

Modeling and Rendering of Metallic Patinas

Julie Dorsey*

Massachusetts Institute of Technology

Pat Hanrahan†

Stanford University

Abstract

An important component that has been missing from image synthesis is the effect of weathering. In this paper, we present an approach for the modeling and rendering of one type of weathering — metallic *patinas*. A patina is a film or incrustation on a surface that is produced by the removal of material, the addition of material, or the chemical alteration of a surface. Oxidation, sulphidization, and painting are examples of phenomena that produce patinas.

We represent a surface as a series of layers. Patinas are simulated with a collection of operators, such as “coat,” “erode,” and “polish,” which are applied to the layered structure. The development of patinas is modulated according to an object’s geometry and local environmental factors. We introduce a technique to model the reflectance and transmission of light through the layered structure using the Kubelka-Munk model. This representation yields a model that can simulate many aspects of the time-dependent appearance of metals as they are exposed to the atmosphere or treated chemically. We demonstrate the approach with a collection of copper models.

CR Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.3.6 [Computer Graphics]: Methodology and Techniques.

Additional Key Words and Phrases: weathering and appearance, material models, time-dependent phenomena, reflection models.

1 Introduction

All materials have an inherent tendency to change in appearance or composition when exposed to the physical and chemical conditions of the surrounding environment. The rate of change is dependent on the material’s characteristics and the degree of its exposure. The deterioration, decay, and change in appearance of materials due to the effects of the surrounding environment are generally termed *weathering*. Specific examples of weathering include the corrosion of metals, efflorescence on stone and brick, fungal attack on organic materials [28], and the wear and tear of everyday life.

Techniques for realistic image synthesis have advanced dramatically in recent years. However, a common criticism of such images is that they look too ideal, and therefore animators and modelers go to great lengths to create a more natural look. An outstanding example of this approach is the techniques to simulate wear and tear in the recent movie *Toy Story* produced at Pixar [9]. Many texture maps per surface were used to model scuffs, dirt, gouges and so on. Unfortunately, this use of texture maps is labor intensive and somewhat ad hoc. Each texture map must be hand-painted and combined using special shaders. It is also very difficult to properly account for

*Room NE43-213, 545 Technology Square, Cambridge, MA 02139.
<http://graphics.lcs.mit.edu/~dorsey>

†370 Gates Computer Science Building 3B, Stanford, CA 94305-4070.
<http://www-graphics.stanford.edu/~hanrahan>

Permission to make digital or hard copies of part or all of this work or personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1996 ACM-0-89791-746-4/96/008...\$3.50



Figure 1: Example of a real patina.

certain effects, such as a spill that crosses the boundary between two patches. There is clearly a need for modeling and rendering tools that make it is easier to create naturally worn surfaces.

In addition to its importance in computer graphics, visualization of weathering effects has broad applicability to a variety of additional fields. Moreover, practitioners in other fields have studied weathering effects for many years, and their theories and studies are good sources from which to draw.

- For many applications it is useful to predict how a material will look in the future. For example, in architecture and preservation, it is important to understand how the surfaces of buildings change over time. Since buildings may stand for hundreds of years, much of their final appearance is dominated by weathering effects [22, 24].
- Conversely, there are applications where modeling weathering interactions is applicable to the inverse problem of understanding the history of an object from its current appearance. Restoration fundamentally involves returning a weathered object to its true initial appearance [7].
- Many materials are created pre-weathered. These materials are treated in various ways to simulate the process of weathering. Since such treatments can be highly desirable (from prewashed jeans to fake antiques), modeling them is important [24].

Our long range goal is to develop easily controllable models of weathered materials for computer graphics. This task involves the identification of the basic physical processes underlying changes in appearance and the development of appropriate computer models. Since weathering involves the action of many environmental forces over time, this will require simulating these processes. Such simulations will also give us the ability to visualize changes over time. We expect that different materials may require different modeling effects. For example, stone, wood and metals weather quite differently because their chemistry and material structures are very different. In this paper we consider only the modeling and rendering of metallic patinas as a starting point.

Metals are particularly susceptible to weathering interactions and often develop a characteristic *patina*. The term patination is generally reserved for effects involving the chemical alteration of the

surface resulting in a change in color. It may describe the results of either deliberately applied craft processes or natural corrosion [16].

Metals in general — and the finished metal surface in particular — begin to change under the influence of the atmosphere, or the local chemical environment, as soon as they are exposed. Copper and its alloys are particularly interesting, as they have broad aesthetic as well as practical applications.

1.1 Previous Work

To date, the simulation of weathering effects has been given little explicit attention in the computer graphics literature. Related work exists in three areas: procedural textures and fractal surface growth models, specific weathering models, and layered surface representations.

Procedural textures [6] can be used to build up complex patterns that often resemble natural effects. For example, the “shade tree” concept of Cook [4] allowed arbitrary procedures to define a different shading model for each surface, as well as lighting and atmospheric optics. Perlin [26] described an entire procedural language to define textures and laid the foundation for the stochastically-generated textures that permeate rendering today. Finally, Turk [30] and Witkin and Kass [33] introduced synthetic texture models inspired by biochemical processes. In this paper, we use procedural textures in a new way: to vary parameters of a physically-inspired model of material properties over time. We draw on fractal surface growth models, used mainly in physics and various branches of engineering. These models are concerned with the morphology of various pre-formed interfaces and with the dynamics of how the morphology develops over time [1].

