited to the region between the two vectors that define the component colors since negative intensities of colors are physically impossible.

When colors are specified in XYZ space, the magnitude of X, Y, and Z is often not as important as the proportions between them. Two colors with the same proportions of X,Y, and Z differ only in brightness. Because of this, it is useful to define a pair of numbers, x and y, called chromaticity coordinates (Williamson). Chromaticity coordinates for a color represented by tristimulus values in XYZ space are defined as follows

$$x = \frac{X}{X+Y+Z} \qquad y = \frac{Y}{X+Y+Z}. \tag{9}$$

These formulae are actually just a projection of the color onto the plane $X+Y+Z=1$. An achromatic or colorless (grey) color has $X = Y = Z$ so the chromaticity coordinates are $x = 1/3$ and $y = 1/3$. The spectral or monochromatic colors form a horseshoe-shaped curve (Figure 4). Note that having the brightness Y and the x and y chromaticity coordinates, it is possible to get the original XYZ tristimulus values back.

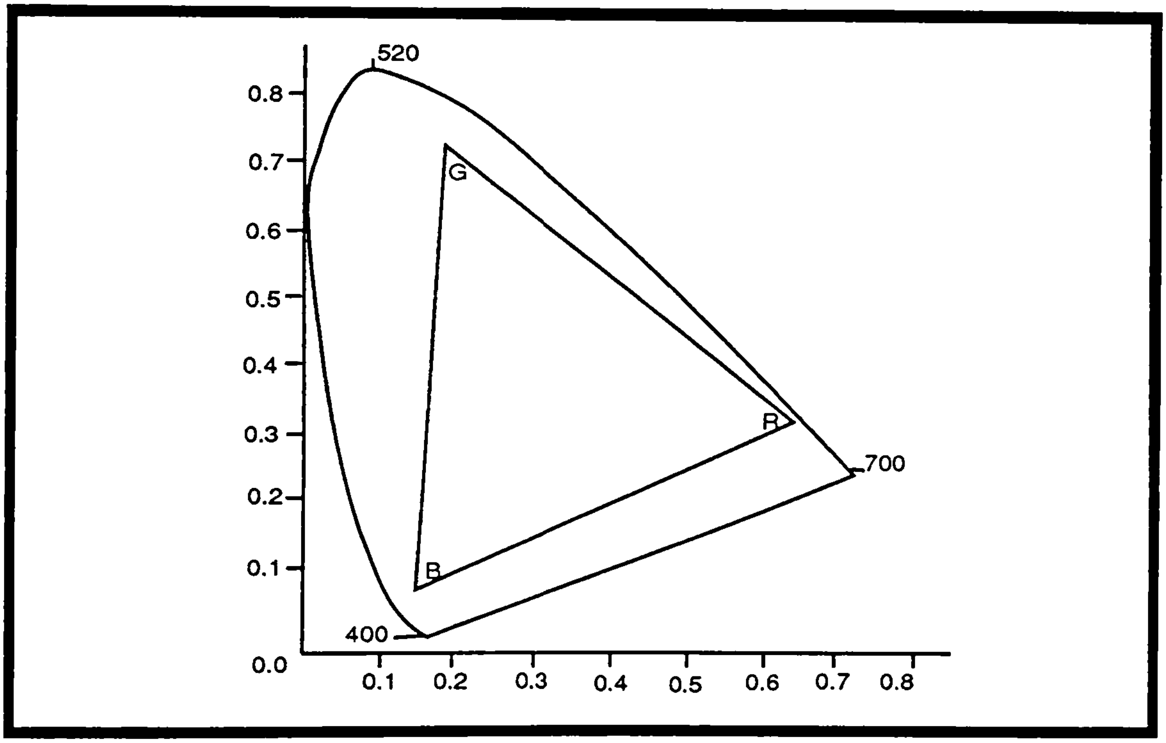

Figure 4: CIE Chromaticity diagram showing the coordinates of the red, green, and blue primaries for a color monitor and its gamut. (Williamson 1983).

Mixing colors on a chromaticity diagram forms colors that lie on a line between the two points. This result can be seen by considering how colors mix in the XYZ color space. Mixing two colors in the XYZ color space forms colors that lie in a plane and that plane goes through the origin. Transforming to chromaticity coordinates limits the colors to the plane $X + Y + Z = 1$. The range of colors on the chromaticity diagram is limited to the projection of the first plane onto the $X + Y + Z = 1$ plane, which forms a line. The fact that the colors cannot have negative intensities implies that the range of colors lies between the two points of the colors being mixed. If three colors are used, the range of mixtures lies in the triangle formed by the three points. This can be seen by mixing two of the colors first, and then mixing the third color with the mixture of the first two.

For reproducing color on a monitor, three primaries must be chosen. One of these primaries is made up of the longer wavelengths and is designated R for red. The second is composed of mid-range wavelengths and is designated G for green. The last is composed of the short wavelength and is designated B for blue. This color space is generally called the RGB color space, even though the R, G, and B primaries may be different from device to device. These RGB primaries are vectors in the XYZ color space and define a parallelpiped in the space since each of the primaries has a maximum intensity. This parallelpiped is the subspace of colors or gamut that these primaries can reproduce. The gamut defined by any three primaries will always be a subset of the total color space since every color in the gamut is in the color space, but there exist colors outside of the gamut and within the color space. This is easy to see on a chromaticity diagram (Figure 4). The three primaries are points on the diagram and form a triangle. Mixing various amounts of the primaries allows any color within the triangle to be reproduced, but colors outside the triangle always exist and are outside of the gamut of the primaries. The problem of reproducing out of gamut colors will be discussed in section 4.

Looking at a pair of points, it would be convenient to have a metric that tells us how different these colors appear to the viewer. The simplest of these would be to use Euclidean distance in the XYZ or other linear color space. After checking a few pairs of points, however, it becomes apparent that two colors in one area of the color space appear distinctly different, while two colors in another area of the color space differing by the same Euclidean distance appear the same. It would be nice to have a color space in which an equal distance in the space resulted in an equal perceptual change. It is generally assumed that no linear transform will give a perceptually uniform color space. To date all perceptually uniform color spaces require non-linear transforms. The CIE has adopted a uniform color space called $L^*u^*v^*$ space. The following transform (CIE 1978) takes XYZ colors to $L^*u^*v^*$

$$\begin{aligned} L^* &= 116(\tfrac{Y}{Y_n})^{\frac{1}{3}} - 16 \quad \tfrac{Y}{Y_n} > 0.01 \\ u^* &= 13L^*(u' - u'_n) \\ v^* &= 13L^*(v' - v'_n) \end{aligned} \tag{10}$$

with:

$$u' = \frac{4X}{X+15Y+3Z} \qquad v' = \frac{9Y}{X+15Y+3Z}$$
$$u'_n = \frac{4X_n}{X_n+15Y_n+3Z_n} \quad v'_n = \frac{9Y_n}{X_n+15Y_n+3Z_n} \tag{11}$$

where X_n, Y_n, and Z_n are the coordinates of the white point, the point designated as white by the person doing the transform. This color space has the property that L^* measures perceived brightness rather than the intensity of the light. It also has the property that all achromatic colors have $u^* = v^* = 0$.

