Sources, shadows and shading

- But how bright (or what colour) are objects?
- One more definition: Exitance of a source is
  - the internally generated power radiated per unit area on the radiating surface
- Similar to radiosity: a source can have both
  - radiosity, because it reflects
  - exitance, because it emits

- General idea:

Radiosity leaving =
Exitance + Radiosity due
to incoming light

- But what aspects of the incoming radiance will we model?

Radiance due to point sources

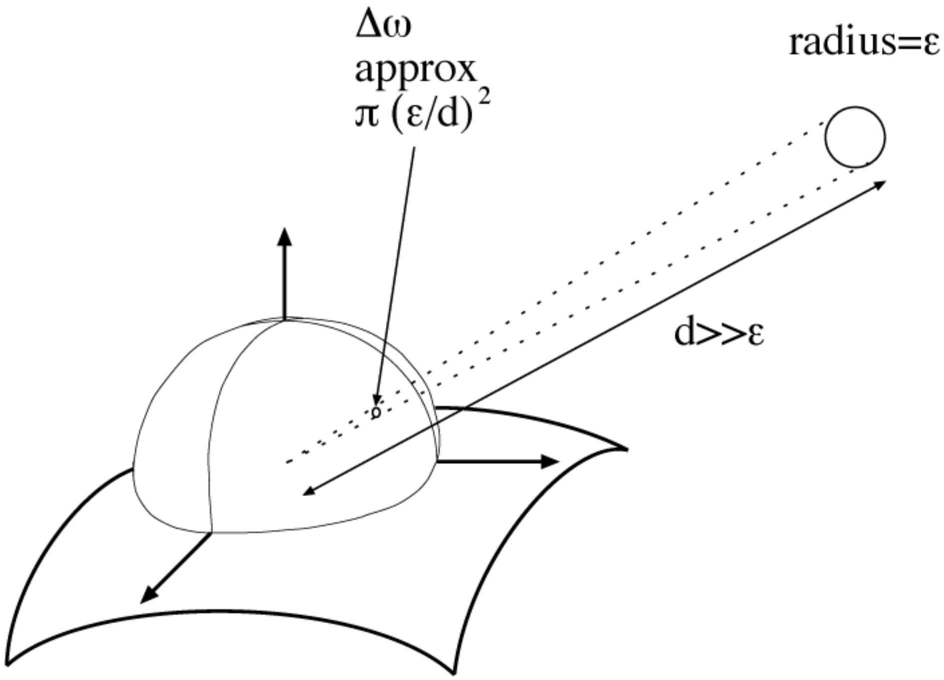

- small, distant sphere radius ε and exitance E, which is far away subtends solid angle of about

$$\pi\left(\frac{\varepsilon}{d}\right)^2$$

Standard nearby point source model

$$\rho_d(x)\left(\frac{N(x) \bullet S(x)}{r(x)^2}\right)$$

- N is the surface normal
- rho is diffuse albedo
- S is source vector - a vector from x to the source, whose length is the intensity term
  - works because a dot-product is basically a cosine
- r(x) is distance from surface point to source --- term occurs because source “looks smaller” as we move away--or, alternatively, its energy is spread out over a larger surface.

Standard distant point source model

- Issue: nearby point source gets bigger if one gets closer
  - the sun doesn't for any reasonable earthly notion of closer
- Assume that all points in the model are close to each other with respect to the distance to the source. Then the source vector doesn't vary much, and the distance doesn't vary much either, and we can roll the constants together to get:

$$\rho_d(x)\left(N(x) \bullet S_d(x)\right)$$

Shadows cast by a point source

- A point that can’t see the source is in shadow
- For point sources, the geometry is simple