Starting with procedural texture models, several researchers have attempted to simulate related weathering effects. Becket and Badler [2] modeled surface imperfections through texture specification and generation techniques, which are based on fractal subdivision and simple distribution models. Blinn modeled the appearance of dusty surfaces, given the thickness of the dust layer [3]. More recently, Hsu and Wong [15] introduced functions for simulating dust accumulation that attempt to mimic the dust adherence process. In addition, Miller [23] proposed a set of algorithms for local and global accessibility shading; this type of shading yields visual effects that resemble tarnish on surfaces. We take a different approach — focusing in considerable detail on one particular material and its changes in appearance due to weathering.

A particularly promising method for modeling the appearance of materials is to treat a surface as a set of layers. For example, a common application of RenderMan is to develop surface descriptions as a series of layers [31]. Another area of interest is the modeling of reflection and transmission of light through layered surfaces. Hanrahan and Krueger [13] present a model for subsurface scattering in layered surfaces in terms of one-dimensional linear transport theory. This model is useful for rendering common layered materials such as skin, snow, and dust. In this work, we build on such layered representations by developing a set of tools for modeling and rendering surface and subsurface structures as a function of time.

1.2 Overview

We begin by briefly reviewing the physical basis of metallic patinas. We then present a representation of a metal surface as a stack of layers and propose a collection of operators that can be applied to the layered structure to produce a taxonomy of patination effects. We also discuss how patinas can be modulated according to environmental factors and the surface geometry. Next, we offer a method of rendering the layered structure. Last, we demonstrate the approach on several complex models.

2 Physical Basis of Patina Formation

Natural patinas develop primarily as the result of the process of atmospheric corrosion [22, 25]. The atmosphere alters the surface of

a metal, causing gradual changes that quickly tarnish it with a thin, but uneven, dark coloring, and may eventually convert the whole surface into crumbly mineral products. In this section, we briefly describe the principal causes of natural patination. Similarly, artificial patinas are deliberately induced through various surface treatments designed to mimic and exaggerate the natural processes.

The corrosion process forms a complicated system of reacting layers consisting of the metal, corrosion products, surface electrolyte, and the atmosphere. The exact nature of the surface chemistry underlying the development of surface coatings and the factors that influence how they form and break down are major foci of current corrosion research. Unfortunately, the experimental data to support a physical simulation that would predict the appearance of a metal is not yet available. Therefore, in this paper we present a phenomenological model for the development of patinas, based on observed physical behavior.

2.1 Composition and Formation of Copper Patinas

In order to study the development of metallic patinas, we chose copper as a representative metal. The patinas of copper and copper alloys are classic examples of layered structures. When viewed in cross-section (see Figure 3 [8]) multiple layers are distinctly visible to the eye. Vernon and Whitby identified the principal chemical constituents of such natural copper patinas in the 1930s [32]. More recently, Franey and Davis [8] and Graedel et al. [11] have carried out a series of investigations, which have provided a detailed picture of the basic composition and formation mechanisms of natural patinas on copper.

Clean copper surfaces exposed to the atmosphere quickly form a thin layer of dull brown tarnish that gradually changes with time to a reddish brown color, which is indicative of copper oxide, or mineral cuprite [8, 22]. Once this layer is in place, subsequent layers grow much more slowly. The primary chemical constituents of the patina on copper and copper alloys include copper oxides, sulphides, and inorganic and organic copper salts. Copper sulphide is very dark brown, and the sulphide coating forms integrally with the underlying metallic crystal structure. Thus, the sulphide layer often appears relatively shiny. The copper salts, consisting mostly of sulphates, chlorates, and nitrates, come in a wide variety of colors. Copper sulphate causes the characteristic green color of aged copper. Figure 2 contains a key showing the approximate color of some of the most important copper minerals.

2.2 Dependence on Physical Environment

Figure 2 is a schematic diagram of the patination process as it can be outlined for three environments based on Graedel’s observations [10]. This diagram indicates the varying composition of the patina in the environments and average lengths of time for development.

The left column represents four primary stages of patina growth in a marine environment. The first and second stages, which are common to the other two environments, involve the formation of a layer of tarnish and a layer of cuprite at the copper surface. The third stage is characterized by the formation of several copper minerals, with atacamite and other chlorine-containing substances expected to predominate because of the high chloride concentration near the sea. Small amounts of organic matter are also present within the structure. The final stage of the process augments the patina with similar minerals and binding materials.

On land, a patina of a different composition is produced. Both rural and urban air have moderate to high concentrations of sulfur. Thus, the transition from the second to the third stage involves sulphurization atop the cuprite, with some atacamite being formed as well. The initial form of the sulphur compound is primarily the mineral posnjakite; an organic binder is also present in the patina.

The final stage differs in the two environments. In the rural case, where there are only moderate concentrations of acid and oxidized sulfur in the atmosphere, the formation of sulfates proceeds slowly, and a mixed patina of atacamite and brochantite is typical. In urban

areas, there are higher concentrations of sulfur, favoring the formation of the basic sulfate minerals antlerite and brochantite. Organic compounds are also incorporated into the patina as it forms.