There are several ways of comparing the differences between two colors in $L^*u^*v^*$ space. The color difference ΔE is measured by the Euclidean distance

$$\Delta E = [(\Delta L^*)^2 + (\Delta u^*)^2 + (\Delta v^*)^2]^{\frac{1}{2}} \tag{12}$$

This measures total difference in color between two colors. Another measurable property of color is hue. The hue corresponds to the name we give to colors, such as red, orange, yellow, etc. Saturation is also a measurable property of color. It is a measure of how different a color is from neutral or grey. Red and pink are about the same hue, but red is much more saturated than pink. Monochromatic colors are saturated whereas achromatic colors are unsaturated. The difference in hue ΔH between two colors $L_1^*u_1^*v_1^*$ and $L_2^*u_2^*v_2^*$ is given by the formula

$$\Delta H = [(\Delta E)^2 - (\Delta L^*)^2 - ((u_1^{*2} + v_1^{*2})^{\frac{1}{2}} - (u_2^{*2} + v_2^{*2})^{\frac{1}{2}})^2]^{\frac{1}{2}}. \tag{13}$$

The difference in saturation (ΔS) between two colors is given by the formula

$$\Delta S = \frac{(u_1^{*2} + v_1^{*2})^{\frac{1}{2}}}{L_1^*} - \frac{(u_2^{*2} + v_2^{*2})^{\frac{1}{2}}}{L_2^*}. \tag{14}$$

3 Interference Phenomena

3.1 Introduction

Interference phenomena are found in many places. Perhaps the most common example is the colors produced on the surface of oil slicks and soap bubbles. Interference effects can also be seen on old vases, metal, and peacock feathers. Structural engineers analyze stress by viewing interference patterns in plastic models. In most cases, these interference effects are caused by very thin films which determine the amount of light at each wavelength that is reflected back to the viewer. As the thickness of the film varies, the spectral composition of the light and therefore the color that gets reflected back to the viewer changes as well. Interference effects have not been modeled in computer graphics so this work enlarges the set of natural phenomena that have been simulated in computer graphics. They are another example of using the physical properties of light and the way it interacts with surfaces to more accurately model natural phenomena.

3.2 The Physics of Interference Phenomena

Interference phenomena are the result of the wave-like properties of light. According to the superposition principle, when two waves meet the amplitude at a point of the resultant wave is the sum of the amplitudes of the two component waves at that point. Suppose two waves with the same wavelength and amplitude merge together. If the crest of both waves match up exactly, perfect constructive interference occurs and the resulting wave has twice the amplitude of the two original waves. These two waves are said to have a phase difference of 0. If the crest of one wave matches up with the trough of another, perfect destructive interference results and the waves cancel each other completely. The phase difference in this case is equal to π. Given the phase difference δ of two waves with amplitude A, the resulting amplitude A' is given by

$$A' = 2A\cos\left(\frac{\delta}{2}\right) \tag{15}$$

Because light has wave-like properties, it obeys this rule, and if two light waves merge, they will affect each other (Tipler 1976). Interference phenomena produce phase differences that work to cancel out certain ranges of light waves and strengthen others, changing the color of the light we see.

Interference phenomena are affected by the speed of the light waves through the medium in which they are travelling. Each medium, such as water, air, glass, or oil, has what is called an index of refraction labeled n which is the ratio of the speed of light in a vacuum, c, to the speed of light in the medium, v given by

$$n = \frac{c}{v} \tag{16}$$

A vacuum has an index of refraction $n = 1$ and air can be assumed to have the same index for most calculations. Water, however, has an index of refraction $n = 1.33$ showing that light travels at only seventy-five percent of its usual speed through water. The relationship between the speed v of a wave, its wavelength λ, and frequency f is given by

$$v = \frac{f}{\lambda} \tag{17}$$

When a wave enters a new medium, its speed changes. This means that the wavelength must change because, for a given wave, the frequency is fixed. The new wavelength λ' in the new medium is given by

$$\lambda' = \frac{\lambda}{n}. \tag{18}$$

The index of refraction also determines how the wave is reflected back from the boundary. If, at the boundary of two different mediums, the new medium has a higher index of refraction than the old one, when the wave bounces off of the boundary, it will be inverted ($\delta = \pi$), every trough becoming a crest and every crest a trough. A wave will not be inverted if it reflects off a boundary with a medium having a lower index of refraction. This is important because it makes two different kinds of interference effects possible, one with a phase change, and one without.

When light with wavelength λ strikes a thin film, such as a soap bubble or an

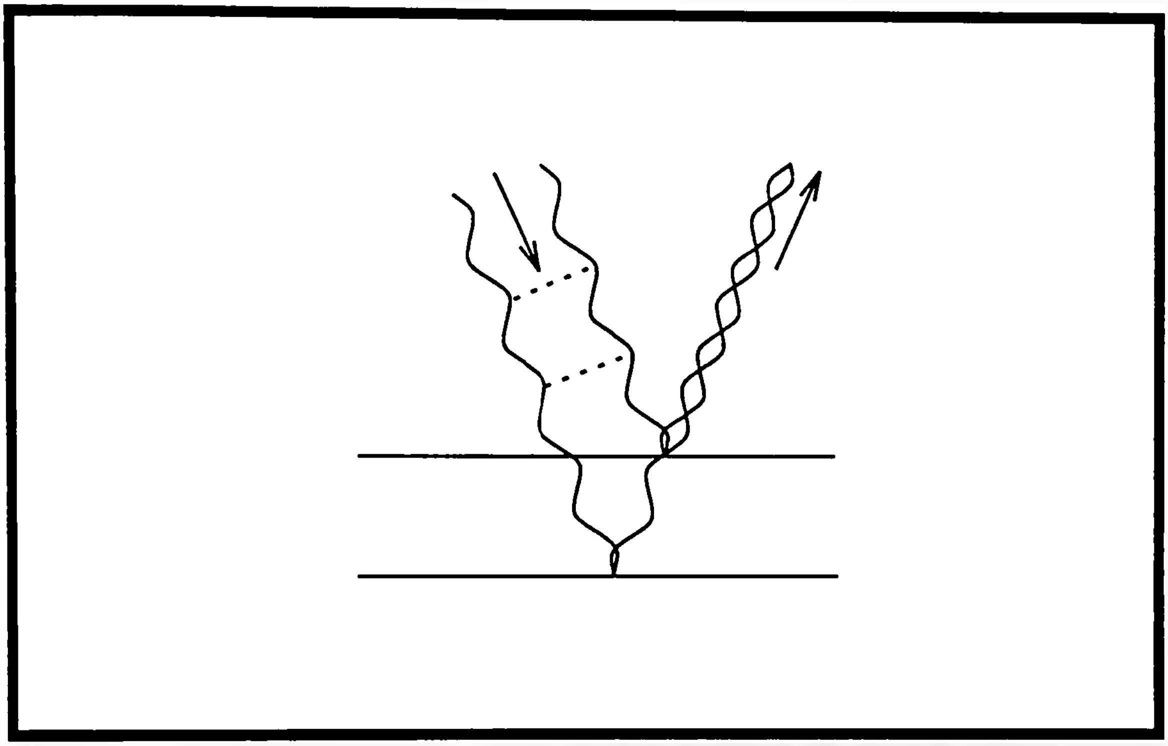

Figure 5: Interference occurs between light wave that has reflected off the bottom of the film and light wave that has reflected off the top.

oil film on water, it is reflected back at both surfaces of the film. If the film has an index of refraction greater than the surrounding medium, the wave will be inverted when it is reflected off of the back face of the film. When the ray reflected from the top surface meets the ray reflected from the bottom surface, they interfere with one another (Figure 5). Assume that the light waves travel perpendicular to the surface and for the moment, neglect the inversion of the wave. Then the extra distance or path difference that the light rays travel is equal to two times the thickness t of the film. There are $2t/\lambda'$ waves within the film where λ' is the wavelength of the light modified by the index of refraction of the surface. If the film has a thickness of $t = 0$, then there are $2 * 0/\lambda' = 0$ waves within the film and the phase difference is 0. If the film has a thickness of $t = \lambda'/4$ then there are $(2 * \lambda'/4)/\lambda' = 1/2$ waves within the film and the phase difference is π. Because of the phase change that happens to the light reflecting off of the back boundary, we need to add a constant of π to these numbers. From these relations, the phase difference of light reflecting off of a thin film is

$$\delta = 2\pi\left(\frac{2t}{\lambda'}\right) + \pi \tag{19}$$

and the amplitude A_r of the resulting wave is

$$A_r = 2A\cos\left(\frac{\delta}{2}\right) = 2A\cos\left(\frac{2\pi tn}{\lambda} + \frac{\pi}{2}\right) \tag{20}$$

where A is the amplitude of the wave of light reflected off of each surface (Tipler 1976). For a film with thickness t of about 0, the resulting amplitude is close to 0 for all wavelengths of light. This causes the top of a soap bubble to appear black, because this is where the bubble is thinnest. This shows that the waves are inverted when they reflect off of the back face and cancel out almost all of the light reflecting off of the front face.