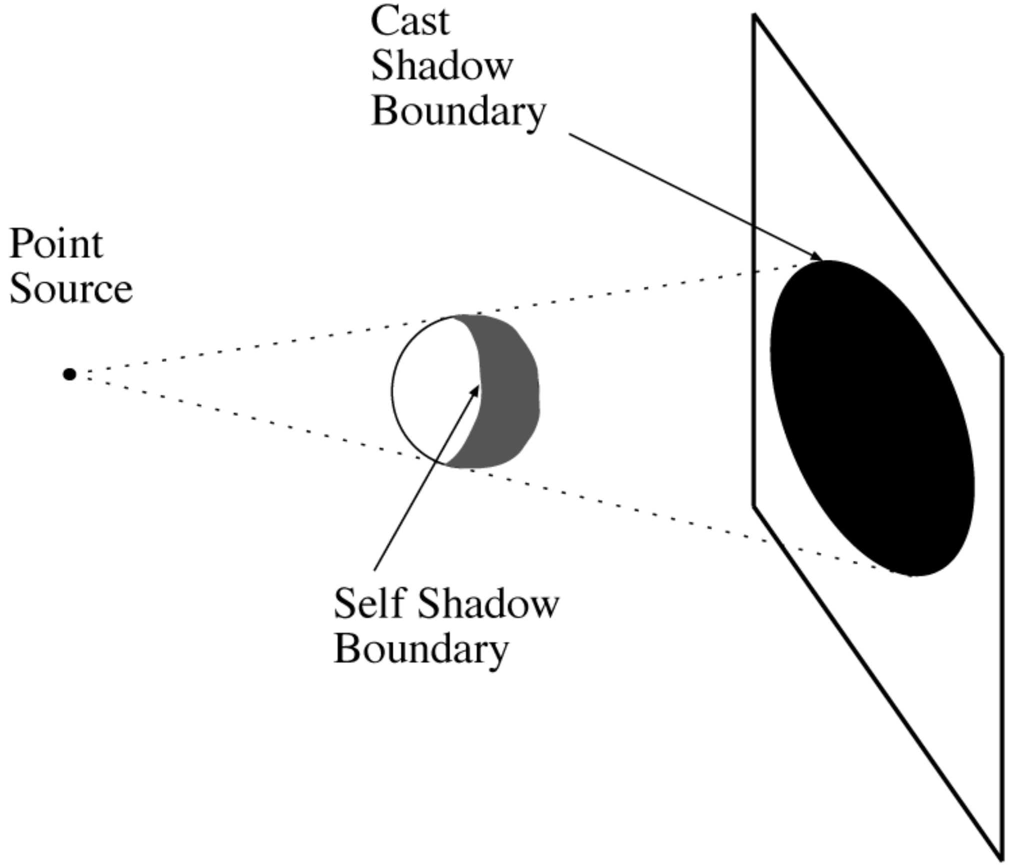

Line sources

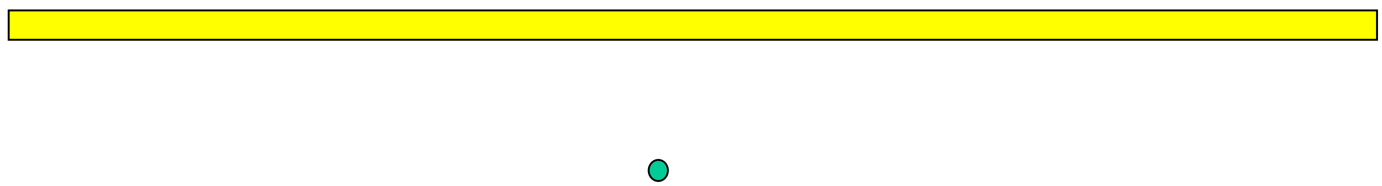

Radiosity due to line source varies with inverse distance,
if the source is long enough (derivation is through integration
of the contributions along the line)

General extended sources

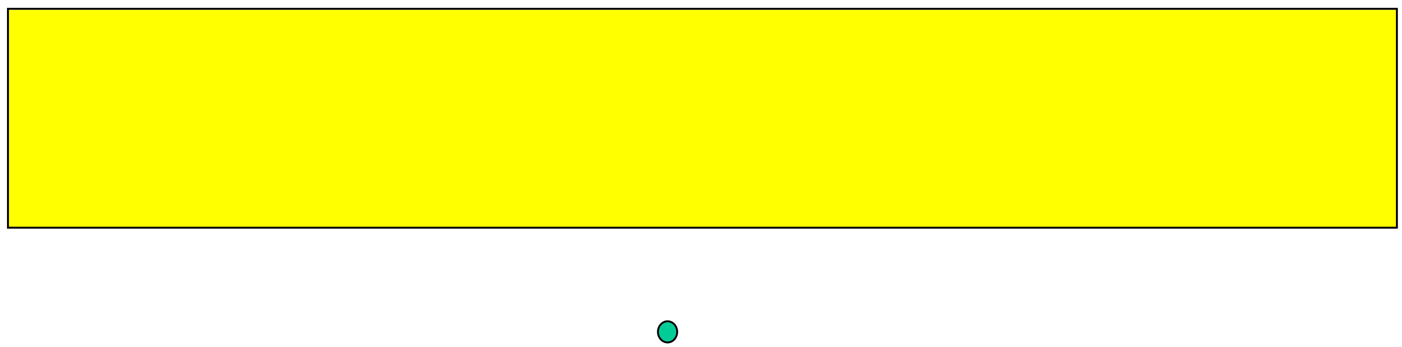

Can be handled by doing the integration (we won’t)

What if the source is large relative to the distance to it?

How about the hemisphere of the sky?

Local shading model

- Assume that all surface radiance is due to sources alone
  - I.e. both diffuse and specular, but no exitance.
- Can use standard point source model for diffuse term
  - either nearby or at infinity
- Common simplification:
  - drop 1/r^2 term from nearby point source (still have direction variation)
- Intensity = Diffuse intensity due to sources + specular term due to sources