In a marine environment, patinas grow in thickness at an average rate of $1\mu\text{m}$ per year. In a rural environment, the rate is less, about $.5\mu\text{m}$ per year; in an urban environment the rate is slightly greater, about $1 - 2\mu\text{m}$. These rates are highest during the first few years of exposure and approach a lower, almost negligible rate after about 20 years [22]. While the preceding section gives a descrip-




Figure 2: A schematic diagram of the processes involved in the growth of copper patinas in marine, rural, and urban atmospheres.

tion of the development of patinas in three different environments, including the constituents of the patina and general growth rates, the development can be further influenced by a specific surrounding environment. Wetness is perhaps the most important factor in the patination process. Patination occurs more rapidly in areas of the surface that retain stagnant water. Thus, horizontal or inclined surfaces patinate more rapidly than vertical surfaces [11]. In addition, exposure to daylight decreases the rate of patination due to evaporation of surface water.

There are many other factors that affect the corrosion of metals, such as the composition of the electrolyte on the surface and the temperature; variations in surface thickness due to abrasions, polishing, and pitting are also relevant. However, the influence of these factors is not well understood [22], and we do not consider them further in this paper.

3 Modeling

As we observed in the previous section, the development of a metallic patina is a process that proceeds in a complex system of layers. In this section we describe a representation of the structure of a layered




Figure 3: A surface patina as a stack of layers. (a) Micrograph of a copper surface showing the layers. (b) Our abstraction of the layered structure.

material and a set of operators that can be applied to this representation to simulate a variety of effects. By writing a script in terms of these operators, the weathering of the material as a function of time may be simulated.

3.1 Layers and Materials

In our representation, a surface contains a stack of n layers (see Figure 3). The “0”-th layer is assumed to be the base material with an essentially infinite thickness. The total thickness of the stack of layers is assumed to be small relative to the area covered by the layer. Each individual layer also has a thickness, although it may be zero at certain points. We have varied only the thickness as a function of position, but other properties could be similarly controlled.

The remainder of the material properties do not depend on position. Each layer consists of a homogeneous material. The surface of each material has the standard set of surface shading parameters such as diffuse and specular colors and an overall roughness or shininess. In addition each material has two volumetric properties that control how light is transmitted and reflected due to subsurface scattering in the layer. The light reflection and transmission properties of layers and the rendering of the layer structure are discussed in Section 4.

3.2 Operators

The development of a patina is implemented as a series of operators on layers. By controlling the sequencing of these operators, a wide variety of effects may be created. Our current implementation includes the following operators:




Figure 4: The coat and erode operators.

- **coat material thickness thickness-map**

The coat operator adds a new layer of material to a surface. The new layer has a specified maximum thickness that is modulated by the thickness map (see Figure 4).

- **erode thickness thickness-map**

The erode operator removes material from a layered surface. The depth of the erosion may be controlled by a thickness map. The erosion operator proceeds through the stack of layers decreasing the thickness of each layer until the desired amount of material has been removed. This provides a means to cut into a material and expose underlying layers. In principle, the amount of an erosion may depend on the hardness of the material in each layer, although our current system does not store hardness as a material attribute (see Figure 4).

- **fill material height height-map**

The fill operator deposits material up to a given absolute height above the base material. This is somewhat akin to filling all the valleys with water until it reaches a given height. This is a simple way to simulate the deposition of material in cracks and crevices.

- **polish height height-map**

The polish operator removes material until a given absolute height above the base material is reached. The result is a smoothing effect across a surface. A variation of the operator removes material until a given area of surface is exposed. This version of the operator was motivated by observing profiles and statistics of rough surfaces after they have been polished [29].

- **offset radius**

The offset operator applies a material to a surface by first applying a thick coat and then removing the part accessible to a sphere of a given radius. The offset surface is computed using techniques described by Miller [23]. (We have implemented both versions of his operator: one computes the largest tangent sphere, the other forms a positive and then a negative offset surface and compares that to the height of the original surface).

We have implemented a simple scripting language to describe the sequential application of different operators. This can be used simply to assemble a layered surface from a set of materials and texture maps. In this case each material and thickness map represents the current characteristics of the layered surface. This makes it possible to give the user control over the ordering of various operators and special effects such as polishing. Alternatively, the operators can be repeatedly applied to simulate the action of the environment on the surface. When used in this way the scripts may be generated by a program that attempts to cycle through various operations occurring through time. In Section 5 we will show several example scripts.

3.3 Texture/Thickness Maps

The operators we have described are able to model a multi-layered patina, but one that is completely uniform and devoid of the variations and richness of detail that any natural process generates. Next, we will explain a few simple, physically plausible methods to modulate these operators across the surface, simulating the detail present in the real aging process. In our system we have used two representations for spatially-varying thickness maps:

1. **Rectangular texture maps.** In this implementation the thickness is controlled by a standard 2D texture map.
2. **Triangulations as texture maps.** This implementation represents thicknesses as values attached to the vertices of a triangulation.

In order to simulate variations in thickness over time, we have implemented a series of fractal surface growth models for use with the layer structure. Many of these approaches are variations of models from the book by Barabási and Stanley [1]. We chose a set of models that deals with the deposition of material and the lateral growth of patches on surfaces. These models are particularly appropriate for modeling the types of patterns that emerge in the corrosion process, as corrosion typically begins with the deposition of water and particulate matter from the atmosphere onto a surface. Patches develop and spread based on the amount of water and other substances on the surface. Our implementation of these models includes several growth rates, such as linear, parabolic, and logarithmic. We provide the following models:

Steady thickening (ST). This model creates a very simple, relatively uniform pattern, which increases in thickness with time. We

sample a surface evenly with a small number of points. An initial thickness is assigned to each of these points, and the thickness at intermediate points is interpolated. Over time, the thickness at each sample point is increased by a user-specified growth rate. A small amount of noise is added to the pattern to keep the appearance natural.