Given this information we can calculate the spectral energy distribution of the light reflecting off of a thin film. The intensity of the wave is equal to the square of its amplitude $I = A^2$. We define the reflectance $\rho(\lambda)$ of a film with thickness t and index of refraction n to be

$$\rho(\lambda) = [K\cos\left(\frac{2\pi tn}{\lambda} + \frac{\pi}{2}\right)]^2 \tag{21}$$

where K is a constant that measures how much light is reflected at each boundary of the film.

Some interference phenomena are caused by a film with air on one side and some other medium with a higher index of refraction than the film on the other side. In this case, the light waves are not inverted when they are reflected off of either the back boundary or the front boundary. The constant π that was added to the previous equation for a film with mediums of lower refractive index on each side should not be used for this case. For this type of interference phenomena, the reflectance $\rho(\lambda)$ for a film with thickness t is defined to be

$$\rho(\lambda) = [K \cos\left(\frac{2\pi tn}{\lambda}\right)]^2 \tag{22}$$

where K again is a measure of how much light is reflected at each boundary of the film.

In both of these types of interference phenomena, the path difference is affected by the angle at which the viewer is looking at the film. When light strikes a film at an angle, the exit point of the part of the wave that enters the film is different from the entrance point. At the exit point, the wave interferes with a different wave that strikes the film surface at this point and reflects off. The light travels farther through the film than before but the path difference ends up being less. This is because the light wave that the exiting ray interferes with has also traveled farther. The light wave that enters the film travels a distance of c before the corresponding point on the other waver reaches the film. The path difference between the two waves is $a + b$ (Nassau 1983) which is equal to $2t \cos\theta$ (Figure 6). For a perpendicular wave, the

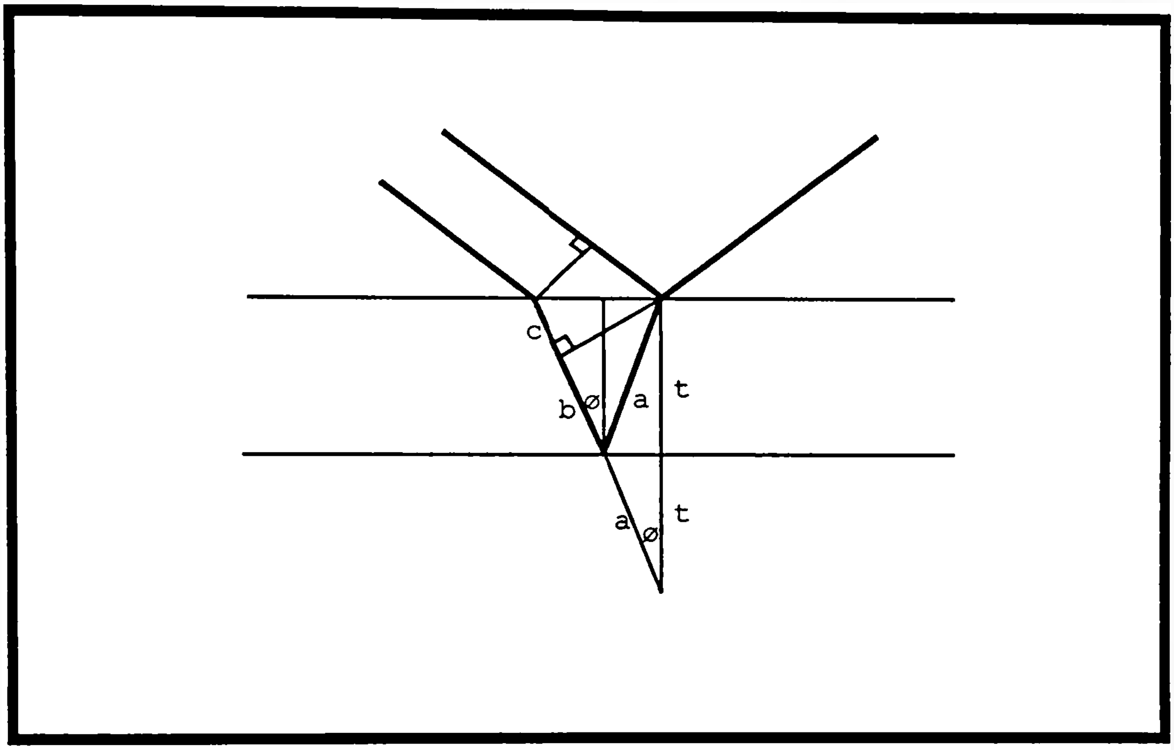

Figure 6: Angle of view causes to a difference in path length $a + b$ that is equal to $2t \cos\theta$.

path difference is twice the thickness. So the the effective thickness t' of the film that gets used in equation 22, the formula for light striking perpendicular to the film is

$$t' = t \cos\theta. \tag{23}$$

3.3 Calculating Newton's Colors

In order for interference effects to be used in computer graphics, it is necessary to determine the spectral energy distribution of the light reflected from the film and to convert the spectral energy distribution to the proper point in the RGB color space of a color television monitor. (A previous attempt to model interference phenomena (Watt 1989) incorrectly performed the calculations directly in terms of the RGB

monitor primaries.) Given the spectral energy distribution $E(\lambda)$ of the light source, and the expression for $\rho(\lambda)$ given in equation 22, the spectral energy distribution can be expressed as a function of t the thickness and n the index of refraction of the film as well as λ the wavelength. Using the matching functions $x(\lambda)$, $y(\lambda)$, and $z(\lambda)$ for XYZ space, we can convert this spectral energy distribution into XYZ tristimulus values and then into into the RGB space of the monitor (Meyer 1987).

For some ranges of interference colors, the RGB tristimulus values become negative. This occurs where the film is relatively thin because in some cases, only very narrow bands of wavelengths have a high intensity. These ranges are closer to monochromatic colors and cannot be displayed on a monitor. The process of deciding what color to display for these out of gamut colors will be covered in Chapter 4.

A table of RGB values for interference effects with a phase change and one for effects without the phase change is precomputed and is indexed by the thickness of the film. This is done so that the expensive processes of calculating the color need only be done once for each thickness. Note that the table can be precomputed for an index of refraction of 1 (Nassau 1983). When a film with thickness t and index of refraction n seen at an angle θ the functional thickness t' is

$$t' = tn\cos\theta \tag{24}$$

Once the final thickness t' is determined, the color for that thickness is looked up in the table. This makes the rendering much more efficient.

3.4 Multiple Film Interference

Another interference phenomena occurs when light strikes a surface that is composed of multiple uniform layers of a transparent substance. This multiple film interference phenomenon causes iridescence, the bright colors on the feathers of peacocks and some other birds, as well as the bright colors on some insect wings and opals. Each film of the surface tends to reinforce certain ranges of wavelengths more than others. The more films that are present, the narrower and brighter the range of reflected light. This repeated reinforcement of limited ranges of wavelengths causes the very bright and almost spectral colors that are found in multiple film interference phenomena. The intensity I at each wavelength λ of light is

$$I = \frac{1}{2}K\left(\frac{\sin 2\pi(l+1)t/\lambda}{\sin 2\pi t/\lambda}\right)^2 \tag{25}$$

where K is the amplitude of the light striking the film, l is the number of layers, and t is the thickness of a layer (Anderson 1942). One of the properties of iridescence is that the colors change as the angle of incidence changes. The effective thickness is a function of θ and n, the index of refraction. The derivation of the change with angle is much more complicated for multiple film interference than for single film interference because of the interaction between films. The effective thickness t_θ is

$$t_\theta = t(n^2 - sin\theta)^2. \tag{26}$$

Often in nature the results do not follow these formulae exactly because the films may not be uniform or the films are composed of interlaced layers of different materials.