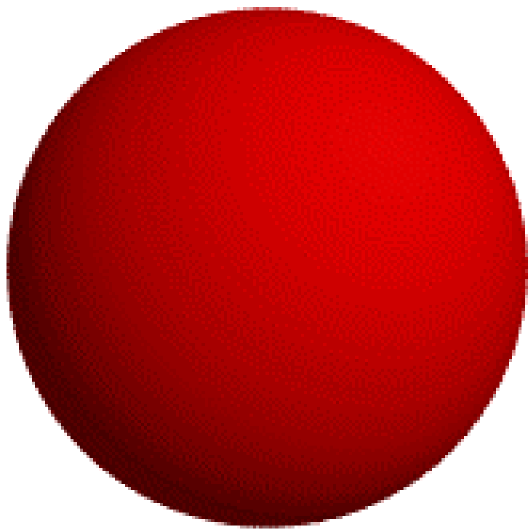

Diffuse Lighting

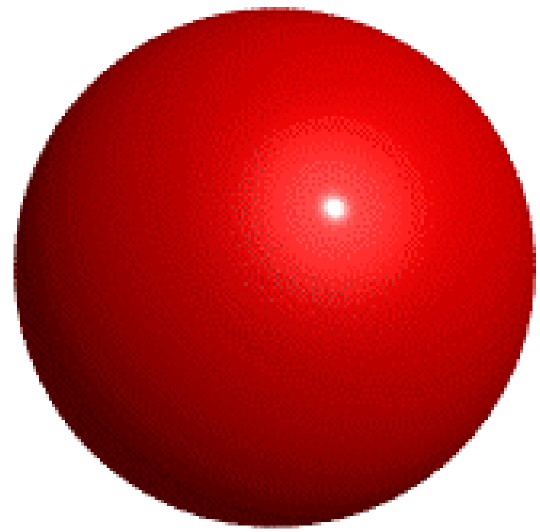

Plus Specular Highlight

from
