High Dynamic Range Imaging: Theory and Practice

SIGGRAPH 2006
Course #5
Sunday, 8:30 – 17:30

Erik Reinhard — University of Bristol
Paul Debevec — USC/ICT
Greg Ward — Anyhere Software
Karol Myszkowski — MPI for Informatik
Helge Seetzen — BrightSide Technologies
Drew Hess — Industrial Light + Magic
Gary McTaggart — Valve Software
Habib Zargarpour — Electronic Arts

Course Abstract

- Current display devices have a limited range of contrast and colors that can be displayed, which is one of the main reasons that most image and video acquisition, processing and display techniques use no more than eight bits per color channel. This course outlines recent advances in high dynamic range imaging (HDRI) - from capture to display - that lift this restriction thereby enabling images to represent the color gamut and dynamic range of the original scene rather than the limited subspace imposed by current monitor technology. In a hands-on approach, this course teaches how HDR images and video can be captured, the file formats available to store them, and the algorithms required to prepare them for display on low dynamic range display devices. The trade-offs at each step are assessed, allowing attendees to make informed choices about data capture techniques, file formats and tone reproduction operators. The course will cover the latest advances in Image-Based Lighting, in which HDR images can be used to illuminate CG objects and realistically integrate them into real-world scenes. In addition, we show the vast improvements in image fidelity afforded by HDRI through practical applications drawn from the film (ILM) and games industries (EA, Valve). Finally, recent advances in display hardware are demonstrated (BrightSide Technologies).

Prerequisites
Intended Audience

- Participants should be familiar with basic techniques in digital photography, traditional 8-bit image editing, and basic computer graphics modeling and rendering. Familiarity with a specific image-editing package or 3D modeling and rendering package would be helpful.

- This course is intended for students, researchers, and industrial developers in digital photography, computer graphics rendering, real-time photoreal graphics, game design and visual effects production (esp. rendering and compositing).

Summary

- This course teaches new techniques in capturing, representing, processing, and displaying High Dynamic Range Images and Video that cover the full range of light in the real world, and thereby enable marked improvements in visual fidelity and photorealism. Applications in lighting, compositing, game design and film, as well as advances in display hardware are covered.

Introduction

Images

- Traditionally we use one byte per pixel per color channel
- What if we had floating point numbers to represent colored pixels?
- This may seem a small issue, but really isn't…