3.5 Modeling Oil

When a drop of oil falls onto a surface it often forms a pattern of concentric circles. On a completely smooth and level surface the oil would spread slowly out over the surface, staying thicker in the middle and becoming very thin at the edges. It forms a mound on the surface with the thickness of the oil being a function of distance from the center. Based on these observations and an illustration from (Minnaert 1954), the thickness was modeled as an exponential function of the square of the distance from the center point. For an oil spill centered at the origin, this formula is

$$t = e^{-k_1(x^2+y^2)}, \tag{27}$$

which is a normal distribution with k_1 a constant that determines how quickly the oil diminishes to a negligible thickness.

This method is a good approximation to an oil slick on a perfectly smooth surface. Oil, however, is most commonly seen on pavement and does not spread out evenly due to the roughness and slope of this surface. Using a turbulence function from Perlin (1985) (see Appendix A), the distance from the center can be randomly perturbed at each point in a way that keeps the oil slick roughly concentric. The resulting formula is as follows

$$t = e^{-k_1(x^2+y^2+k_2 \mathrm{turb}[x,y])} \tag{28}$$

where k_2 is a constant that affects how irregular the spreading of the oil is. The turbulence function returns a self-similar pattern of perturbation which gives a good visual impression of a turbulent flow (Perlin 1985). This formula approximates how each area of the oil slick spreads out at different rate. Further distortions could be done to t in order to model the effects of a sloped surface or other irregularities, but this formula serves as a good representation for a basic oil slick. For some point on the surface, the thickness t is computed based on the distance from the center and the value of the turbulence function at that point. The thickness t is then changed based on the angle between the viewer and the surface as in equation 23. The resulting t is used as an index into a precomputed table of colors and the color in RGB space is finally determined for that point.

In order to get a good picture, the slick needs to lie on an appropriate surface. The most familiar surface for an oil slick is pavement, but pavement is a very complex surface because of all of its irregularities. Using functional textures from Perlin (1985), a simple bumpy surface was first used. This surface did not do a good job of modeling pavement because the bumps were all of the same magnitude and distance apart. To get a more realistic effect, a surface that had an irregular scale of bumps and an irregular amount of roughness to it had to be approximated. This was done by using two separate calls to a noise function (see Appendix A) to determine the scale and the roughness of the final texture. After doing this distortion, very fine bumps were mapped over the result to make the surface appear even rougher (Figure 7).

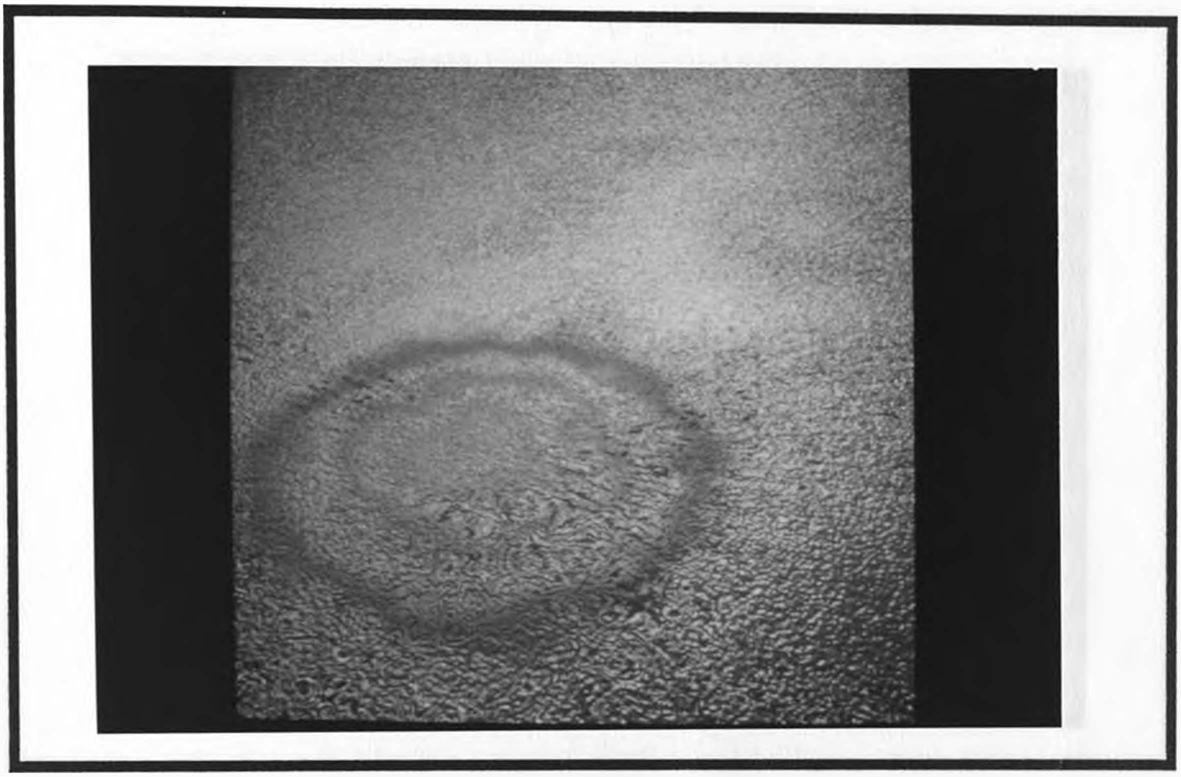

Figure 7: Picture of oil slick generated with physically based modelling techniques.

3.6 Modeling Soap Films

A soap film is another example of interference phenomena. The bright colors that reflect off of a soap bubble are caused by the very thin film of the bubble. As with the oil slick, the light changes phase when it strikes the back face of the film.

Soap films are affected by gravity. The water in the film is being pulled down by gravity and up by the surface tension in the film. In still air the film is thicker near the bottom than the top, and therefore the color at each point on the film is a function of its height. This relationship causes bands of color to appear on a soap film after it has stabilized. In modeling soap bubbles, height was normalized to a zero to one range, the result from a call to the turbulence function was added to the

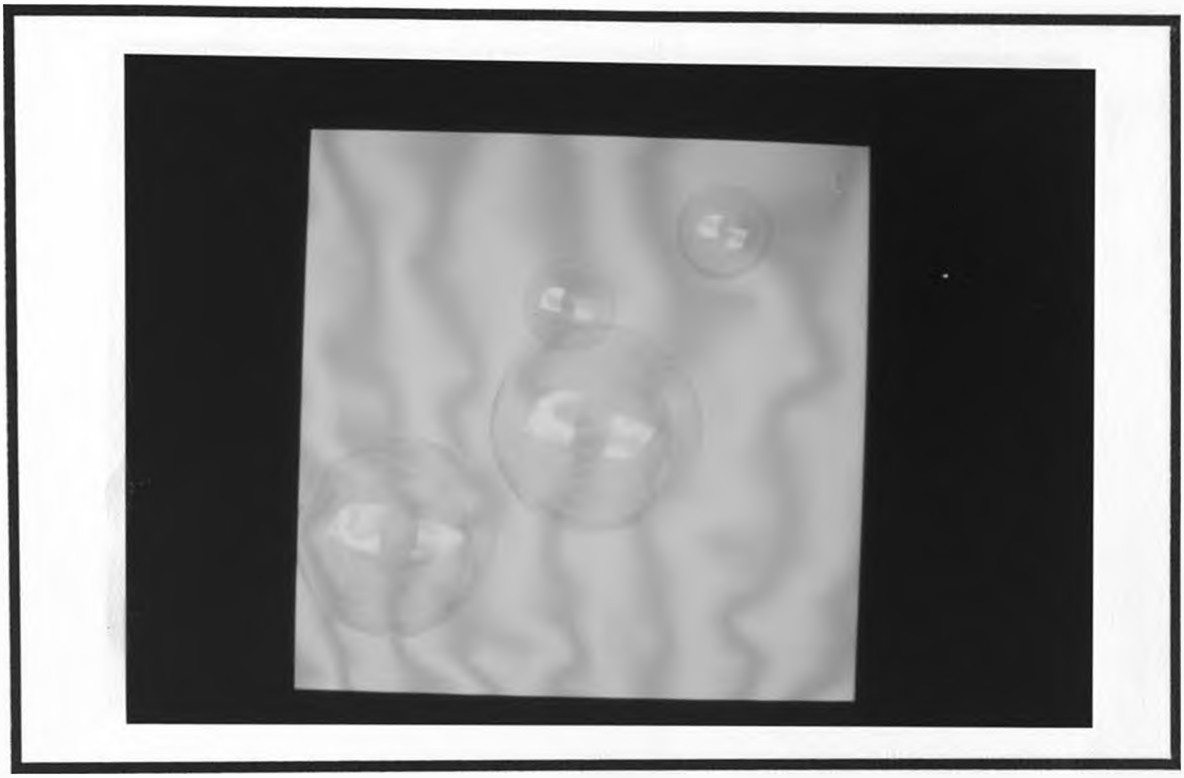

Figure 8: Picture of soap bubbles generated with physically based modelling techniques.

height, and the thickness was taken as being directly proportional to the height. The turbulence approximates the effects of air currents and other irregularities that affect the thickness of the film. Using a simple linear function for the thickness is probably incorrect, but because of the large amounts of turbulence that are added, it is not noticeable (Figure 8).

