



Light Transport



Surface Reflection

- When light hits a surface some is absorbed, the rest is reflected and transmitted (never mind for now)
- The reflected light is what we see
- Reflection is not simple and varies with material
 - the surface's micro structure defines the details of reflection
 - variations produce anything from bright specular reflection (mirrors) to dull matte finish (chalk)



Illumination

- Light Sources emit light
 - EM spectrum (color)
 - Position and direction
- Surfaces reflect light
 - Geometry (position, orientation, micro-structure)
 - Absorption
 - Transmission
 - Reflectance
- Illumination determined by the interactions between light sources and surfaces





Types of Light Sources

- Ambient: equal light in all directions
 - a hack to model inter-reflections



- Directional: light rays oriented in same direction
 - good for distance light sources (sunlight)



- Point: light rays diverge from a single point
 - approximation to a light bulb




More Light Sources



- **Spotlight:** point source with directional fall-off
 - intensity is maximal along some direction D , falls off away from D
 - specified by color, point, direction, fall-off parameters
- **Area Source:** Luminous 2D surface
 - radiates light from all points on its surface
 - generates soft shadows

Diffuse Reflection

- Simplest kind of reflector (also known as *Lambertian Reflection*)
- Models a matte surface -- rough at the microscopic level
- Ideal diffuse reflector
 - incoming light is scattered equally in all directions
 - viewed brightness does not depend on viewing direction
 - brightness *does* depend on direction of illumination




- Lambert's Law $I_{diffuse} = k_d I_{light} \cos \theta$
 $= k_d I_{light} (N \bullet L)$


I_{light} : Light Source Intensity

k_d : Surface reflectance coefficient in [0,1]

θ : Light/Normal angle



Examples of Diffuse Illumination



Above is the same sphere lit diffusely from different lighting angles


Ambient + Diffuse Reflection

- CG started using the Lambertian model and then added more terms as extra effects were required

$$I_{d+a} = k_a I_a + k_d I_{light} (N \cdot L)$$

I_a : Ambient light intensity (global)

k_a : Ambient reflectance (local)



- This is diffuse illumination plus a simple ambient light term
 - a trick to account for a background light level caused by multiple reflections from all objects in the scene

Further Simple Illumination Effects

- Light attenuation:

- light intensity falls off with the square of the distance from the source - so we add an extra term for this

$$I_{d+a} = k_a I_a + f_{att} k_d I_{light} (N \bullet L) \quad \text{where} \quad f_{att} = \frac{1}{d^2}$$

with d the light source to surface distance

- Colored lights and surfaces:


- just have three separate equations for RGB (or CIE, YIQ, etc.)

- Atmospheric attenuation:

- use viewer-to-surface distance to give extra effects
- the distance is used to blend the object's radiant color with a "far" color (e.g., a nice hazy gray)

Specular Reflection

- Shiny surfaces change appearance when viewpoint is varied
 - specularities (highlights) are view-dependent
 - caused by surfaces that are microscopically smooth
- For shiny surfaces part of the incident light reflects coherently
 - an incoming ray is reflected in a single direction (or narrow beam)
 - direction is defined by the incoming direction and the surface normal
- A mirror is a perfect specular reflector
 - approximate specular reflectors give fuzzy highlights



Phong Illumination

- One function that approximates specular falloff is called the *Phong Illumination* model


$$I_{\text{specular}} = k_s I_{\text{light}} (\cos \phi)^{n_{\text{shiny}}}$$

ϕ : Angle between reflected light ray **R** and viewer **V**



k_s : Specular reflectance

n_{shiny} : Rate of specular falloff


- no physical basis, yet widespread use in computer graphics




Computing the Reflected Ray



Project L onto N



Double length of vector



Subtract L

Phong Illumination Curves


- The specular exponents are often much larger than 1; values of 100 are not uncommon.

$$I_{specular} = k_s I_{light} (\cos \phi)^{n_{shiny}}$$


ϕ : angle between line of sight and perfect reflection

k_s : Specular reflectance


n_{shiny} : Rate of specular falloff



Phong Illumination



Moving the light source




Changing n_{shiny}

Putting It All Together

$$I = k_a I_a + f_{att} I_{light} \left[k_d \cos\theta + k_s (\cos\phi)^{n_{shiny}} \right]$$

- Combining ambient, diffuse, and specular illumination
- For multiple light sources
 - Repeat the diffuse and specular calculations for each light source
 - Add the components from all light sources
 - The ambient term contributes only once
- The different reflectance coefficients can differ.
 - Simple “metal” : k_a and k_d share material color, k_s is white
 - Simple plastic: k_s also includes material color
- Remember, when cosine is negative lighting term is zero!

Some Examples



Where do we illuminate ?

To this point we have discussed how to compute an illumination model at a *single* point on a surface. But, at which points on the surface is the illumination model applied? Where and how often it is applied has a noticeable effect on the result.

