Now Playing:

This Is a Fire Door Never Leave Open
The Weakerthans
from *Left and Leaving*
Released July 25, 2000

Movie:
HDR Far Cry

Available online at:
http://farcry.filefront.com/file/Far_Cry_HDR_video;54636

Color and
High Dynamic Range
(HDR) Imaging

Rick Skarbez, Instructor
COMP 575
November 20, 2007

Announcements

- Programming Assignment 4 (Ray tracer) is due today by 11:59pm
  - Any questions?
- You need to arrange to talk to me before December 1 for your project update

---

Path Tracing

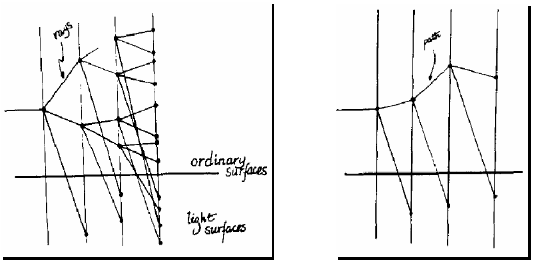

Trace one ray in one (random) direction,
and one ray to a light

Path Tracing Method

- At each ray intersection
  - Generate one ray based on diffuse / specular / transmissive coefficients
    - Not random; proportional to distribution
  - Also, generate one random ray per light
- Need a lot of rays per pixel
  - Kajiya used 40

Path Tracing Results

Ray Traced Image

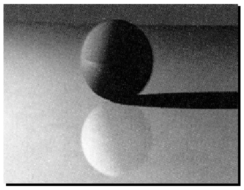

401 minutes

Path Traced Image

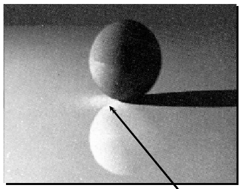

533 minutes

Note the light scattered off the diffuse sphere

Caustics

- Caustics occur when many light rays are reflected/refracted onto a single point
  - Anyone ever burn anything using a magnifying glass?
    - Caustics from Latin *causticus*, “burning”

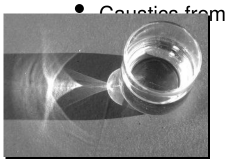

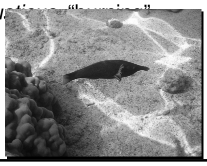

Caustics

- The feature that clearly distinguishes real global illumination solutions
  - Need to trace an envelope of rays from a diffuse surface through/off a curved reflective/refractive surface
    - Ray tracers can’t do it, because they can’t follow all the rays from a diffuse surface
    - Radiosity can’t do it, because it doesn’t have reflection/refraction

---

BDPT Results

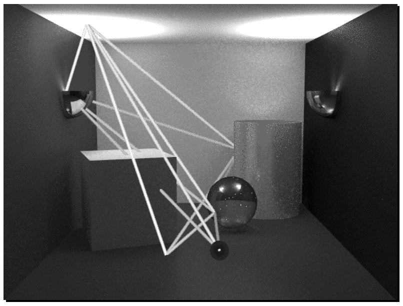

Bi-Directional Path Tracing Results

(a) Bidirectional path tracing with 25 samples per pixel

(b) Standard path tracing with 56 samples per pixel (the same computation time as (a))

Metropolis Light Transport

- Metropolis is a method for importance sampling paths
  - Instead of sampling paths randomly, identify a “good” path, and then sample paths that are slight perturbations from that path

Metropolis Results


(b) Metropolis light transport with 250 mutations per pixel [the same computation time as (a)].

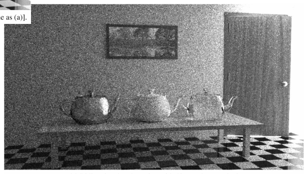
(a) Bidirectional path tracing with 40 samples per pixel.

Only light in this scene comes through the crack in the doorway

Metropolis Results

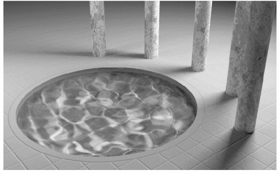
(b) Metropolis light transport with 100 mutations per pixel [the same computation time as (a)].