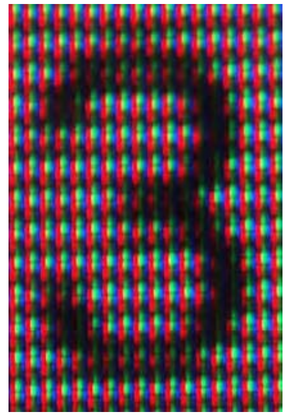

HDR Image Capture

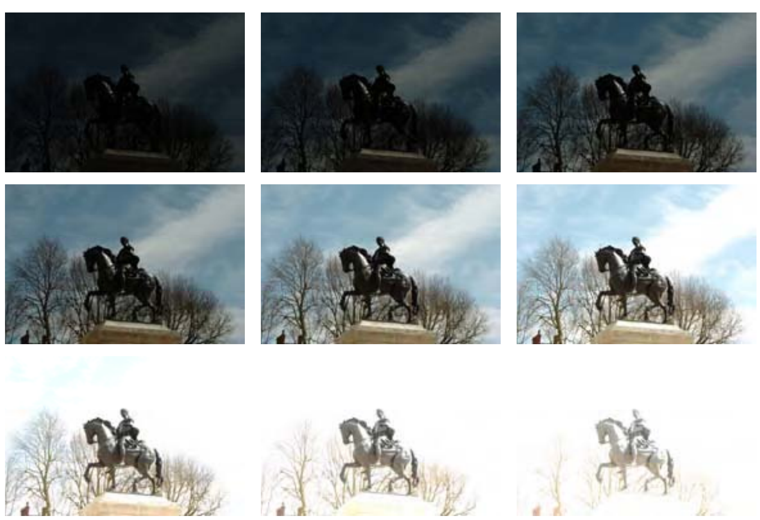

HDR Image Capture

HDR File Formats

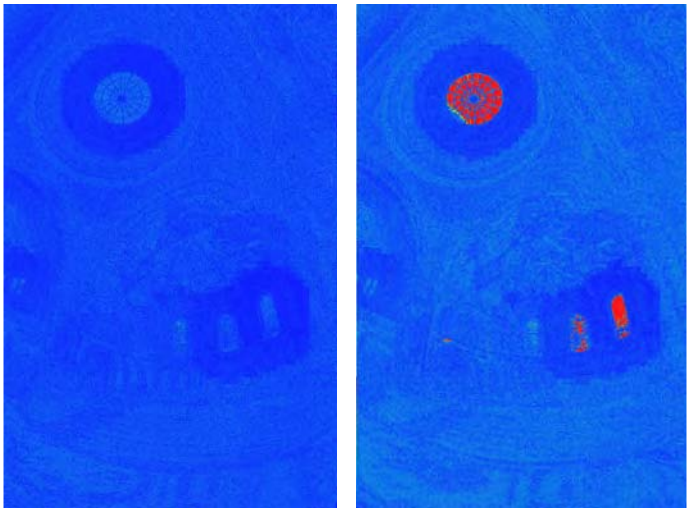

HDR Video

SIGGRAPH2006

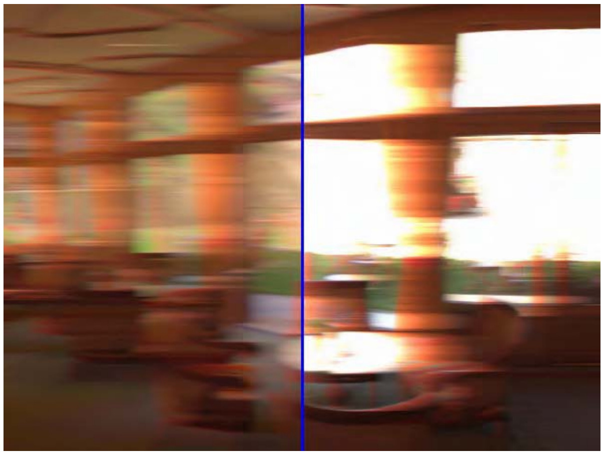

Tone Reproduction

SIGGRAPH2006

Image-Based Lighting

SIGGRAPH2006

HDR Display Devices

SIGGRAPH2006

Industrial Light and Magic

SIGGRAPH2006

Valve Software

SIGGRAPH2006

Electronic Arts

SIGGRAPH2006

Course Schedule (Morning)

SIGGRAPH2006

– Introduction – Erik Reinhard
(nearly finished)

– Taking HDR Images – Paul Debevec
8:40 – 9:15

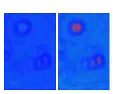

– HDR Image Representation – Greg Ward
9:15 – 10:00

– Break
10:00 – 10:15

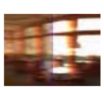

– HDR Video and Applications – Karol Myszkowski
10:15 – 10:55

Course Schedule (Morning)

SIGGRAPH2006

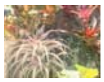

– Tone Reproduction Operators – Erik Reinhard

10:55 – 11:30

– HDR Image-Based Lighting – Paul Debevec

11:30 – 12:15

– Lunch Break

12:05 – 12:15

Course Schedule (Afternoon)

SIGGRAPH2006

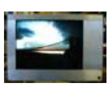

– HDR Display Devices – Helge Seetzen

1:45 – 2:45

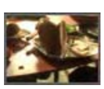

– HDR at ILM – Drew Hess

2:45 – 3:30

– Break

3:30 – 3:45

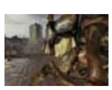

– HDR at Valve – Gary McTaggart

3:45 – 4:30

Course Schedule (Afternoon)