http://www.geocities.com/SiliconValley/Horizon/6933/shading.html

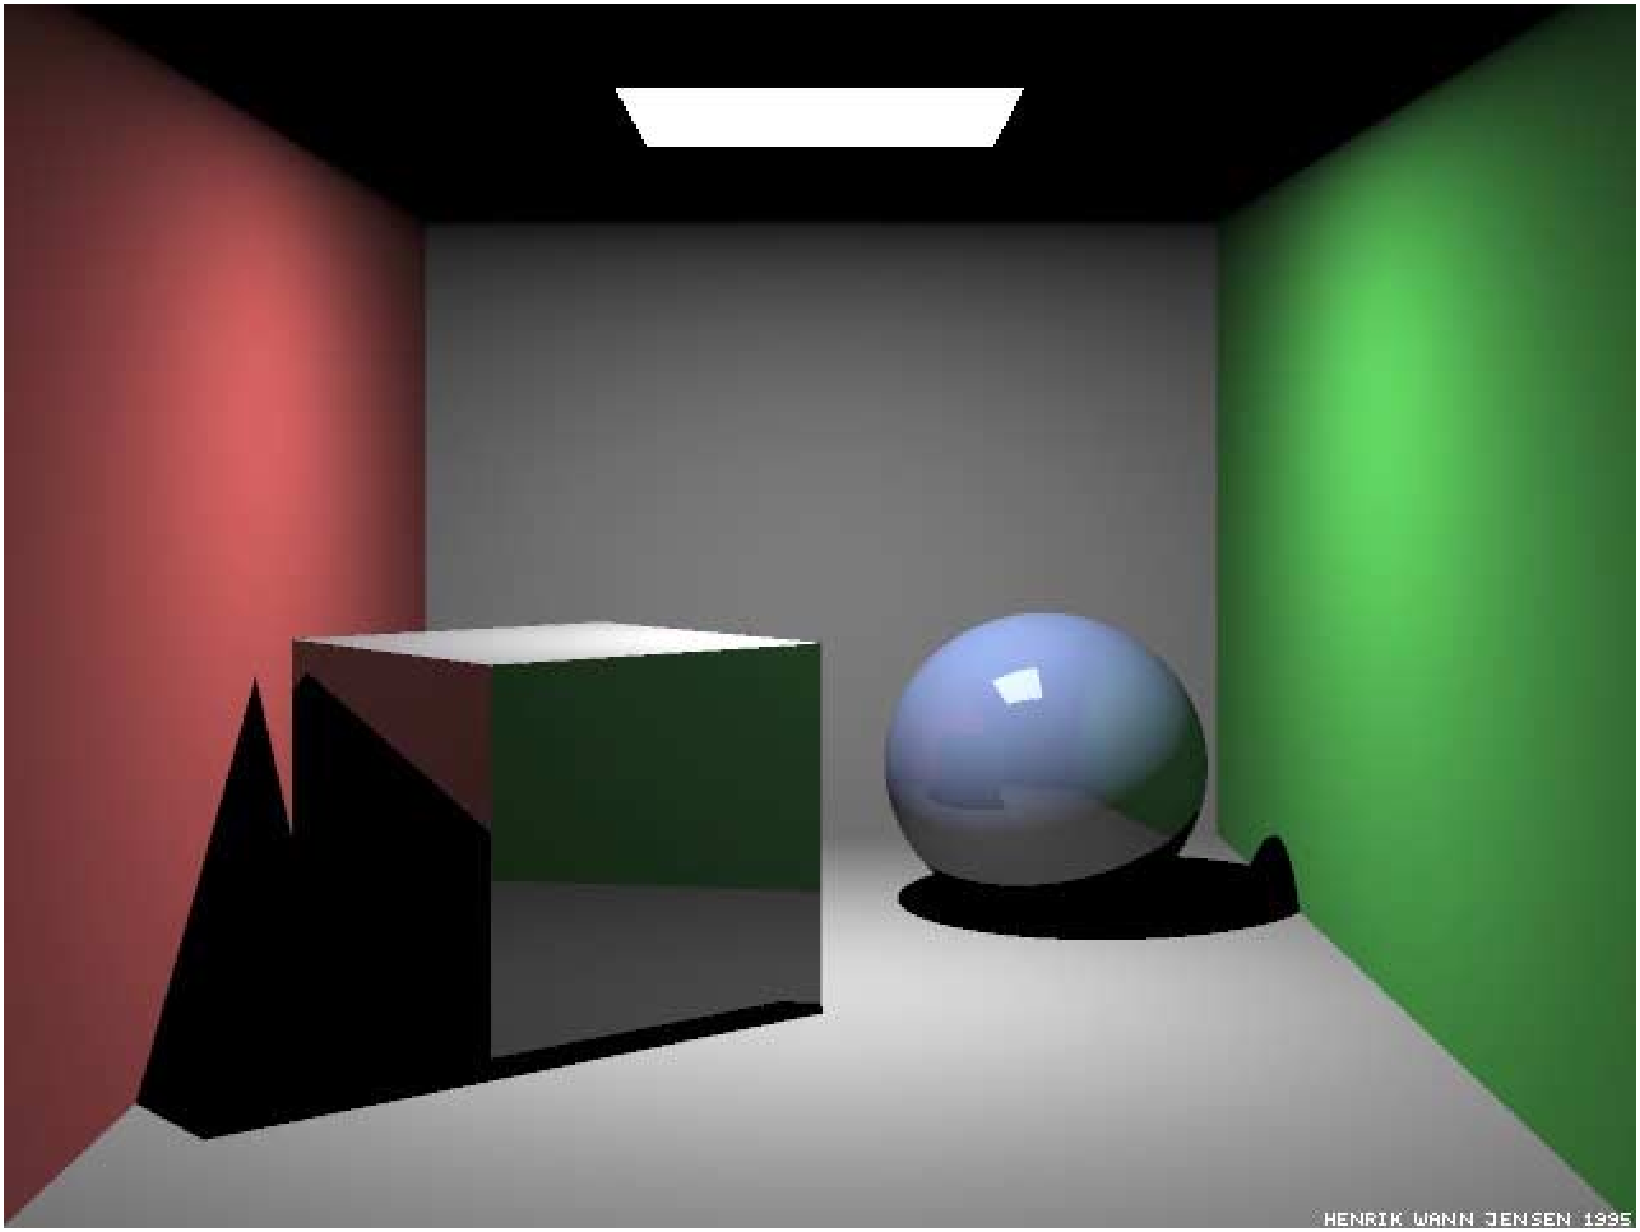

Ray-traced Cornell box, due to Henrik Jensen,
http://www.gk.dtu.dk/~hwj