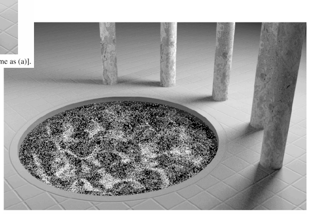
(a) Path tracing with 210 samples per pixel.

There are specific mutations to capture caustics.

Photon Mapping

Henrik Wann Jensen

- This is a two pass algorithm:
  - The <u>photon mapping</u> pass traces "photons" along rays from the light, and distributes them in the environment
    - The illumination data is stored in a <u>photon map</u>
  - The rendering pass traces rays from the eye, and reads back the illumination from the photon map to create the image

Photon Mapping Results

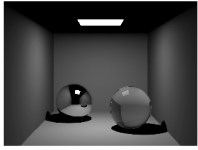
Figure 2.16: Ray traced Cornell box with sharp shadows.

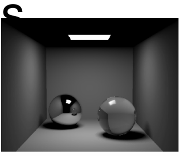
Figure 2.17: Ray traced Cornell box with soft shadows.

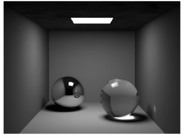
Figure 2.18: Cornell box with caustics.

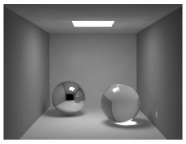
Figure 2.19: Cornell box with global illumination.

Photon Mapping Results

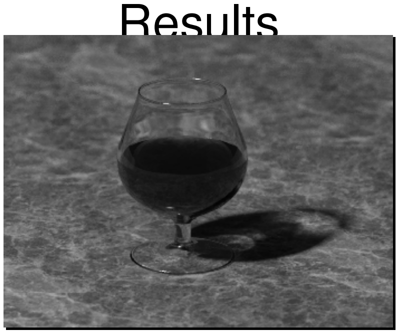

Global Illumination Review Over

- Any questions?

Understanding Color

- We've talked (briefly) about color corresponding to the wavelengths of light
- However, in everything we've actually done, we've represented color as a triple
  - One [0-255] value each for red, green, and blue
  - Are there any problems with this?
    - We'll come back to this later

Color Spaces

- We can think of RGB defining a <u>color space</u>
  - In this case, a 3-dimensional color space

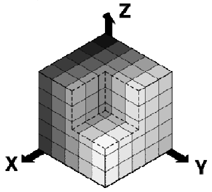

Problems with RGB

- Representing color with an RGB triple causes some problems
  - Perceptual difference
  - Luminance coupling
  - Color gamut

Perceptual Differences in RGB

- We saw that we can use RGB to define a 3D space
  - Ideally, distances in this space should *mean* something
    - For example, we could say how similar two colors are by Euclidean distance in color space
  - This does not happen with RGB

Perceptual Differences in RGB

- RGB was designed to correspond to cone sensitivities in the human eye
  - We have 3 types of cones, and each is sensitive to a different range of wavelengths
    - Roughly: red, green, and b
- However, this mapping do not give a color space with perceptually similar colors near to each other

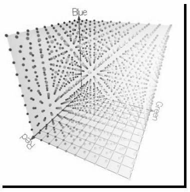

Human Visual Perception

- The human eye has two types of detectors
  - Rods and cones
- Rods, basically, only detect luminance and are the dominant detector in low light
- Cones detect color
  - Three types of cones: red, green, and blue (more or less)

Find me Find me

Human Visual Perception

- As you can see in the images below,the “red” cone doesn’t just see red, etc.
- The left image shows the range of wavelengths each cone responds to
- The right image shows their relative “brightnesses”

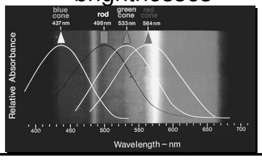

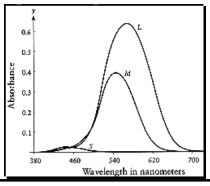

Human Visual Perception

- To humans (with normal vision), green appears much brighter than other colors

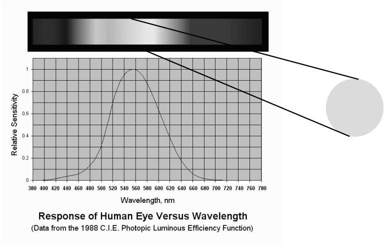

Response of Human Eye Versus Wavelength

(Data from the 1988 C.I.E. Photopic Luminous Efficiency Function)