Figure 5: (a) Random deposition; the bent arrows indicate RD with surface relaxation. (b) Ballistic deposition.

Random deposition (RD). Random deposition is the simplest growth model that we utilize. From a randomly chosen position over the surface, a particle falls vertically until it reaches the top of the surface under it, whereupon it is deposited. We implement this model by randomly choosing a position i and increasing its height $h(i, t)$ by one, where t denotes the time step. A variation of this model, random deposition with surface relaxation, allows the deposited particle to diffuse along the surface up to a finite distance, settling when it finds the position with the lowest height (see Figure 5a). Due to the relaxation process, the final interface — or surface of the layer — will be smooth, compared to the model without relaxation, which is extremely rough.

Ballistic deposition (BD). In the Ballistic deposition model, a particle is also released from a randomly chosen position above a surface and follows a straight vertical trajectory until it reaches the surface, whereupon it sticks (see Figure 5b). The height of the interface at that point i is increased to $\max[h(\text{neighboring_points}, t), h(i, t) + 1]$. Growth is defined quantitatively through a simple function that calculates the *mean height* of the surface. If the deposition rate (number of particles arriving at a position) is constant, the mean height increases linearly with time. In addition, the *interface width*, which describes the roughness of the surface, is defined by the rms fluctuation in height. The key difference between the RD and BD models is that the RD interface is uncorrelated — i.e. the thickness at each point on the surface grows independently, since there is no mechanism that can generate correlations along the interface. In BD, the fact that particles are capable of sticking to the edge of neighboring points leads to lateral growth.




Figure 6: A lattice of blocked and unblocked cells (left), early stage of pattern formation (middle), late stage of pattern formation (right).

Directed percolation depinning (DPD). Starting with a collection of initial patches on a surface, we develop an interface that grows in all directions in two-dimensions and increases in thickness. We begin with a simple 2D lattice and mark a percentage of cells as *blocked* and others as *unblocked*. Over time, patches are much more likely to advance onto an unblocked cell than a blocked cell (the rules for it are much less stringent). In our version of the model, unblocked cells imply concentrations of moisture on the surface and thus stimulate growth.

Figure 6 shows a representative, initial lattice (where black indicates blocked and white indicates unblocked) and a collection of patches at early and mature stages. In this case, the seed patches are 10×10 squares. The patches grow in the direction of their principal neighbors (cardinal directions) based on a simple probability function and according to whether the given cell is blocked. Note that the pattern develops much more readily in the bottom half of the image, reflecting the lower percentage of blocked cells. Over time as the initial patches spread across the surface, the system inserts additional seed patches at unoccupied positions according to a given probability; these patches grow according to the same rules. The above process creates the bottom level or overall pattern. The upper levels (accounting for variations in thickness) are filled in as follows: the probability of a cell appearing in level $l + 1$ is proportional to the number of cells in level l that support it. Variations in thickness are denoted by gray levels; white is solid and black is void. Other initial patterns and lattices could be used, and such components could be varied as a function of time.

In our current system the development of the layer structure through time is controlled in two ways. In addition to scripting different operations, specific growth models, as described above, can be used to generate thickness maps through time. These growth models can be controlled by environmental factors and the geometry of the surface. Finally, the layer structure is output as a series of material properties and texture maps, which are passed to the rendering system. This allows additional control over the final appearance of the surface.

The use of development models represents something of a black box in that it is possible to use other procedural models or scanned patterns to generate the thickness maps. In addition to the approaches described above, we also make use of standard procedural texturing approaches, such as noise and turbulence functions [6]. We have found the above models to be especially useful. However, as additional information is learned about the development of patinas, more exact models could be used, without affecting the overall patina modeling framework.

4 Rendering

The appearance of a layered structure is a result of light interacting with the surface and subsurface. Surfaces arise at the interfaces between layers. Light is reflected from the surface depending on the surface roughness and the specular color. The surface reflection gives the layer a glossy appearance. Light is also reflected from and transmitted through the interior depending on the absorption and scattering properties of the material. The subsurface reflection gives the layer a matte or diffuse appearance. The fact that light is also transmitted through the stack of layers causes the colors of different layers to be mixed; it also allows underlying materials to remain partially visible.

4.1 Single Layer

To model reflection and transmission through a single layer we use the Kubelka-Munk (KM) model [21]. The KM model is widely used in the paint, printing, and textile industries to compute diffuse colors due to subsurface scattering [17]. The book by Kortum is an excellent source of information about the KM model [18].

The KM model was recently introduced to computer graphics by Haase and Meyer [12]. In that work the color of a thick layer of paint consisting of several pigments was quantitatively modeled. Given the relative concentrations of several pigments, they were able to predict the final appearance of the paint. The KM model was compared to the standard additive or subtractive color models used in computer graphics and was found to be more accurate. They also used the KM model to estimate the pigment concentrations needed to match a given color. In this work we use the KM model to predict the color effects due to variations in layer thickness and to predict the color of a stack of layers of different pigments.




Figure 7: Positive and negative flux density within a layer.

The KM model corresponds roughly to 1-dimensional volume radiosity. Figure 7 shows the transport of light within a layer. The KM model assumes that the light distribution is directionally isotropic, but varies as a function of depth. The light distribution therefore is described by the flux density (energy per unit area), or volume radiosity, in the inward or positive, and outward or negative directions, denoted by B_+ and B_- .