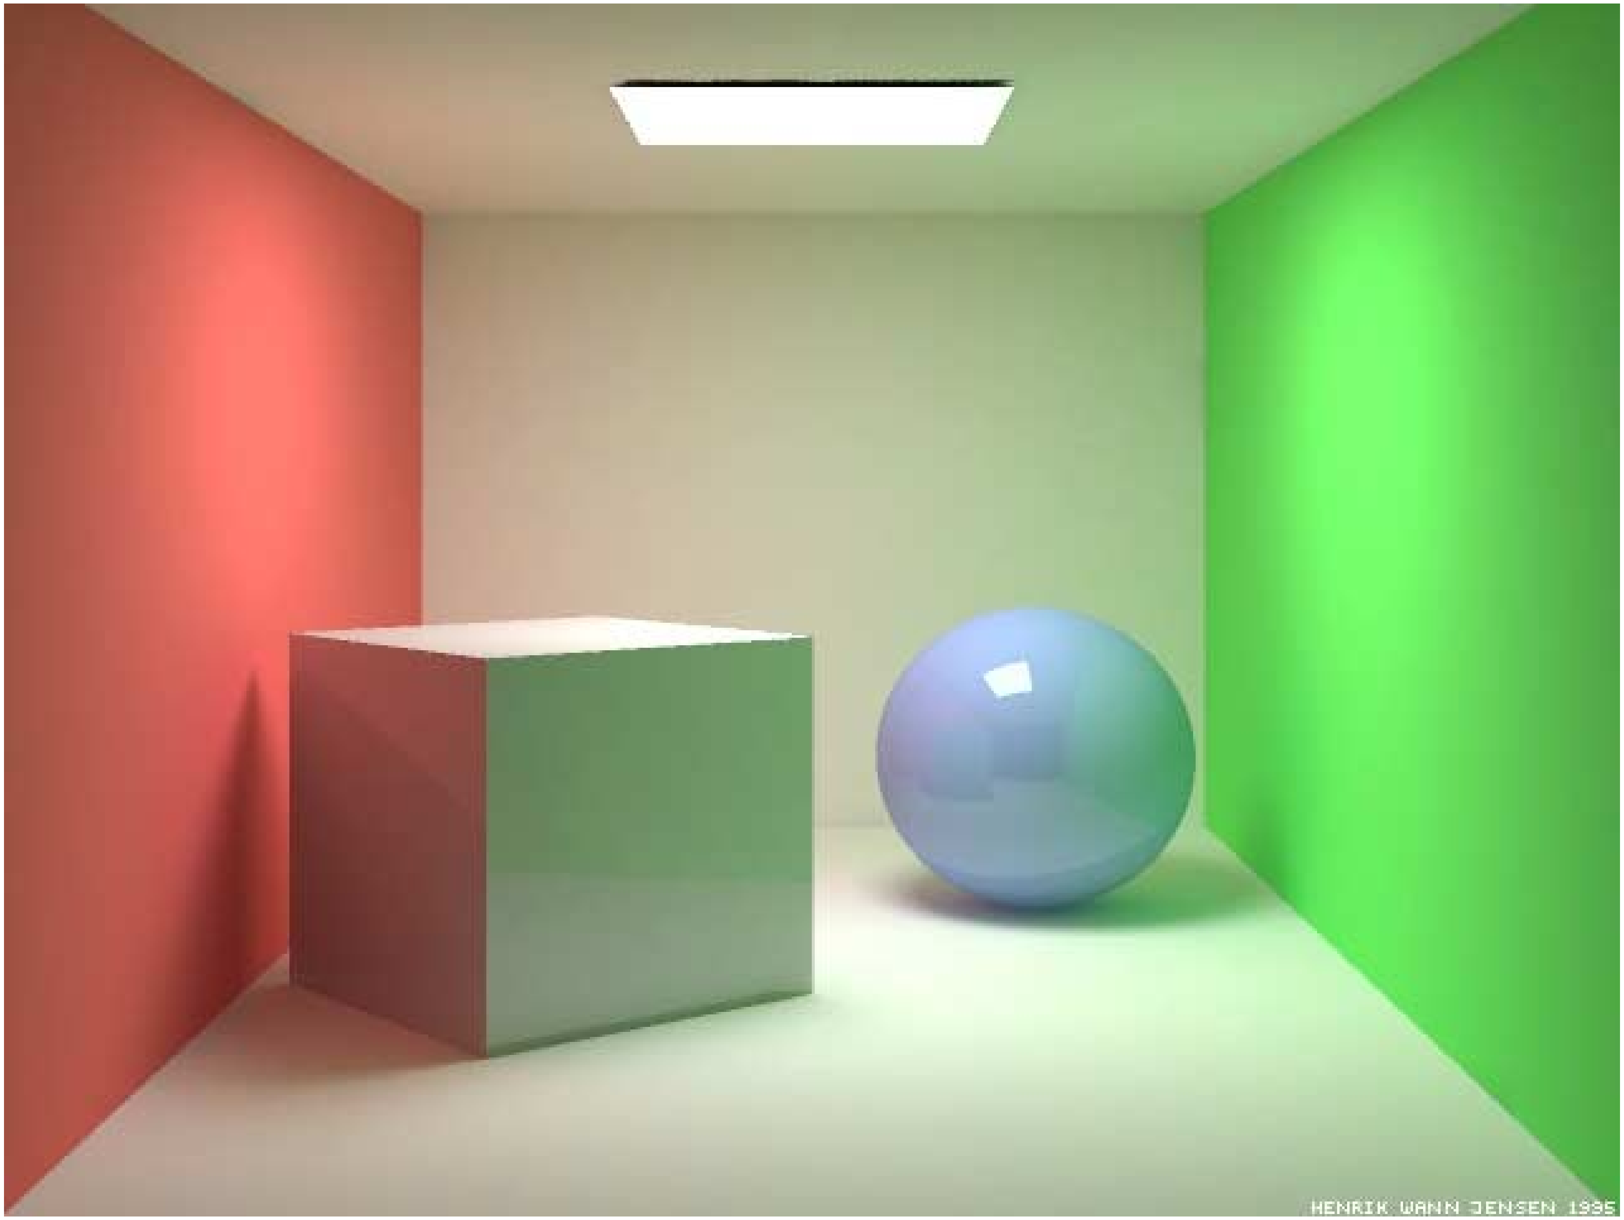

Radiosity Cornell box, due to Henrik Jensen, http://www.gk.dtu.dk/~hwj, rendered with ray tracer