Chrominance and Luminance

- When we talk about color, we are really combining two values
  - Chrominance
    - What color something is
  - Luminance
    - How bright that color is

Luminance Coupling

- RGB doesn’t provide any way to change the color without changing brightness
  - Why is this? (Remember how cones are sensitive)
- Example: If you reduce the blue value by 1 bit, and increase the green by 1 bit, the luminance doesn’t stay the same
  - Green is perceptually “brighter” than blue

Color Gamut

- The gamut is the rang of colors that can be described with a giver representation
- Problem: We can see colors that RGB cannot describe!

RGB gamut is white triangle

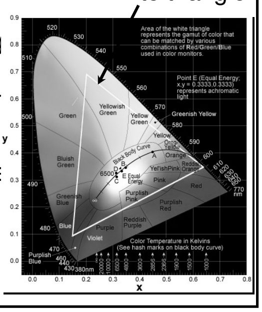

Other Color Spaces

- So, what else could we use to avoid some of these problems?
  - Luv (or Yuv)
  - CIE Lab
  - HSV
  - CMYK

Luv (or Yuv) Color Space

- L: a single value representing luminance
- u,v: two colors for chrominanc
  - Define a position in the color space to the right
  - u = Blue - L
  - v = Red - L
- Component (HD) video cables use YPbPr, which is a scaled version of Yuv

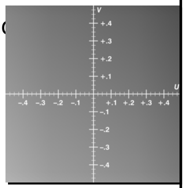

CIELab Color Space

- Very perceptually linear
  - Great for measurements and user studies
- Derived from the CIE 1931 calibration space, shown here

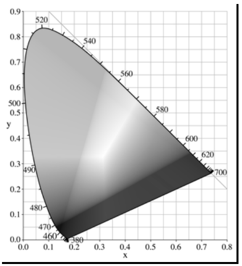

HSV Color Space

- Hue, saturation, and value
  - Hue: color
  - Saturation: “vividness”
  - Value: brightness
- Popular in graphics applications

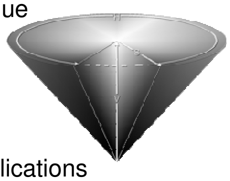

CMYK Color Space

- Not really used for computer graphics
  - However, far and away the most popular for printed materials
- CMYK is a subtractive (*i.e* like paints) color space
  - As opposed to RGB, which additive (*i.e.* like light)

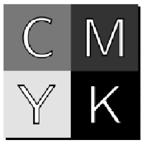

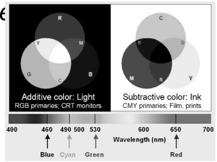

CMYK Color Space

- Why do we need subtractive col for printing?
  - Paper is already white (maximum value), so adding ink can only make the image darker
- Black (K) is separate
  - Because no combination of cyan, magenta, and yellow can generate a true black

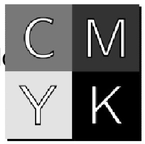

olor Spaces Revie

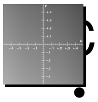

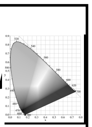

- These were just a sampling of possible color spaces
- There are equations to easily switch between spaces
  - However, some colors that are within the gamut of one space may not be in the gamut of another
- Consider what properties you need when choosing a color space

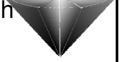

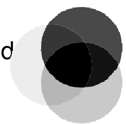

Another Problem

- So, we just got done discussing some of the problems inherent in the RGB color space
  - However, there is another problem with color that we haven't even talked about yet
    - Any guesses?
- Resolution!
  - We can only represent 255^3 = 16 million colors

Poor Quality Panoramas

A Better Panorama

Dynamic Range

- Computer monitors and digital cameras have limited dynamic range
  - 8 bits [0,255] to 12 bits [0,4095]
- The real world has MUCH greater dynamic range
  - The difference between sunlight and moonlight is on the order of 10000x
  - Some scenes can contain even a wider range