As light propagates through the volume it may be absorbed or scattered, according to the following coupled differential equations

$$\begin{aligned} \frac{\partial B_+}{\partial z} &= -(K + S)B_+ + SB_- \\ \frac{\partial B_-}{\partial z} &= SB_+ - (K + S)B_- \end{aligned}$$

where K is the absorption per unit length, and S is the backscattering per unit length. Backscattering decreases the energy flow in one direction, thereby causing an increase in the opposite direction.

These equations have an analytical solution for a homogeneous layer of thickness d [19]. Given these solutions it is possible to compute the reflectance R and transmittance T through a layer:

$$\begin{aligned} R &= \frac{B_-(0)}{B_+(0)} = \frac{\sinh bSd}{a \sinh bSd + b \cosh bSd} \\ T &= \frac{B_+(d)}{B_+(0)} = \frac{b}{a \sinh bSd + b \cosh bSd} \end{aligned}$$

Here $a = (S + K)/S$ and $b = \sqrt{a^2 - 1}$. Note that in general K and S are functions of wavelength, hence the calculations involved in the KM model must be performed separately for each color sample.

The major problem in using the KM theory is the determination of the layer parameters K and S . One method is based on the formula for the reflectance of an infinitely thick layer:

$$R_\infty = 1 + \frac{K}{S} + \sqrt{\frac{K^2}{S^2} + 2\frac{K}{S}}$$

This formula can be inverted and the ratio K/S computed from R_∞ . Thus, measuring the reflectance of an optically thick sample allows one to compute K/S . Another common method for determining the KM parameters is to measure the reflectance of a thick layer and then the reflectance of a layer of known thickness over a background material of known reflectance. A good discussion of different methods for determining the KM parameters in the context of computer graphics is contained in Haase and Meyer [12].

In our work we have estimated the KM parameters by matching synthesized color samples to physical copper and mineral specimens, as well as photographs, such as those found in Hughes and Rowe [16]. We generally match R_∞ to the sample, determine K/S from the above formula, and assume a value for S . This approach has enough accuracy to capture the range of colors and color effects that we are attempting to model. A more careful study involving detailed comparisons with physical samples under controlled weathering conditions would require more accurate radiometric and colorimetric measurements. However, this is beyond the scope of this study.

4.2 Multiple Layers

Kubelka also extended his model to account for the reflectance and transmittance, including all scattering events, due to a stack of layers [20]. The reflectance and transmittance of two layers may be merged to yield an equivalent reflectance and transmittance:

$$R = R_1 + \frac{T_1^2 R_2}{1 - R_1 R_2}$$
$$T = \frac{T_1 T_2}{1 - R_1 R_2}$$

This compositing process can be repeated, sequentially combining two layers into a single layer, to account for the reflection and transmission through a stack of n layers. We normally do this front to back, although it can be done in any order. Thus it is essentially no more difficult to compute the diffuse reflectance and transmittance, including all possible transport paths, for a stack of n layers than for a single layer. These formulas may also be used to combine a stack of layers with an opaque base layer. (Note that most formulations of the KM model explicitly include composition with a background layer R_g .)

We call this process *subsurface compositing* to differentiate it from the normal compositing operators widely used in computer graphics [27]. Subsurface compositing differs from normal compositing in that visible light is always assumed to be transmitted through each layer twice, which gives rise to a T^2 factor in front of each reflectance. This lowers the visibility of underlying layers. Subsurface compositing also accounts for the effect of multiple scattering which leads to a $1/(1 - R_1 R_2)$ factor which increases the contribution from each layer. These are important effects when viewing a stack of adjoining layers; it differs from the typical case in computer graphics where a set of disjoint partially transparent reflecting surfaces are being viewed. In this case it is incorrect to only consider one-dimensional interreflection, and the light source is not necessarily emanating from the eye.

4.3 BRDF

The above techniques are used during rendering to compute an approximation to the BRDF of each point on the layered surface. We assume the final BRDF consists of a diffuse term and a mixture of glossy terms. The diffuse reflectance of a stack of layers is computed exactly as described above, and needs no further explanation.

Accounting correctly for glossy reflection is more complicated, so we make several assumptions to simplify rendering. Glossiness is caused by light reflection at the interface between layers. Traditionally in computer graphics there is only a single interface between air and the material — the surface itself — and this interface causes the glossy appearance; in a layered structure there are multiple interfaces. In our system we perform the following two steps. First, glossy reflection at each interface is modeled using the standard computer graphics model for shiny surfaces; that is, with a specular color and a microfacet distribution function parameterized by the surface roughness, $C_s(N \cdot H)^{1/r}$. Each interface inherits its properties from the material below it, hence, each interface may have a different roughness and color. Second, to account for attenuation due to absorption and scattering in the layers above the interface, the specular color is multiplied twice by diffuse transmittance of the intervening layers calculated using the KM equations. The output glossy BRDF consists of a set of surface roughnesses and attenuated specular colors. The renderer needs to be modified to sum over this set of microfacet distributions when computing the final glossy reflection component. If the final specular color of one of the glossy terms is 0, that usually implies that it has been covered by an opaque material.

This method of computing glossy reflection is an approximation for several reasons. First, it is not strictly correct to use the diffuse transmittance to attenuate the incoming and outgoing radiance because even with a homogeneous layer the transmittance will be a

function of direction. Second, the above model does not account for scattering events in the intervening layers before or after the glossy reflection. In general, accounting for this would require the full solution of the one-dimensional transport equation, as was done in Hanrahan and Krueger [13]. The precise nature of the errors introduced by these approximations needs further study. But our approximations capture quite well all the major visual effects that we hoped to achieve.