3.7 Modeling the *Morpho* Butterfly

Iridescent effects are found on the wings of *Morpho* butterflies. The vanes of this butterfly's wings have twelve horizontal mullions which create a multiple film interference effect. The spacing of these mullions was found to be about 155nm (Anderson 1954).

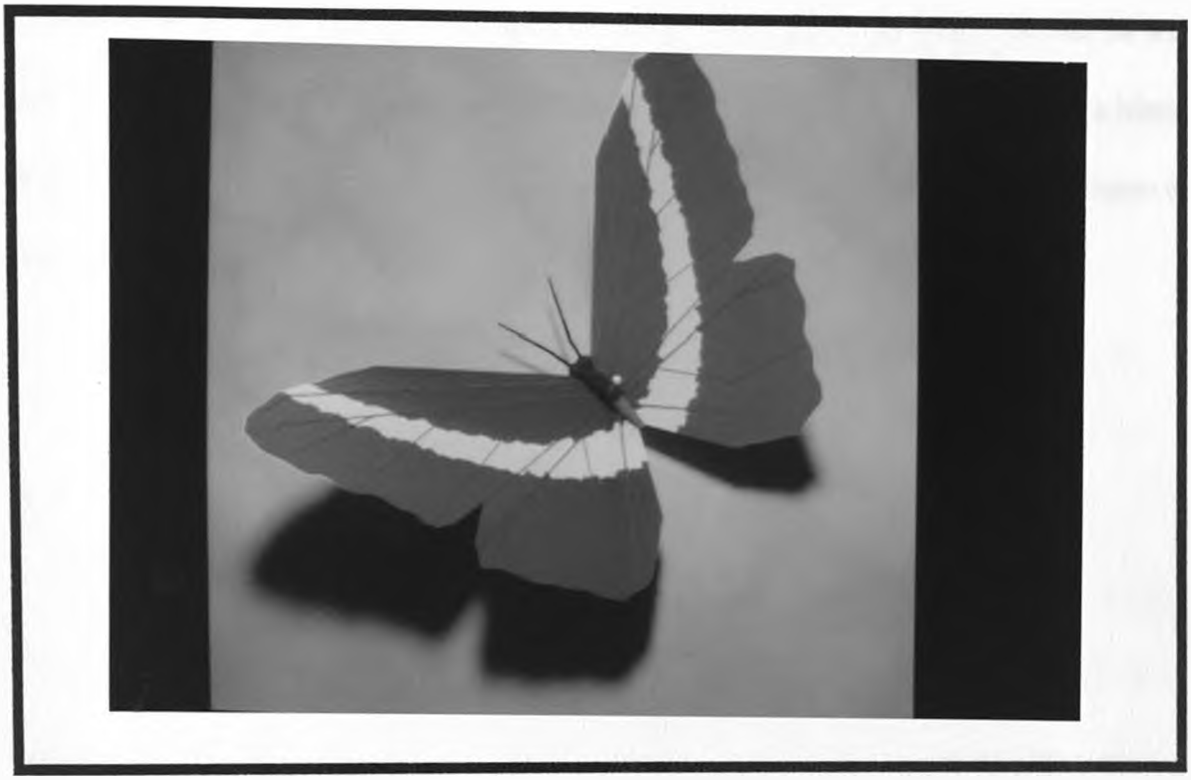

Figure 9: Image of a *Morpho* butterfly showing the iridescence of the wings.

The mullions have some thickness as well, which accounts for the perceived specular maximum at approximately 400nm, implying a thickness of 200nm. Because of the thickness of the mullions, the effective thickness t_{theta} of the film is a function of the angle between the incident ray and the surface normal, θ, the spacing between the mullions d_1, the thickness of the mullions d_2, and the refractive index of the mullions, n (Anderson 1954)

$$t_\theta = d_1 \cos\theta + d_2(n^2 - \sin^2\theta)^{\frac{1}{2}}. \tag{29}$$

A table referenced by angle was generated for $d_1 = 150$ and $d_2 = 50$. This table was used to determine the color at each point on the butterfly's wings. The surface was made slightly bumpy to keep the wings from appearing perfectly smooth (Figure 9).

The mullions are not always the same thickness, they vary somewhat between 50 nm and 150 nm, so the actual colors reflected from the wings of the butterfly are a blend of multiple film interference colors resulting from the variations in the thicknesses of the mullions (Anderson 1954).

4 Out of Gamut Colors

4.1 Introduction

The gamut of a color device is the set of colors that it can produce. The primaries used in the device determine the gamut and when the primaries of two devices differ, the gamuts do not match up exactly. This gamut mismatch problem causes ranges of color on one device that can not be reproduced on the other. These colors are called out of gamut colors.

Gamut mismatch occurs in several different circumstances. One of the situations in which gamut mismatch is a major problem is between color monitors and color printers. The vivid, highly saturated colors that a monitor can display well are impossible to reproduce using a color printer, because color printers use different primaries and a different process for creating color. This problem makes the final printed results look much different than the display on the monitor. Gamut mismatch also occurs between monitors, causing dissimilarity between the same picture seen on different monitors. This is because the red, green, and blue phosphors vary slightly between monitors. The gamut mismatch that resulted from modeling interference phenomena was a third mismatch, that between nature and the monitor. Interference effects produce colors that are almost spectral and which lie outside of the gamut of almost all reproduction devices.

In all of these gamut mismatches, adjustments need to be made to the the colors that are out of the destination gamut. The color that is chosen as a replacement for the out of gamut color must be in the destination gamut and should be as close as possible to the original color. One potential solution is to clip the colors outside the destination gamut to the boundary of the gamut. This involves finding a color on the boundary of the gamut that is close to the original out of gamut color. Another idea is to compress the input gamut and all colors in the image down to fit within the destination gamut. This involves scaling the colors down by some amount that is determined by the input gamut.

4.2 A Solution to Out of Gamut Colors

In a recent paper (Gentile 1989) thirteen different techniques for dealing with out of gamut colors were compared. All but one of these techniques (the control) were performed in the perceptually uniform $L^*u^*v^*$ space. The control was done in RGB space and used the Euclidean distance formula to select the closest color to the out of gamut color. Both clipping and compression techniques were tested, keeping various subsets of lightness, saturation, and hue constant while performing the color adjustment. Some of the compression techniques analyzed the ranges of colors in the image before deciding what to do. Considering only techniques which did no analysis on the picture, the best techniques were clipping with a constant lightness, or clipping with a constant lightness and hue. Only one of the image dependent techniques was rated better than this. Image dependent techniques are usually not feasible because they have to analyze the entire picture before they do anything to bring the out of gamut colors back inside the gamut.

The problem with all of the techniques investigated in Gentile is that they operate in $L^*u^*v^*$ space. $L^*u^*v^*$ space is not computationally efficient for large pictures because of the non-linear transformation that is required for every color in the image. As pictures get larger and people demand faster printing or display with better quality, the problem will only become more acute.