Capturing Greater Dynamic Range

- To capture this greater dynamic range with digital cameras, we can capture multiple bracketed images
  - Bracketing means taking multiple pictures of the same scene with different camera settings
    - *i.e.* different exposure times or aperture sizes
- To capture it with computer graphics, can just do lighting calculations with more bits

Stanford Chapel

Shortest Exposure

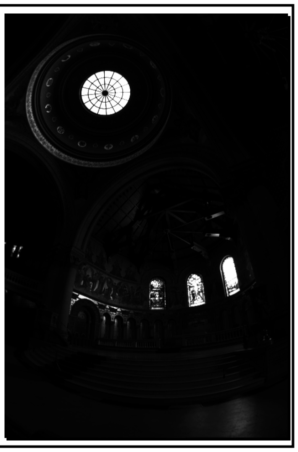

All HDR Data from http://debevec.org/

Stanford Chapel

Short Exposure

Stanford Chapel

Longer Exposure

Stanford Chapel

Longest Exposure

HDR Image Generation

- Given the set of bracketed images, and knowledge of the exposure times
  - We can reconstruct the most likely underlying HDR signal
- The “response curve” shown here is native to the camera

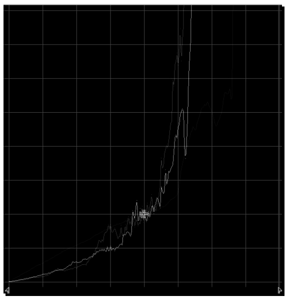

HDR Image Generation

- From here on out, the discussion assumes that we already have the underlying HDR image
  - Downloaded from the internet, or
  - Constructed from bracketed images by some other software package, or
  - Generated by a computer graphics application

HDR Image Generation

- We have a problem here
  - Does anyone see it?
    - These images have too much dynamic range to be drawn on our display!
- The process of fixing this is called <u>tone mapping</u>

What Can We Do?

Solution #1:
Log Mapping

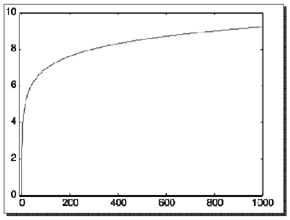

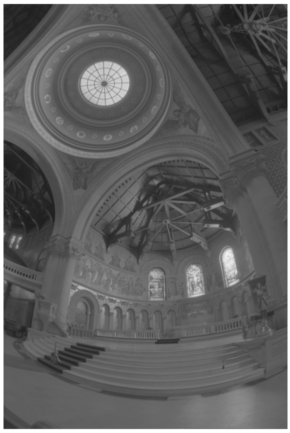

Log Mapping

- Log mapping is a two-step process
  - Take the log of the signal
  - Scale the new signal to use the entire 0-255 range
- Note that if you switch the order, you get a different result
  - Log is non-linear
- Note also that you cannot have any 0 values in your signal
  - log(0) = ∞

How can we do better?

- Retinex theory
  - Edwin Land, 1971
  - Basically, states that the human visual system is really bad at detecting absolute differences, and really good at detecting relative differences
    - Gradual changes in luminance aren't noticed
    - Sharp changes are

What Can We Do?

Solution #2:
Low Frequency Attenuation

- Reduce Low Frequencies
  - i.e. gradual changes
- Keep High Frequencies

What Can We Do?

Solution #2:
Low Frequency Attenuation

- Reduce Low Frequencies
  - i.e. gradual changes
- Keep High Frequencies

Low Frequency Mask

What Can We Do?

Solution #2:
Low Frequency Attenuation

- Reduce Low Frequencies
  - i.e. gradual changes
- Keep High Frequencies

High Frequency Mask

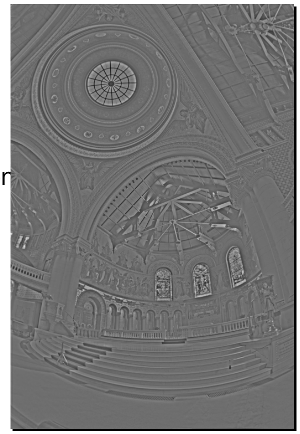

What Can We Do?