Illuminating can be a costly process involving the computation of and normalizing of vectors to multiple light sources and the viewer.

Flat Shading

The simplest shading method applies only one illumination calculation at one point for each primitive. Which one? Usually the centroid. For a convex facet the centroid is given as follows:

$$centroid = \frac{1}{vertices} \sum_{i=1}^{vertices} \bar{p}_i$$



This technique is called *constant or flat shading*. It is often used on polygonal primitives.

Facts:

- For point light sources the direction to the light source varies over the facet
- For specular reflections the direction to the eye varies over the facet
- Facet facts are clearly visible

Gouraud Shading


The Gouraud Shading method applies the illumination model on a subset of surface points and interpolates the intensity of the remaining points on the surface. In the case of a polygonal mesh the illumination model is usually applied at each vertex and the colors in the triangles interior are linearly interpolated from these vertex values during polygon rasterization.

Facet artifacts are still visible.

Gouraud Shading

- Gouraud shading interpolates illumination values at each vertex using *screen-space interpolation*.
- Screen-space interpolation is *wrong!*
- However, you usually will not notice because the transition in colors is very smooth.
- There are some cases where the errors in Gouraud shading become obvious.
 - When switching between different levels-of-detail representations
 - At "T" joints.

A "T" joint



Phong Shading

In Phong shading (not to be confused with Phong's illumination model), the surface normal is linearly interpolated across polygonal facets, and the Illumination model is applied at every point.

A Phong shader assumes the same input as a Gouraud shader, which means that it expects a normal for every vertex. The illumination model is applied at every point on the surface being rendered, where the normal at each point is the result of linearly interpolating the vertex normals defined at each vertex of the triangle.



Phong shading will usually result in a very smooth appearance, facet artifacts can only be seen along silhouettes.

Comparison



faceted shading

Gouraud shading

Phong

More About Shading

- Approximating Phong Shading using Gouraud Shading?
- Texture mapping gets you further
 - Assign radiance based on an image

Transmission with Refraction

- Refraction:
 - the bending of light due to its different velocities through different materials
- Refractive index:
 - light travels at speed c/n in a material of refractive index n
 - c is the speed of light in a vacuum
 - varies with wavelength hence rainbows and prisms

MATERIAL	INDEX OF REFRACTION
Air/Vacuum	1
Water	1.33
Glass	about 1.5
Diamond	2.4

Snell's Law

- Light bends when passing from a material of index n_1 to one of index n_2 *Snell's law* gives the angle of refraction:

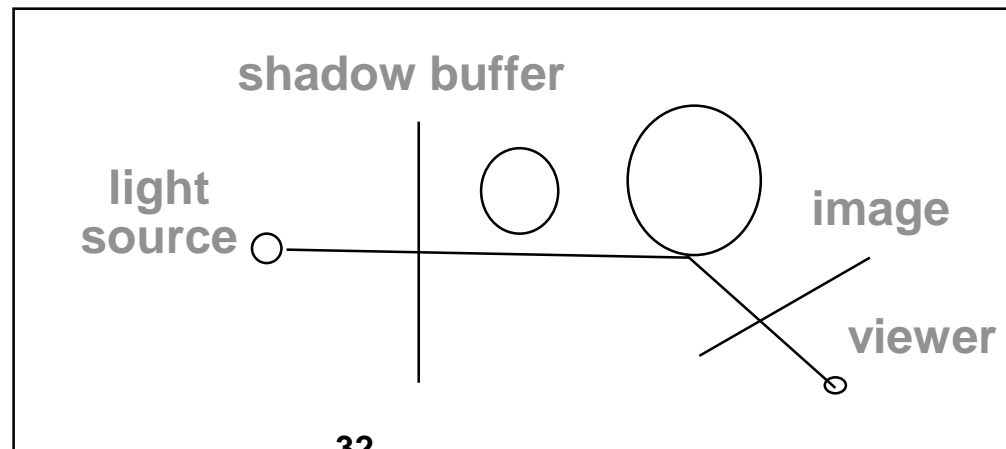
$$n_1 \sin \theta_1 = n_2 \sin \theta_2$$

where θ_1 and θ_2 are the angles from perpendicular

- When traveling into a denser material (larger n), light bends to be more perpendicular (e.g. air to water) and vice versa
- When traveling into a less dense material total internal reflection occurs if $\theta_1 > \sin^{-1}(n_2/n_1)$

Shadows

- Shadows occur where objects are hidden from a light source
 - omit any intensity contribution from hidden light sources
- Working out what is hidden is simply a visibility problem
 - can the light source see the object?
 - use the z-buffer shadow algorithm:
 - » run the algorithm from the light source' s viewpoint
 - » save the z-buffer as the shadow buffer
 - » run the real z-buffer algorithm, transforming each point into the light source' s coordinates and comparing the z value against the shadow buffer
 - » (Yeah, I know we haven' t done z-buffers yet)