Phong's model of specularities

- There are very few cases where the exact shape of the specular lobe matters.
- Typically:
  - very, very small --- mirror
  - small -- blurry mirror
  - bigger -- see only light sources as "specularities"
  - very big -- faint specularities
- Phong's model
  - reflected energy falls off with

$$\cos^n(\delta\vartheta)$$

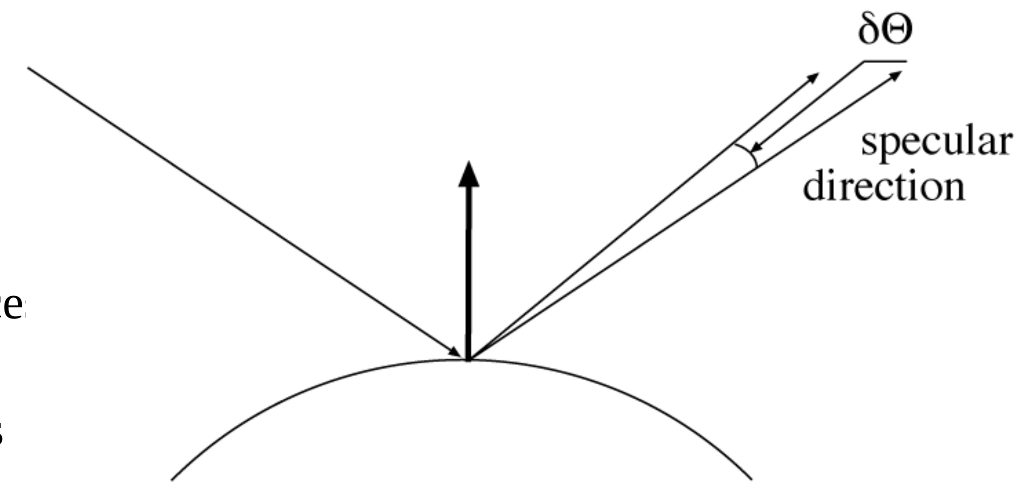

Flat shading

- compute shading value inside polygon using interpolate
- Flat shading
  - use polygon normal to shade
  - Adv:
    - fast -- one shading value per polygon
  - Disadv:
    - inaccurate -- looks blocky

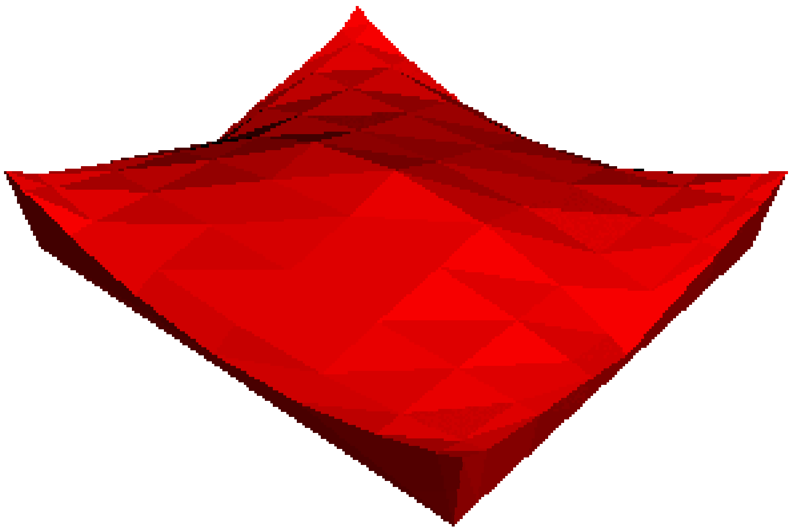

Figure from http://freespace.virgin.net/hugo.elias/graphics/x_polygo.htm

Interpolated Shading

- Gouraud shading
  - use normal at each vertex of polygon
  - shade these, and linearly interpolate
  - Adv:
    - fast
    - much smoother
  - Disadv:
    - specularities get lost

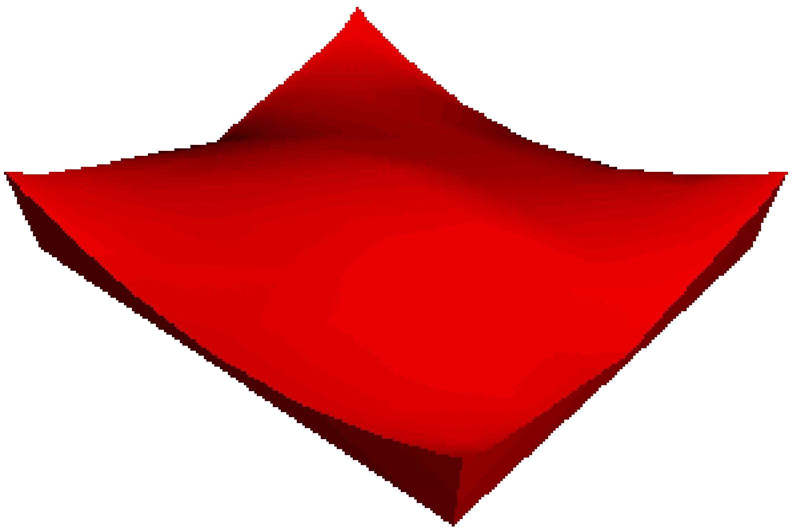

Phong Shading

- Interpolate normals
  - and then shade
  - Adv:
    - high quality, narrow specularities
  - Disadv:
    - expensive