Solution #2:
Low Frequency Attenuation

- Reduce Low Frequencies
  - i.e. gradual changes
- Keep High Frequencies

Retinex HDR

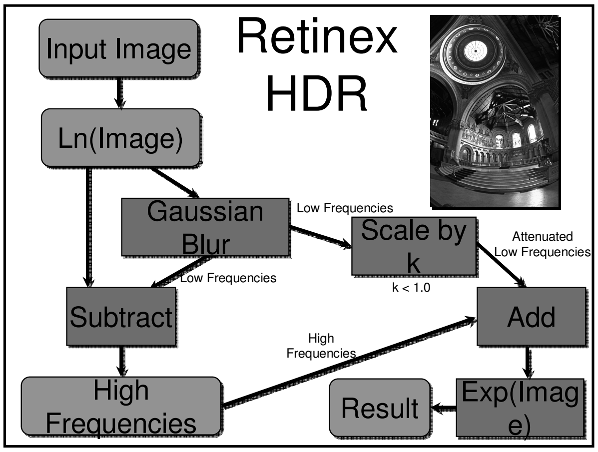

Input Image

Ln(Image)

Gaussian Blur

Low Frequencies

Scale by k

k < 1.0

Attenuated Low Frequencies

Low Frequencies

Subtract

Add

High Frequencies

High Frequencies

Exp(Image)

Result

How can we do even better?

- Maybe Gaussian filters aren't the best tool
  - Blur across edges, obscuring high frequency detail
- Can use an edge-preserving filter
  - I won't go into the math
  - Basically, the filter can recognize when it encounters an edge, and not blur across it

Bilateral Filter Example

Original Gaussian Bilateral

Bilateral HDR

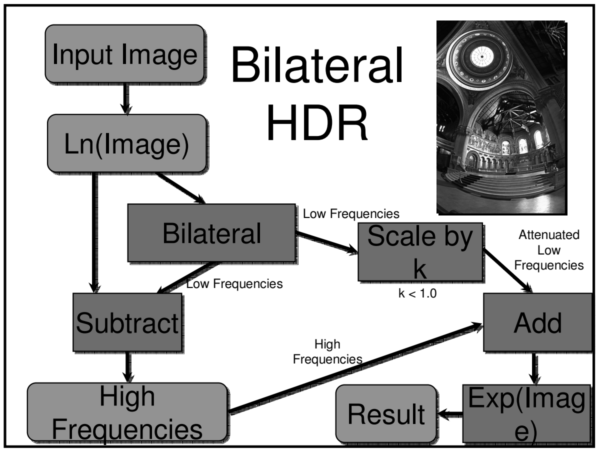

Input Image

Ln(Image)

Bilateral

Low Frequencies

Scale by k

k < 1.0

Attenuated Low Frequencies

Low Frequencies

Subtract

Add

High Frequencies

High Frequencies

Exp(Image)

Result

Bilateral Filter HDR

Durand & Dorsey, SIGGRAPH 02

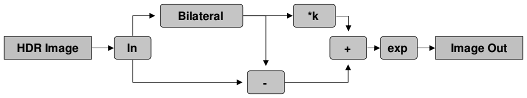

HDR Image

ln

Bilateral

*k

-

+

exp

Image Out

What Can We Do?

Solution #3:
Bilateral Tone Mapping

Durand & Dorsey 02

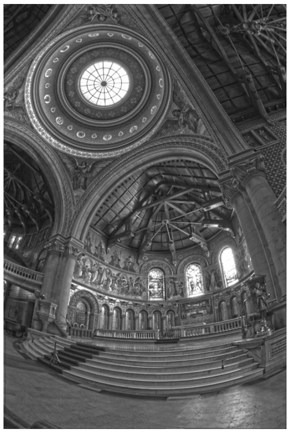

HDR Results

Global Compression
Local Tone Mapping
Local Tone Mapping

Log Mapping

Retinex
Tone Mapping

Bilateral
Tone Mapping

Sample HDR Images

Image from Wikipedia Commons

Sample HDR Images

Image from Wikipedia Commons

Sample HDR Images

Image from Wikipedia Commons