What is needed is an algorithm which will work in XYZ space or the RGB space of a particular monitor and still have the same properties as the constant lightness and hue $L^*u^*v^*$ space algorithms. By holding lightness and hue constant, what is changing is the saturation of the color. The color is being pulled in towards the neutral colors that lie on the diagonal line between black and white in the color space. In XYZ space lightness is a function of Y alone, so we could pull back in towards the neutral diagonal while keeping Y constant (Figure 10). This will desaturate the color and bring it within the gamut of the output device. In XYZ space, colors with constant hue do not lie on a straight line from the neutral diagonal to the edge of the gamut, but are curved somewhat (Cornsweet 1970). No linear transform will make the lines

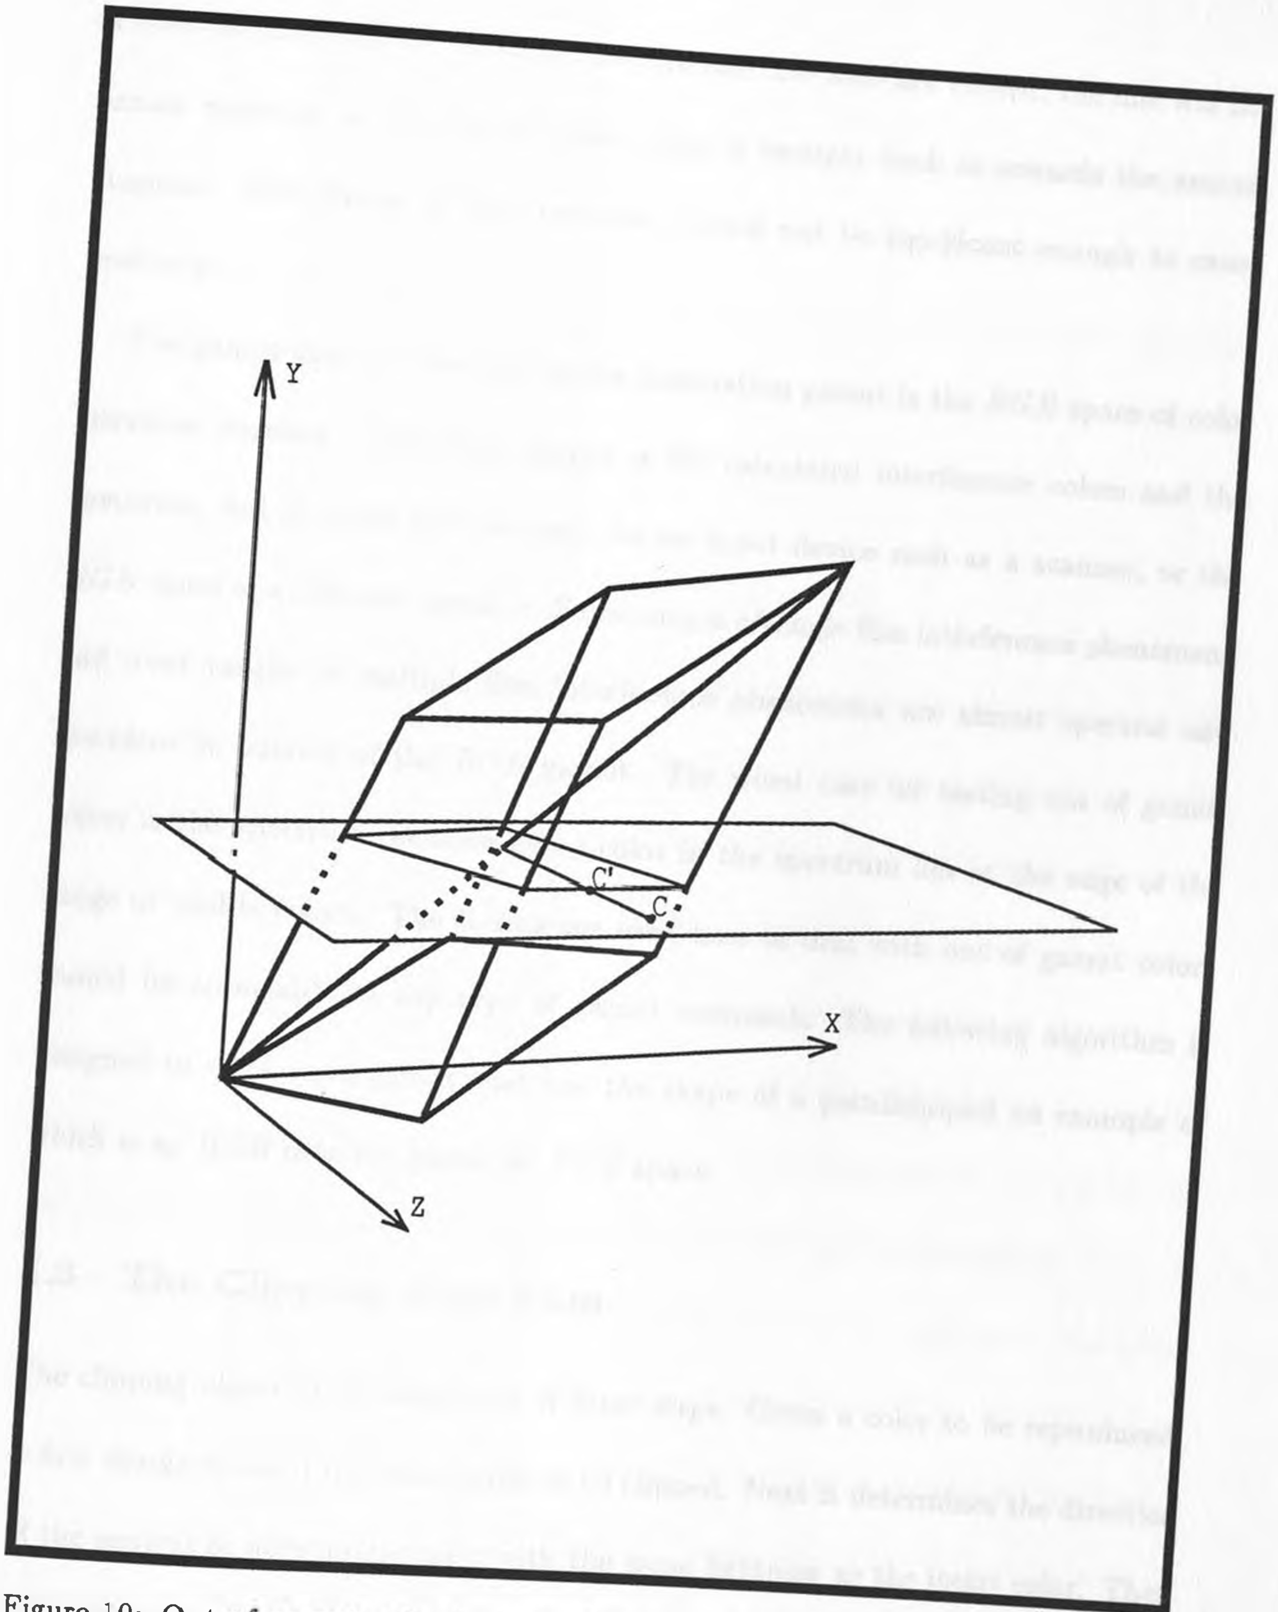

Figure 10: Out of gamut color is approximated by following a line of constant hue and lightness in towards the neutral diagonal.

of constant hue straight. Because the constant hue lines are curved, the hue will not remain constant as the out of gamut color is brought back in towards the neutral diagonal. The change in hue, however, should not be significant enough to cause problems.

The gamut that will be used as the destination gamut is the RGB space of color television monitor. The input gamut is the calculated interference colors and the spectrum, but it could just as easily be an input device such as a scanner, or the RGB space of a different monitor. Some ranges of single film interference phenomena and most ranges of multiple film interference phenomena are almost spectral and therefore lie outside of the RGB gamut. The worst case for testing out of gamut colors is the spectrum, because every color in the spectrum lies at the edge of the range of visible colors. The techniques used here to deal with out of gamut colors should be applicable to any type of gamut mismatch. The following algorithm is designed to work on a gamut that has the shape of a parallelpiped an example of which is an RGB monitor gamut in XYZ space.