SIGGRAPH2006

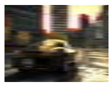

– HDR at EA – Habib Zargarpour
4:30 – 5:15

– Discussion - All
5:15 – 5:30

Course Notes

SIGGRAPH2006

Slides (clickable links):

Taking HDR Images – Paul Debevec

HDR Image Representation – Greg Ward

HDR Video and Applications – Karol Myszkowski

Tone Reproduction – Erik Reinhard

HDR Image-Based Lighting – Paul Debevec

Course Notes

SIGGRAPH2006

Slides:

- HDR Display Devices – Helge Seetzen
- HDR at Industrial Light + Magic – Drew Hess
- HDR at Valve – Gary McTaggart
- HDR at Electronic Arts – Habib Zargarpour

- Contrast Processing – Rafal Mantiuk
- Lightness Perception – Grzegorz Krawczyk

Course Notes

SIGGRAPH2006

Papers

Greg Ward and Maryann Simmons, 'Subband Encoding of High Dynamic Range Imagery', ACM Symposium on Applied Perception in Graphics and Visualization, pp 83-90, 2004

Greg Ward and Maryann Simmons, 'JPEG-HDR: A Backwards Compatible, High Dynamic Range Extension to JPEG', Proceedings of the 13th color imaging conference, 2005

Paul Debevec, 'Image-Based Lighting', IEEE Computer Graphics and Applications, 2002

Course Notes

SIGGRAPH2006

White Papers:

- Technical Introduction to OpenEXR
- OpenEXR File Layout
- Reading and Writing OpenEXR Image Files with the IlmImf Library
- High Dynamic Range Rendering in Valve's Source Engine

Links to External Resources

SIGGRAPH2006

- HDR mailing list (Greg Ward)

  http://www.radiance-online.org/mailman/listinfo/hdri
- Web site for book on High Dynamic Range Imaging

  http://www.hdrbook.com
- Light Probe Gallery (Paul Debevec)

  http://www.debevec.org/Probes
- Annotated List of HDR File Formats (Greg Ward)

  http://www.anyhere.com/gward/hdrenc/hdr_encodings.html
- OpenEXR website

  http://www.openexr.org

Links to External Resources

SIGGRAPH2006

- HDRShop (Paul Debevec)

  http://www.hdrshop.com
- Photosphere and related Mac and Unix software (Greg Ward)

  http://www.anyhere.com

Acknowledgments

SIGGRAPH2006

- Thanks to Rafal Mantiuk for providing slides on his tone reproduction operator
- Thanks to Grzegorz Krawczyk for providing extra slides on his model of lightness perception for tone reproduction
- Thanks to Yuanzhen Li for applying her tone reproduction algorithm to one of our images to enable a comparison

HDRI at SIGGRAPH 2006

- Papers Session: HDR and Systems
  - Tuesday 10:30 – 12:15
- BoF: High Dynamic Range Imagery
  - Boston Convention and Exhibition Center
  - Tuesday 17:00 – 18:00, Room 251

Taking HDR Images

Paul Debevec

Taking HDR Images

Paul Debevec
University of Southern California
Centers for Creative Technologies

SIGGRAPH 2006 Course #5
High Dynamic Range Imaging: Theory and Applications
Sunday, July 30th, 2006