Finally, we output the final thickness of the stack of layers by summing the thicknesses of each layer. This final thickness may be used to either perturb the normals or displace the surface during final rendering.

5 Results

To demonstrate the modeling and rendering approaches described in the paper, we show results from several complex simulations.

5.1 Copper Strips

Figure 8 depicts the weathering of copper strips exposed to marine, rural, and urban environments respectively over the course of a thirty year period (in six year increments). This example shows the possibilities for the development of natural patinas in distinct environments. In all cases, the basic development of the patina was derived from the background material presented in Section 2. The patterns used in the layer maps were created though the fractal surface growth models described in Section 3.3; in the scripts below, the models are identified with the algorithm name, growth rate, step number, and total number of steps. Names such as `marine_patina_3` signify mixtures of the substances outlined in Figure 2 and are treated as materials with the various operators; the integer indicates the specific stage in this chart. The development of the patina was simulated using the layer structure and combinations of operators. The leftmost strips represent the appearance of pure copper — the starting point for the simulations.

Marine environment. The top row is a simulation of changes due to a typical marine environment. Here, we used the BD and DPD models to vary the thickness of the layers as a function of time. The BD model yields spotty patterns that are characteristic of marine environments, which often leave uneven coatings of moisture and salts on surfaces. The DPD model provides a patchy but more uniform coating of patina in the latter stages. In the last stage, we erode away a small portion of the top layer. The script for the last strip in the series is the following:

```
new copper;
coat tarnish_1 0.35 texture(BD.linear_1.20);
coat cuprite_2 1.2 texture(DPD.linear_5.40);
coat marine_patina_3 3.0 texture(BD.linear_10.20)
coat marine_patina_4 1.8 texture(DPD.linear_20.40);
erode 0.5 texture(BD.linear_5.20);
render maps;
```

Rural environment. The middle row is a simulation of effects common in a typical rural environment. Note that it takes longer for the green patina to begin to appear here. Patinas in rural environments tend to be quite spatially homogeneous. To capture this quality, we used the ST pattern generator to vary the thickness of the layers in the sequence. The script for the last strip in this series is the following:

```
new copper;
coat tarnish_1 0.25 texture(ST.cubic_1.20);
coat cuprite_2 1.0 texture(ST.cubic_5.20);
coat rural_patina_3 1.0 texture(ST.cubic_10.20);
coat rural_patina_4 1.1 texture(ST.cubic_15.20);
render maps;
```



Figure 8: Copper time lines: marine environment (top); rural environment (middle); urban environment (bottom).

Urban environment. The bottom row is representative of changes due to an urban environment. Here, the films are very thick. Note that in this case, the green copper salts begin to appear much earlier in the sequence. Half way into the sequence, the surface is covered with a thick, almost continuous, coating. In this sequence, we used the DPD model to generate the variations in thickness. This model gives rise to fairly continuous patches that grow laterally on the surface. It is common for urban patinas to be very soluble and to develop coatings of soot/dirt. To capture this quality, we erode away some of the patina and add a layer of dirt on the surface using the RD model on the last strip in the sequence. Here is the script for this strip:

```
new copper;
coat tarnish_1 0.55 texture(RD.linear_1.40);
coat cuprite_urban_2 1.0 texture(RD.cubic_10.40);
coat patina_urban_3 4.0 texture(DPD.cubic_5.40);
coat patina_urban_4 3.0 texture(DPD.cubic_10.40);
erode 0.8 texture(DPD.linear_30.40);
coat dirt 1.9 texture(BD.linear_5.20);
render maps;
```

Figure 9 illustrates the layers and the final BRDF parameters for the urban series. Note that the layer structure records the evolution of the strip from bottom to top. The initial development (bottom row) features output from the RD algorithm without surface relaxation. The remaining rows show patterns created with the DPD model. In the top row, the dirt pattern created with the BD model is visible. It is interesting to observe that each layer has constant surface properties; all the variations in color and texture arise from varying the thickness of the different layers.

5.2 Buddha

Figure 11 shows the development of the patina on a small statue of a buddha. The buddha model was created from a Cyberware scan and consists of approximately 60,000 small evenly sized triangles. In these pictures the table and wall are rendered using conventional texturing techniques, and the buddha is rendered using the techniques described in this paper.

In this experiment the various stages of the development of the patina are simulated in the RenderMan shading language. A shader

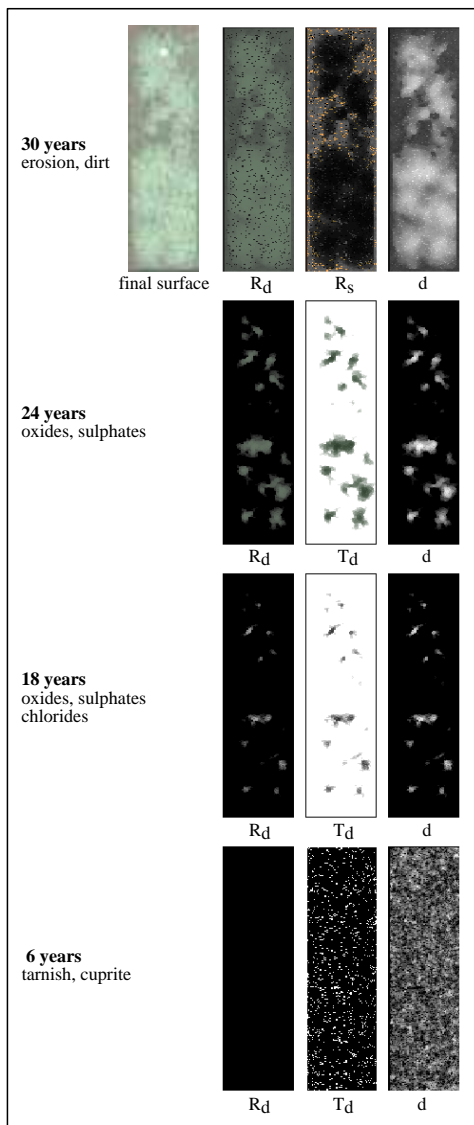


Figure 9: Layer properties and the final BRDF parameters for the urban series.