4.3 The Clipping Algorithm

The clipping algorithm is composed of three steps. Given a color to be reproduced, it first checks to see if the color needs to be clipped. Next it determines the direction of the neutral or achromatic color with the same lightness as the input color. Then the intersection with the monitor parallelpiped must be found for a ray from the color to the neutral color. Most of the work in this algorithm is done in finding the direction to bring the ray into the gamut and doing the intersections with the gamut boundaries. If the color already exists in the RGB space of the destination, but the values are either too large or negative, a very similar but faster algorithm with identical properties will work. Both of these algorithms are variations on the Cohen-Sutherland clipping algorithm for line segments.

4.3.1 Checking Whether In or Out of the Gamut

The first step involves checking to see if the color lies within the output gamut. The monitor parallelpiped is defined by six planes, three of which pass through the origin, and three of which are parallel to the first three. For each pair of planes, the point can either be above both of them, between them, or below both of them. If the point is between them, it could lie in the gamut. If the point lies outside of the planes it is outside the gamut and it can only intersect the plane that lies between it and the gamut (Figure 10). If for each pair of planes the point lies between the planes, the point must be within the gamut and the color does not need to be clipped.

The symmetry of the gamut makes the comparisons more efficient. The plane equations for the plane through the origin and the plane opposite it are $Ax+By+Cz = 0$ and $Ax + By + Cz = D$ respectively. Because the planes are parallel, A, B, and C are the same for each plane. Checking to see which side of the planes the point is on requires putting the coordinates of the point into the plane equation and checking if the result is less than zero. If the result is greater than zero and less than D for all pairs of planes, the point is within the gamut and nothing needs to be done with it. Otherwise the point of intersection will need to be found for the planes where the result is less than zero or greater than D.

In RGB space, this entire process is made much simpler by the fact that the gamut forms a cube with sides parallel to the axes. This means that the plane equations for the sides have the form, $x = 0$ or $x = D$ with similar equations for the other four sides. Instead of putting the point into a plane equation with three variables, all that is necessary is to check each coordinate of the point against 0 and then against the maximum value for the gamut. The results of the comparisons are stored so that only the necessary intersections will be done.

4.3.2 Finding the Neutral Color

The out of gamut color must be pulled back towards the center of the gamut defined by the neutral diagonal running from black through grey to white. The grey point lies on the line from white to black and has the same lightness as the point being clipped. In XYZ space, keeping lightness constant is easy because lightness is just a function of Y. The neutral diagonal, however, is not the line defined by $x = y = z$ since this is not always white in the output gamut. The white point is the location in XYZ space produced when R, G, and B are at maximum intensity. Given the white point of the output gamut, the intersection of the neutral diagonal with the plane $Y = y_p$ is found, where y_p is the luminance Y of the color being clipped. The vector from the input color to this grey point is used to find the point on the edge of the gamut that is closest to the input color.

In RGB space, keeping the lightness constant is more difficult because lightness is a function of all three values R, G, and B. Using the RGB to XYZ transformation matrix (Meyer 1987), it is possible to find the Y value for the input point. If r, g, and b are the coordinates in RGB space, the Y value is $Y = Ar + Bg + Cb$ where A, B, and C are from the transformation matrix. With Y constant, this is the formula for a plane. The neutral diagonal in RGB space is the line with $r = g = b$. Let w be a point on this line, then $Ar + Bg + Cb = Y = Aw + Bw + Cw$ and $w = \frac{Ar+Bg+Cb}{A+B+C}$. The point with $r = g = b = w$ is the point on the grey line with the same lightness as the out of gamut color.

4.3.3 Intersections

Once the direction along which to bring the out of gamut color back into the gamut is found, it is necessary to find where on the edge of the gamut the point of intersection lies. In the process of checking whether or not the input point was within the gamut, the results told which planes the point could intersect as it is pulled back into the gamut. There is a possibility of up to three intersections that need to be done. The intersections are computed and the one furthest from the input point is the output color. The furthest point is used because otherwise the point will still be outside at least one pair of planes. The intersections done in RGB space are a little simpler because the space is a cube and the planes are parallel to the axes.

4.3.4 Analysis

These algorithms need to be analyzed in terms of both speed and accuracy. Both the XYZ and RGB space algorithms give the same point in the output gamut but it is easier to convert from XYZ space to $L^*u^*v^*$ space and hence easier to analyze the accuracy of the XYZ space algorithm. I tested the clipping algorithms on the spectral colors. This is the worst case since every color in the spectrum is monochromatic and as far out of the gamut of the monitor as possible. Numerically, the lightness remains fixed for all colors that get clipped. In general, the saturation of the clipped colors goes down, and the hue changes only a little. For some ranges of colors, however, the hue changes are more severe. These ranges correspond to regions where the plane defining a gamut edge is closer to parallel with the plane of constant lightness. This fact makes sense because in these regions, the color will travel farther towards the grey line before it intersects the gamut. In most practical situations, however, the input colors will not be monochromatic to the extent that the spectrum is. Visually, the spectrum looks better when clipped so that lightness remains constant and hue changes as little as possible (Figure 11). .

The XYZ space and RGB space clipping algorithms are more efficient than converting a point into $L^*u^*v^*$, clipping, and then converting back. A 1,000 by 1,000

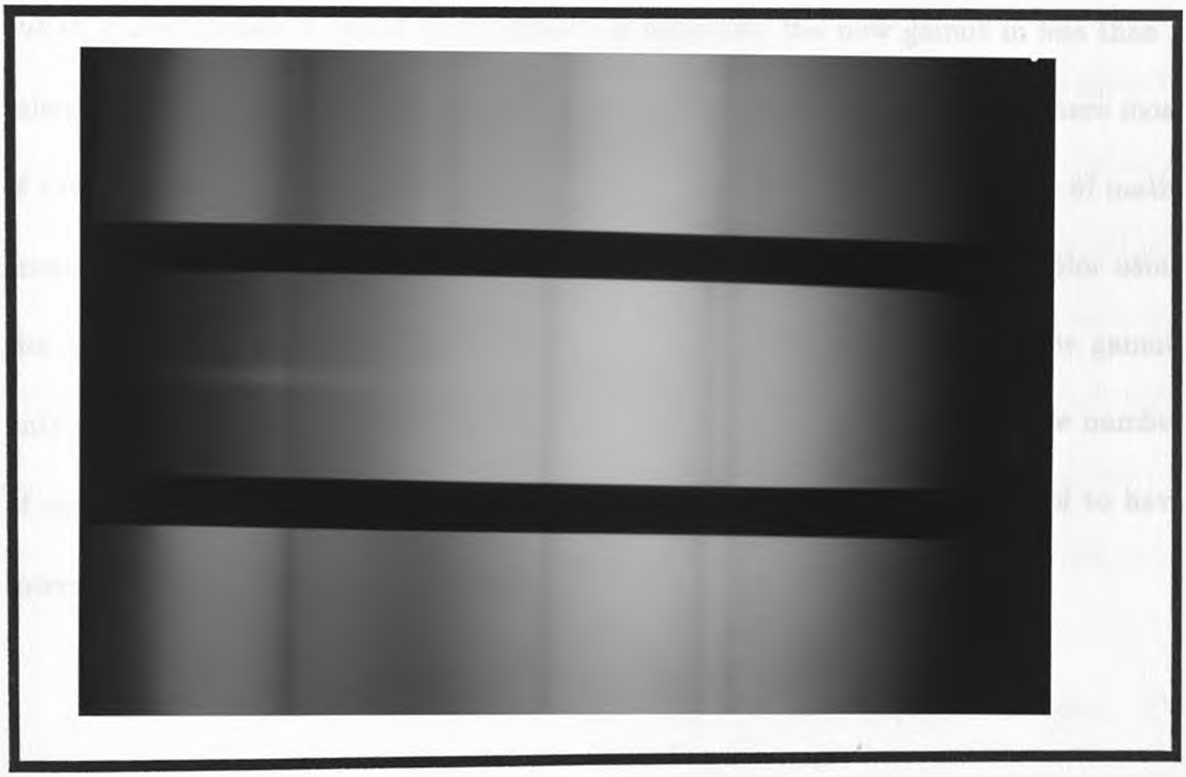

Figure 11: The spectrum clipped three ways (from top to bottom): discarding negative values, keeping hue and lightness constant, and keeping only hue constant.