How OpenGL Simulates Lights

- Phong lighting model
 - Computed at vertices
- Lighting contributors
 - Surface material properties
 - Light properties
 - Lighting model properties

Surface Normals

- Normals define how a surface reflects light


```
glNormal3f( x, y, z )
```

- Current normal is used to compute vertex' s color
- Use *unit* normals for proper lighting
 - » scaling affects a normal' s length

```
glEnable( GL_NORMALIZE )
```

or

```
glEnable( GL_RESCALE_NORMAL )
```



Material Properties

- Define the surface properties of a primitive

```
glMaterialfv( face, property, value );
```

GL_DIFFUSE	Base color
GL_SPECULAR	Highlight Color
GL_AMBIENT	Low-light Color
GL_EMISSION	Glow Color
GL_SHININESS	Surface Smoothness

- separate materials for front and back

Light Properties

```
glLightfv( light, property, value );
```

- *light* specifies which light

» multiple lights, starting with
GL_LIGHT0

```
glGetIntegerv( GL_MAX_LIGHTS, &n );
```

- *properties*

» colors

» position and type

» attenuation

An Example: OpenGL Materials

```
GLfloat white8[] = {.8, .8, .8, 1.},
           white2  = {.2, .2, .2, 1.},
           black   = {0., 0., 0.};
GLfloat mat_shininess[] = {50.}; /* Phong exponent */

glMaterialfv(GL_FRONT_AND_BACK, GL_AMBIENT, black);
glMaterialfv(GL_FRONT_AND_BACK, GL_DIFFUSE, white8);
glMaterialfv(GL_FRONT_AND_BACK, GL_SPECULAR, white2);
glMaterialfv(GL_FRONT_AND_BACK, GL_SHININESS, mat_shininess);
```


An Example: OpenGL Lighting

```
GLfloat white[] = {1., 1., 1., 1.};
GLfloat light0_position[] = {1., 1., 5., 0.}; /* directional light (w=0) */

glLightfv(GL_LIGHT0, GL_POSITION, light0_position);
glLightfv(GL_LIGHT0, GL_DIFFUSE, white);
glLightfv(GL_LIGHT0, GL_SPECULAR, white);
glEnable(GL_LIGHT0);


glEnable(GL_NORMALIZE); /* normalize normal vectors */
glLightModeli(GL_LIGHT_MODEL_TWO_SIDE, GL_TRUE); /* two-sided lighting*/

glEnable(GL_LIGHTING);
```

Color Spaces

RGB



1. convenient for display (CRT uses red, green, and blue phosphors)
2. not very intuitive



Color Spaces

Color Matching


R, G, B is mixed to reproduce the perceptual equivalent of any wavelength. A problem exists, however, because sometimes the red light needs to be added to the target before a match can be achieved. This is shown on the graph by having its intensity, R , take on a negative value.



Color Spaces

HSV

1. an intuitive color space
2. H is hue - what color is it? S is saturation or purity - how non-gray is it? V is value - how bright is it?
3. H is cyclic therefore it is a non-linear transformation of RGB





Color Spaces

CIE XYZ

- Often it is convenient to work in a 2D color space. This is commonly done by projecting the 3D color space onto the plane $X+Y+Z=1$, yielding a *CIE chromaticity diagram* (shown on the right). The projection is defined as:

$$x = \frac{X}{X+Y+Z} \quad y = \frac{Y}{X+Y+Z} \quad z = \frac{Z}{X+Y+Z}$$

- A linear transform of RGB used by color scientists




Additive vs. Subtractive Color

- Working with light: additive primaries
 - Red, green and blue components are added by the superposition property of electromagnetism
 - Conceptually: start with black, primaries add light
- Working with pigments: subtractive primaries
 - Typical inks (CMYK): cyan, magenta, yellow, black
 - Conceptually: start with white, pigments filter out light
 - The pigments remove parts of the spectrum

dye color	absorbs	reflects
cyan	red	blue and green
magenta	green	blue and red
yellow	blue	red and green
black	all	none


- Inks interact in nonlinear ways--makes converting from monitor color to printer color a challenging problem
- Black ink (K) used to ensure a high quality black can be printed

The Eye





- The image is formed on the *retina*
- Retina contains two types of cells: *rods* and *cones*
- Cones measure color (red, green, blue)
- Rods responsible for monochrome night-vision


The Fovea



1.35 mm from retina center



8 mm from retina center



Cones are most densely packed within a region of the retina called the *fovea*

- **Three types of cones: S,M,L**
 - Corresponds to 3 visual pigments
- **Roughly speaking:**
 - S responds to blue
 - M responds to green
 - L responds to red
- **Note that these are not uniform**
 - more sensitive to green than red
- **Colorblindness**
 - deficiency of one cone/pigment type