was written that modeled a three-layered surface: base copper, a tarnished layer, and a green patina. Parameters related to thickness are shown as a function of time in Figure 10. The thickness of different layers also depends on position and other factors as follows:

- The thickness of the tarnish layer was computed using two functions. The first, labeled tarnish in Figure 10, models tarnishing due to atmospheric processes. This parameter does not vary spatially. The second parameter, labeled polish, shows the decrease in tarnish thickness due to polishing the buddha. This parameter does vary spatially as determined by the accessibility map (the accessibility map is computed as a preprocess, and therefore there is a single accessibility value per triangle) [23]. Thus, tarnish appears in cracks and crevices, and shiny copper appears in exposed areas.
- The green patina consists of a steadily thickening layer and a set of steadily growing random patches. The steadily thickening layer depends on the local wetness of the surface; the wetter the surface the thicker the patina. The wetness is controlled by a precomputed exposure map, that gives the average irradiance due to the sun and sky received by that part of




Figure 10: Development of patina on the buddha statue.

the surface, and a simple function that decreases the average wetness of inclined surfaces. Another simple procedure increases the wetness along the base of the statue where water is likely to accumulate; this effect changes through time under control of the base parameter. The patches of the patina are controlled by a random fractal surface growth process. To show the flexibility of the system, the surface growth pattern is computed using the standard noise function in the shading language. The thickness of these various components of the patina change through time as shown in Figure 10. The curve labeled body controls the thickness of the steadily thickening layer, the curve labeled base controls the thickness along the bottom, and the curve labeled bump controls the thickness of the random growth process.

5.3 Towers

Figure 12 shows two copper-covered towers of different ages in an urban environment. The tower on the left displays a dark cuprite coating due to a few years of exposure; the tower on the right has a green patina typical of several decades of exposure.

For this simulation, we modulated the appearance of the patina based on several functions. As in the previous example, the wetness is controlled by a precomputed exposure map and a simple function that decreases the average wetness on inclined surfaces. For example, inclined areas that face south have thinner layers of patina, reflecting the fact that they would be dryer. We also used an accessibility map to vary the thickness of the patinas across the surfaces.

The new tower to the left is covered with a thin layer of cuprite. In addition, some simple texture maps were created with a noise function and combined with the layer structure to simulate the staining due to the flow of water over the top of the tower. Modeling such washing and staining effects is the subject of a separate work [5].

In the tower on the right, there is a thin layer of cuprite on the base copper. This is augmented with a fairly uniform layer of patina. We also added a thin layer of dirt, which is common to urban patinas. All layers are varied by simple noise functions.

6 Summary and Discussion

We have presented an approach for the modeling and rendering of one type of weathering — metallic patinas. A surface is represented as a set of layers. This representation resembles the underlying structure of the physical model of a patina. A set of operators was introduced that can simulate a wide variety of weathering effects. We also presented an approach to modeling the reflection and transmission through the layer structure using the Kubelka-Munk model. The modeling and rendering approach is capable of simulating a variety of metallic patinas.

Although we attempted to model the development of patinas on surfaces based on the available physical evidence, our model is still a phenomenological one. We believe that an exact model is not possible at this time, as the more general problem of the atmospheric corrosion of metals is not yet fully understood. However, the new layered model would appear to have great usefulness in computer graphics due to the ease with which it is possible to give the designer




Figure 11: *A sense of time*. On the left is a sequence of images showing the aging of a statuette. Time progresses from top to bottom. The larger image above illustrates the buildup of both the underlying smooth copper sulphide tarnish and the rough green patina.




Figure 12: Two copper-covered towers. (Left) After a few years of exposure. (Right) After several decades of exposure.

control over a wide variety of processes important in determining the final appearance of an object.

This representation of a material is significant, as it can accommodate the time-varying nature of a surface and its exposure to a complex class of weathering effects. We believe that this general model will be applicable to other materials as well. The problem of creating physically-based models of real materials, which can incorporate variations over time, is an important topic for computer graphics. This is a difficult problem, as the factors that determine the changes in appearance of materials operate simultaneously and are not completely understood. These difficulties notwithstanding, many interesting research directions remain. One area is in the design of models for other materials, such as wood or plastics. Another topic is the development of more complex operators, which can simulate dynamic processes, such as erosion. Last, we would like to add interactive operators to the system. These would be especially useful for creating patterns due to various brushes used in artificial patination techniques.