www.debevec.org

Dynamic Range in the Real World

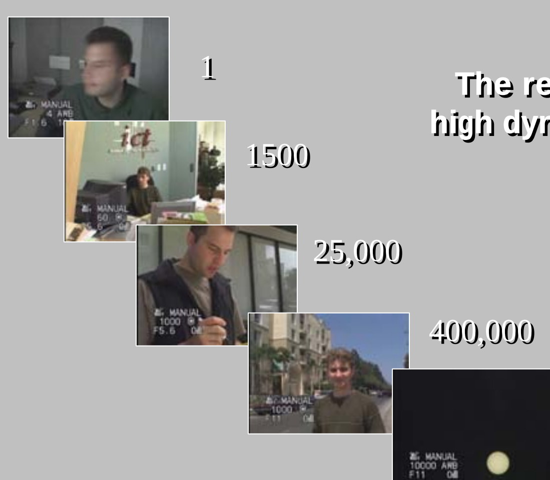

1

1500

25,000

400,000

2,000,000,000

The real world is high dynamic range.

High-Dynamic Range Photography

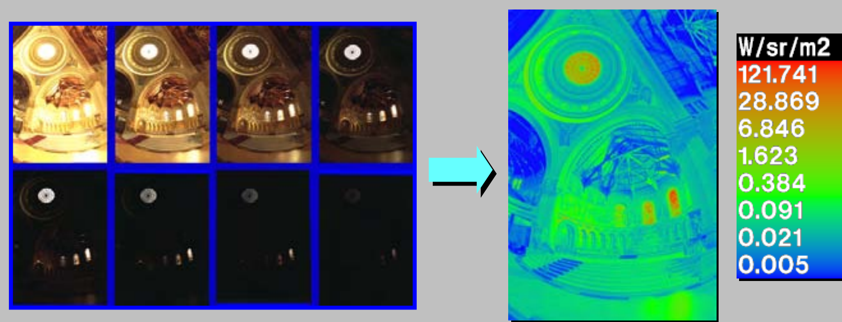

300,000 : 1

Visualization: Greg Ward

Allows one or more images of a scene to be converted to a high dynamic range image that is proportional to radiometric units.

Instead of integer values 0-255, pixel values can range from 0 into the millions, with floating-point precision.

Gray Scale Calibration Charts

Twenty 0.1 density increments =
26% more reflective each step = 1/3 stop
Be careful to minimize specular reflections

<= Gretag MacBeth ColorChecker Chart bottom row reflectivity:
88%, 58%, 35%, 20%, 8%, 3%