Section	comparisons	+ or -	*	/
Check if In or Out	3 - 6	0	0	0
Finding Grey Point	0	5	4	0
Finding Intersections	2	4 - 6	3	1-3
Totals	5 - 7	9-11	7	1-3

Table 1: Mathematical operations required to clip a color.

picture contains the possibility of 1,000,000 colors. If it were composed of completely out of gamut colors it could be transformed back into the new gamut in less than a minute based on profiling techniques on a Sparc workstation. A picture where most of the colors are within the gamut takes considerably less time. The number of mathematical operations that are needed to check and clip one out of gamut color using the new algorithm is shown in Table 1. Note that if the color is within the gamut, only the initial six comparisons need to be made. There is a range on the number of operations because there is a possibility of one to three planes that need to have intersection checks performed on them.

5 Conclusion

A physically based approach to modeling the natural world gives good results for interference effects as well as other natural phenomena. Three different examples of interference effects were simulated. In the process, an efficient and perceptually accurate algorithm for clipping out of gamut colors was developed. Because this approach to modelling is consistent with the laws of physics, it provides a much more flexible framework for simulations. By changing the synthesis process from an inaccurate visual approximation to a physical simulation, the simulation can easily be placed into a more complex simulation of nature and expanded to include the interactions of the rest of the simulation with the interference phenomena. The factors in the simulation that affect the surfaces that interference colors occur on can be modelled directly by placing them into the physically based interference model.

The physically based approach generates spectral energy distributions instead of tristimulus values for a specific type of display. Using these device independent techniques to determine color allows for uniformity of color between display devices. A device independent representation for color permits color to be used in design systems and simulations with the assurance that when a change of display device or medium takes place, the color and therefore the simulation or model will remain as accurate as possible. Without this uniformity, the simulation or design loses its validity when it is transferred to another device or medium.

The obvious cost for this technique is an additional layer of complexity in the modeling process because the colors are generated in wavelength space and must to be transformed into RGB space. Another problem occurs because the colors that are produced are not necessarily within the gamut of a particular display device. The colors produced by single and multiple film interference phenomena are often close to spectral colors and therefore out of the gamut of all monitors. This problem of out of gamut colors is important not only in the case of color monitors. Out of gamut colors are a major problem in obtaining a satisfactory hardcopy of a picture displayed on a color monitor because of the difference in color reproduction properties between the monitor phosphors and the printing inks.

Out of gamut colors can be brought back very efficiently into the gamuts of color monitors by desaturating the color until it enters the monitor parallelpiped. This desaturation of the out of gamut color can be done with a variation of the Cohen-Sutherland clipping algorithm. By performing the clipping in RGB or XYZ space instead of the perceptually uniform $L^*u^*v^*$ space, the lightness can be held constant and the color desaturated, but the hue will no longer remain constant. The changes in the hue will not be large even in pathological cases. The advantage is that there are no expensive transformations into $L^*u^*v^*$ space followed by a complicated clipping technique and then the inverse transformation back into RGB or XYZ space. Hardware implementations of clipping algorithms are also becoming very common and it would be a simple modification to place this out of gamut clipping algorithm into hardware as well. Using an algorithm which keeps lightness constant and minimizes hue distortion is an acceptable compromise between accuracy and efficiency.

A Procedural Textures

Texture mapping is a way of adding complexity to surfaces. There are two main types of texture mapping. One type takes a two dimensional set of data, such as a picture of wood grain, and wraps it onto the surfaces of an object. This is analogous to veneering or wall papering a surface. This method has problems when it is applied to arbitrarily shaped objects. The second type is called procedural textures. Textures defined this way take each point on the object and assign surface properties to the point based on a function of the x, y, z coordinates of the point and generally a random factor. This form of texture mapping sculpts the object out of the texture, so there are no problems with trying to wrap the texture onto an arbitrarily shaped object. The surface parameters that can be changed include color, specular color, surface normal, and incident ray.

The randomness of procedural textures often comes from a noise function. The noise function used for the images in this thesis was taken from (Perlin 1985). It is based on a three dimensional lattice of values ranging from zero to one. The lattice defines a continuous field with a roughly regular frequency. In other words there is a maximum number of oscillations between high and low values between any two adjacent points on the lattice. The value $t = \text{noise}(x, y, z)$ at any point (x, y, z) is found by using a cubic interpolation from the values of the lattice at the sixty-four integer points surrounding it. The gradient of this noise function gives a continuous vector field which can also be used to alter the surface characteristics. The gradient of the noise function is called dnoise. For example, a bumpy surface can be simulated by adding some fraction of the vector returned by dnoise to the surface normal at each point.

For some uses, it is convenient to have a function which has $1/f$ or fractal-like properties. This function is called turbulence since its results have some of the qualities of a turbulent flow. The value $t = \text{turb}(x, y, z)$ is a summation

$$t = \sum_{i=0}^{6} \frac{1}{2^i} \text{noise}(2^i x, 2^i y, 2^i z). \qquad (30)$$

The contribution of each term in the summation is scaled by one over the scaling factor for the point. Turbulence can be used in a variety of ways to simulate such effects as clouds and marble.

B References

Anderson, Thomas F. and Richards, A. Glenn, Jr., (1942), "An Electron Microscope Study of Some Structural Colors of Insects," *Journal of Applied Physics, 13,* 748-758.

Cook, Robert L. and Torrance Kenneth E., (1981), "A Reflectance Model for Computer Graphics," *Computer Graphics, Vol. 15, No. 3*, 307-316.

Commission Internationale de l'Eclairage. (1978), *Recommendations on Uniform Color Spaces, Color -difference Equations, and Psychometric Color Terms.* Supplement No. 2 to Publication CIE No. 15.

Cornsweet, Tom N., (1970), *Visual Perception*, Academic Press.

Gentile, Ronald S. and Allebach, Jan P., (1989), "A Comparison of Techniques for Color Gamut Mismatch Compensation," in *Human Vision, Visual Processing, and Digital Display (SPIE Proceedings Series), 1077,* 342-354.

Klassen, Victor, (1987), "Modeling the Effect of the Atmosphere on Light," *ACM Transactions on Graphics, Vol. 6, No. 3*, 215-247.

Meyer, Gary W., (1986) "Tutorial on Color Science," *The Visual Computer, Vol. 2*, 278-290.

Meyer, Gary W. and Greenberg, Donald P., (1987), "Perceptual Color Spaces for Computer Graphics," in *Color and the Computer* edited by H. John Durrett, Boston: Academic Press.

Minnaert, M., (1954), *The Nature of Light and Color in the Open Air*, New York: Dover Publications.

Musgrave, F. Kenton, (1989), "Prisms and Rainbows: A Dispersion Model for Computer Graphics," *Graphics Interface '89*, 227-234.

Nassau, Kurt, (1983), *The Physics and Chemistry of Color: The Fifteen Causes of Color* New York: John Wiley and Sons.

Perlin, Ken, (1985), "An Image Synthesizer," *Computer Graphics, 19*, 287-296.

Thomas, Spencer W., (1986), "Dispersive refraction in Ray Tracing," *The Visual Computer, Vol. 2*, 3-8.

Tipler, Paul A., (1976), *Physics* New York: Worth Publishers.

Watt, Alan, (1989), *Fundamentals of Three-Dimensional Computer Graphics*, Wokingham: Addison-Wesley.

Williamson, Samuel J. and Cummins, Herman Z., (1983), *Light and Color in Nature and Art*, John Wiley and Sons.