Gouraud

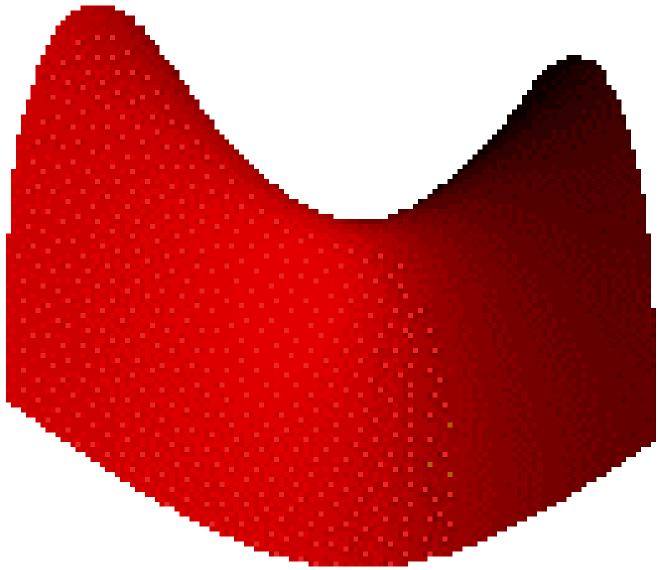

Phong

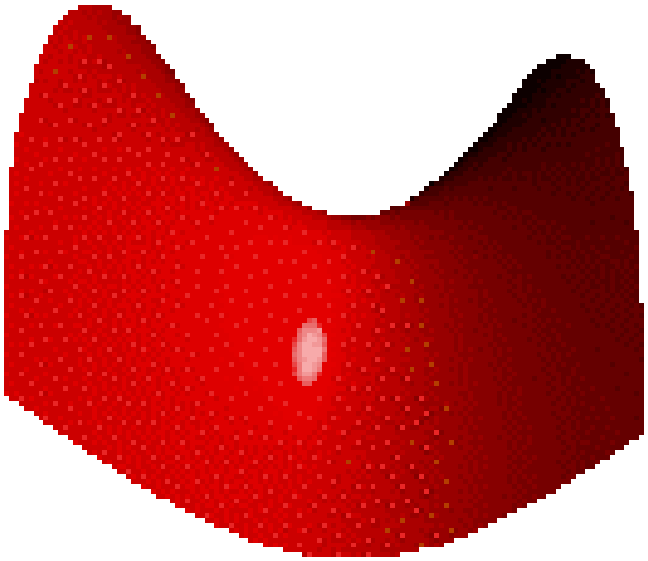

from
http://www.geocities.com/SiliconValley/Horizon/6933/shading.html

What about the color of the light?

So far, we have not dealt with the color of the light--the implicit assumption being that it does not change the color of anything.

This is clearly not the case!

Naïve (but common) model

Consider the color of the light to be specified by its (R,G,B)--technically the color of a perfect uniform reflector (white surface).

Similarly, now specify the albedo as a triple--one for each channel. The color of a Lambertian surface is then:

$$(R, G, B) = (\rho_R S_R,\ \ \rho_G S_G,\ \ \rho_B S_B)(\mathbf{n} \bullet \mathbf{s})$$

Naïve (but common) model

Naïve because we assume that the red part of the light does not interact with green or blue albedos, etc.

(Referred to as the diagonal model)

What about specular surfaces?