Gamma Calibration Chart

After Greg Ward

www.debevec.org/gamma.png

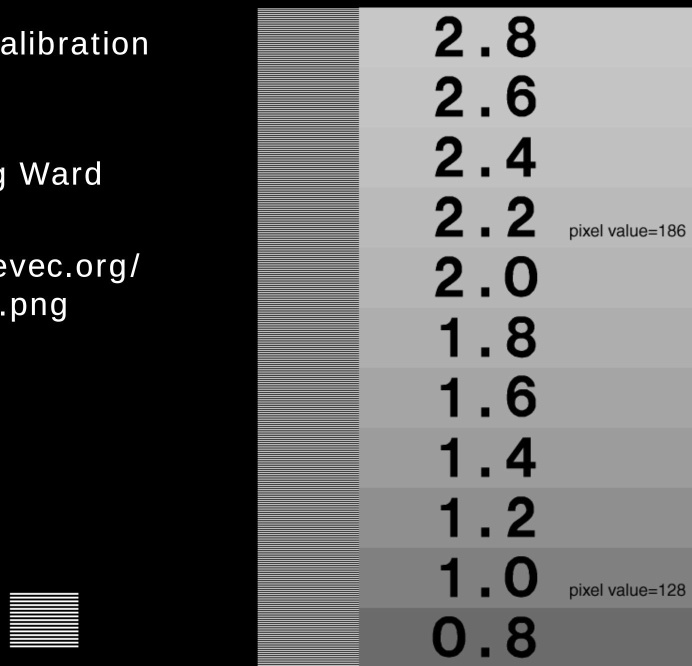

Gamma 2.2 graph

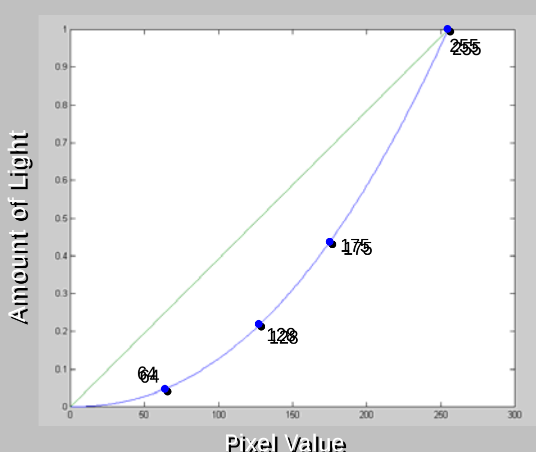

Implications:

128 is less than ¼ as bright as 255

128 is more than 4 times as bright as 64

175 is twice as bright as 128

93 is half as bright as 128

128 + 128 = 175

128 / 2 = 93

See also: http://www.inforamp.net/~poynton/GammaFAQ.html

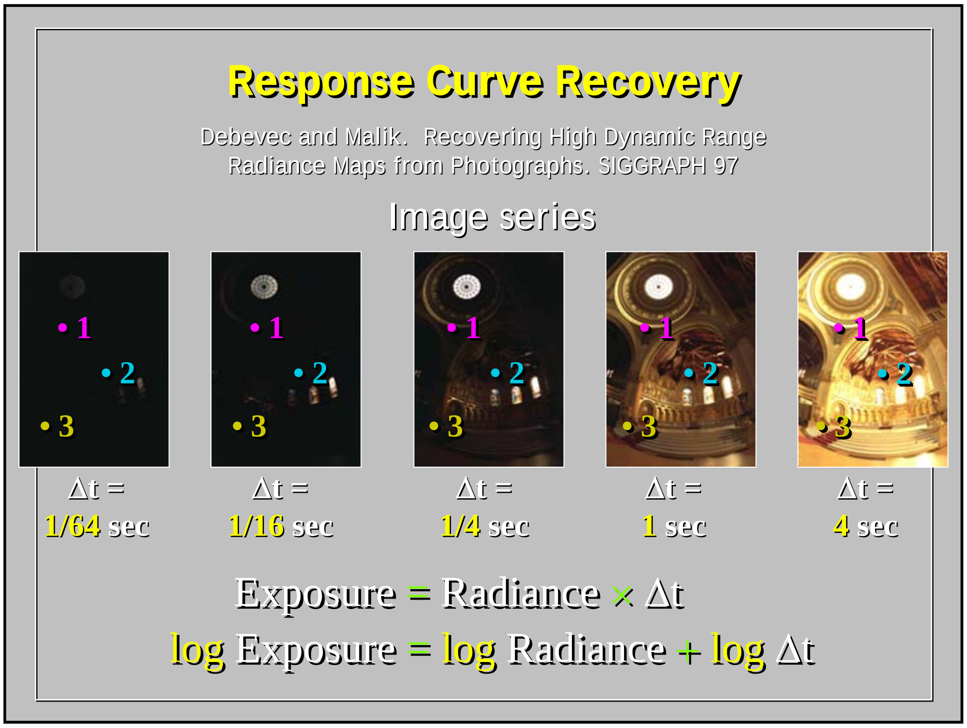
Response Curve Recovery
Debevec and Malik. Recovering High Dynamic Range Radiance Maps from Photographs. SIGGRAPH 97
Image series
• 1
• 2
• 3
Δt = 1/64 sec
Δt = 1/16 sec
Δt = 1/4 sec
Δt = 1 sec
Δt = 4 sec
Exposure = Radiance × Δt
log Exposure = log Radiance + log Δt