Acknowledgements

We would like to thank Jeff Feldgoise for modeling the tower scene and Matt Pharr for rendering it, Brian Curless for scanning the buddha model, Hans Pedersen for providing a geometry translator, Steve Westin for reviewing an early draft of the paper, Craig Kolb for video assistance, Tom Graedel for the copper micrograph, and the anonymous reviewers for helpful suggestions. This research was supported by research grants from the National Science Foundation (CCR-9207966 and CCR-9624172) and the MIT Cabot and NEC Funds, and by equipment grants from Apple and Silicon Graphics Inc.

References

- [1] BARABÁSI, A. L., AND STANLEY, H. E. *Fractal Concepts in Surface Growth*. Cambridge University Press, Cambridge, 1995.
- [2] BECKET, W., AND BADLER, N. I. Imperfection for realistic image synthesis. *Journal of Visualization and Computer Animation 1*, 1 (Aug. 1990), 26–32.
- [3] BLINN, J. F. Light reflection functions for simulation of clouds and dusty surfaces. *Computer Graphics 16*, 3 (July 1982), 21–29.
- [4] COOK, R. L. Shade trees. *Computer Graphics 18*, 3 (July 1984), 223–231.
- [5] DORSEY, J., PEDERSEN, H. K., AND HANRAHAN, P. Flow and changes in appearance. In *Computer Graphics Proceedings (1996)*, Annual Conference Series, ACM SIGGRAPH.
- [6] EBERT, D. S., Ed. *Texturing and Modeling*. Academic Press, New York, 1994.
- [7] FLEMING, S. J. *Dating in Archaeology*. St. Martin's Press, New York, 1977.
- [8] FRANEY, J. P., AND DAVIS, M. E. Metallographic studies of the copper patina formed in the atmosphere. *Corrosion Science 27*, 7 (1987), 659–688.
- [9] FRENCH, L. Toy story. *Cinefantastique 27*, 2 (1995), 36–37.
- [10] GRAEDEL, T. E. Copper patinas formed in the atmosphere – a qualitative assessment of mechanisms. *Corrosion Science 27*, 7 (1987), 721–740.
- [11] GRAEDEL, T. E., NASSAU, K., AND FRANEY, J. P. Copper patinas formed in the atmosphere. *Corrosion Science 27*, 7 (1987), 639–652.
- [12] HAASE, C. S., AND MEYER, G. W. Modeling pigmented materials for realistic image synthesis. *ACM Tran. Graphics 11*, 4 (Oct. 1992), 305–335.
- [13] HANRAHAN, P., AND KRUEGER, W. Reflection from layered surfaces due to subsurface scattering. In *Computer Graphics Proceedings (1993)*, Annual Conference Series, ACM SIGGRAPH, pp. 165–174.
- [14] HANRAHAN, P., AND LAWSON, J. A language for shading and lighting calculations. *Computer Graphics 24*, 4 (Aug. 1990), 289–298.
- [15] HSU, S., AND WONG, T. Simulating dust accumulation. *IEEE Computer Graphics and Applications 15*, 1 (Jan. 1995), 18–22.
- [16] HUGHES, R., AND ROWE, M. *The Colouring, Bronzing and Patination of Metals*. Watson-Guptill Publications, New York, 1991.
- [17] JUDD, D. B., AND WYSZECKI, G. *Color in Business, Science, and Industry*. John Wiley & Sons, New York, 1975.
- [18] KORTUM, G. *Reflectance Spectroscopy*. Springer-Verlag, New York, 1969.
- [19] KUBELKA, P. New contributions to the optics of intensely light-scattering material, part I. *J. Opt. Soc. Am. 38* (1948), 448.
- [20] KUBELKA, P. New contributions to the optics of intensely light-scattering material, part II: Non-homogeneous layers. *J. Opt. Soc. Am. 44* (1954), 330.
- [21] KUBELKA, P., AND MUNK, F. Ein Beitrag zur Optik der Farbanstriche. *Z. tech. Physik. 12* (1931), 593.
- [22] MATTSSON, E. *Basic Corrosion Technology for Scientists and Engineers*. Ellis Horwood Limited, New York, 1989.
- [23] MILLER, G. Efficient algorithms for local and global accessibility shading. In *Computer Graphics Proceedings (1994)*, Annual Conference Series, ACM SIGGRAPH, pp. 319–326.
- [24] MOSTAFAVI, M., AND LEATHERBARROW, D. *On Weathering: The Life of Buildings in Time*. MIT Press, Cambridge, MA, 1993.
- [25] NEWMAN, R. C., AND SIERADZKI, K. Metallic corrosion. *Science 263* (1994), 1708–1709.
- [26] PERLIN, K. An image synthesizer. *Computer Graphics 19*, 4 (July 1985), 287–296.
- [27] PORTER, T., AND DUFF, T. Compositing digital images. *Computer Graphics 18*, 3 (July 1984), 253–259.
- [28] SIMPSON, J. W., AND HORROBIN, P. J. *The Weathering and Performance of Building Materials*. MTP Publishing Co, London, 1970.
- [29] THOMAS, T. R., Ed. *Rough Surfaces*. Longman, New York, 1982.
- [30] TURK, G. Generating textures for arbitrary surfaces using reaction-diffusion. *Computer Graphics 25*, 4 (July 1991), 289–298.
- [31] UPSTILL, S. *The Renderman Companion*. Addison-Wesley, New York, 1990.
- [32] VERNON, W. H. J., AND WHITBY, L. The open air corrosion of copper, a chemical surface patina. *Journal Instit. of Metals 42*, 6 (1932), 181–195.
- [33] WITKIN, A., AND KASS, M. Reaction-diffusion textures. *Computer Graphics 25*, 4 (July 1991), 299–308.