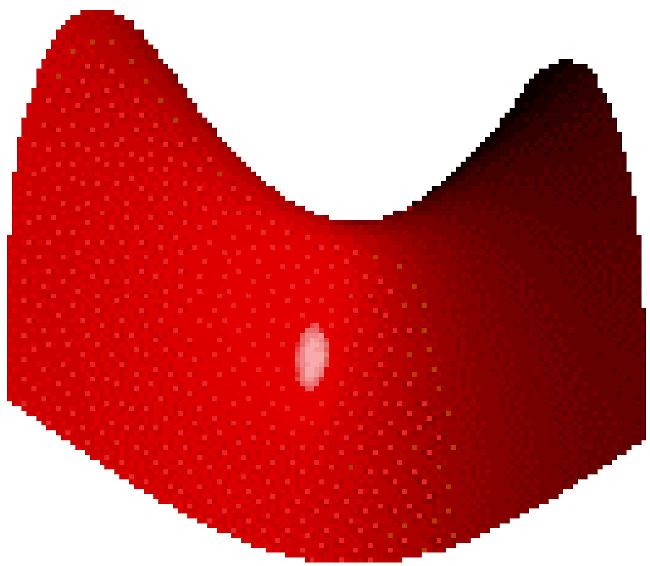

Specular surfaces

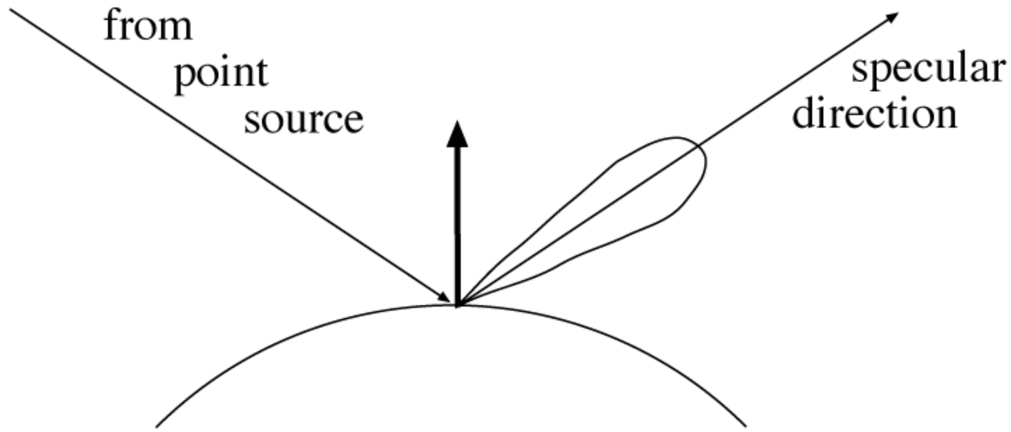

- Important point: The specular part of the reflected light usually carries the color of the **light**
- Technically, this is the case for dielectrics--plastics, paints, glass.
- Important exception is metals (e.g. gold, copper)

The colors of the rainbow

- Light is electromagnetic radiation, occurring at different wavelengths (or photon energies)
- The radiation around us is a mix of these
- Visible portion is about 400 to 700 nm
- Certain applications may require modeling some UV also.
- Light is specified by its spectrum recording how much power is at each wavelength.

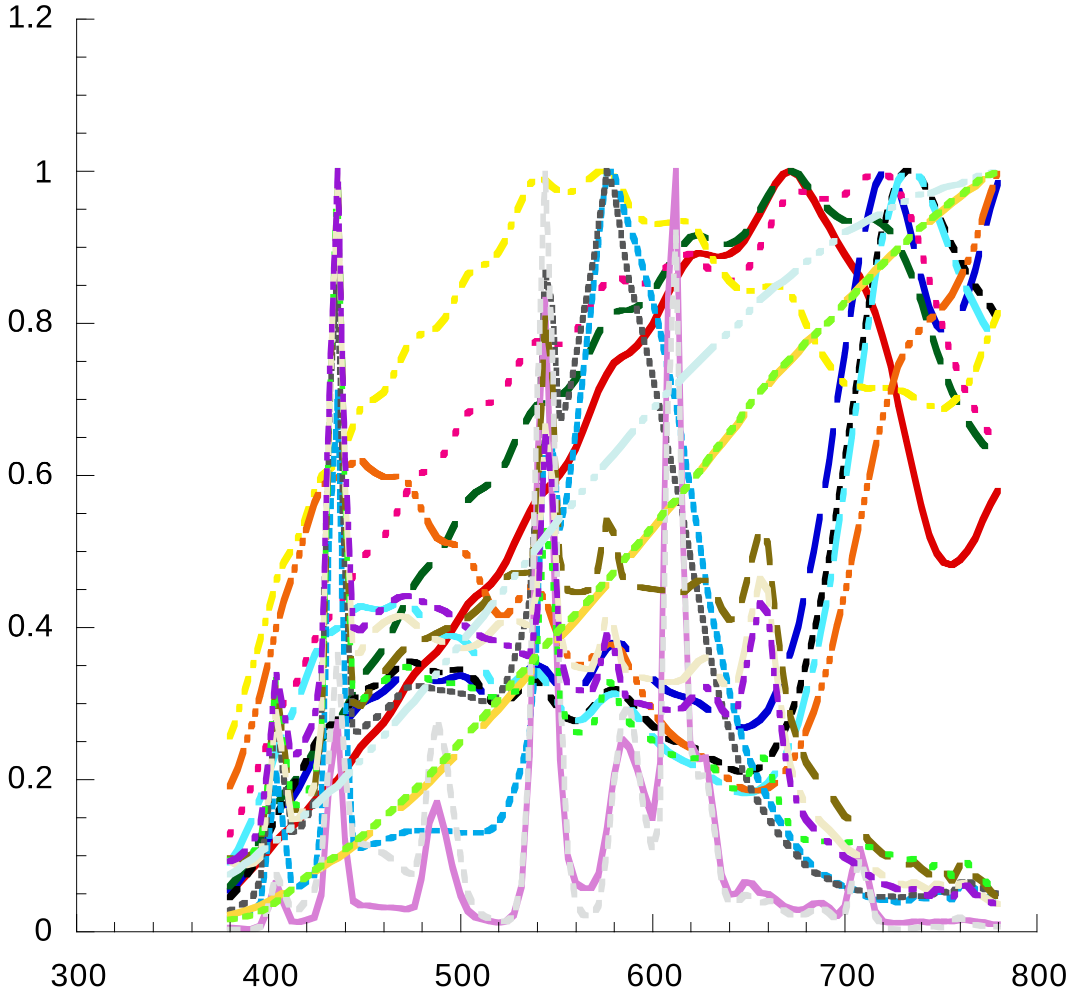
1.2
1
0.8
0.6
0.4
0.2
0
300
400
500
600
700
800

Radiometry for colour

- All definitions are now “per unit wavelength”
- All units are now “per unit wavelength”
- All terms are now “spectral”
- Radiance becomes spectral radiance
  - watts per square meter per steradian per unit wavelength
- Radiosity --- spectral radiosity

Case study

- Dielectric surface, well approximated by a specular part and a Lambertian body part.