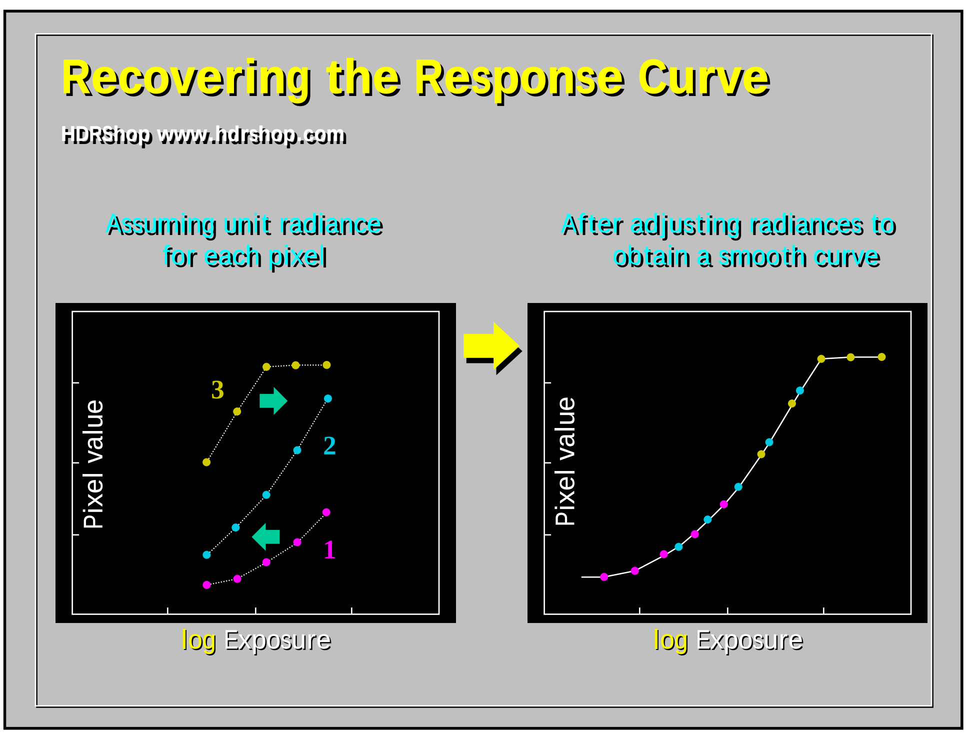
Recovering the Response Curve
HDRShop www.hdrshop.com
Assuming unit radiance for each pixel
After adjusting radiances to obtain a smooth curve
3
2
1
Pixel value
log Exposure
Pixel value
log Exposure

Mitsunaga-Nayar CVPR 99 Response Recovery

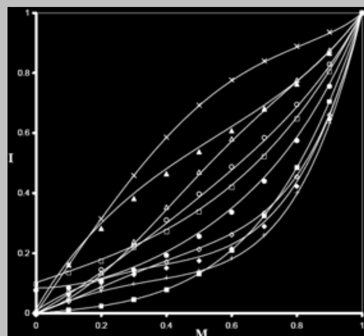

$$I = f(M) = \sum_{n=0}^{N} c_n M^n$$

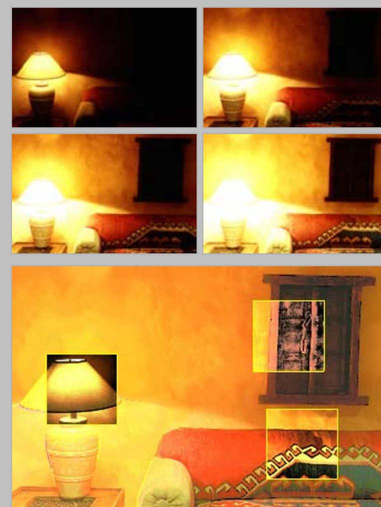

Assumes polynomial form
Estimates unknown exposure times
Vignetting compensation
http://www.cs.columbia.edu/CAVE/

Grossberg and Nayar – Response recovery from image histograms

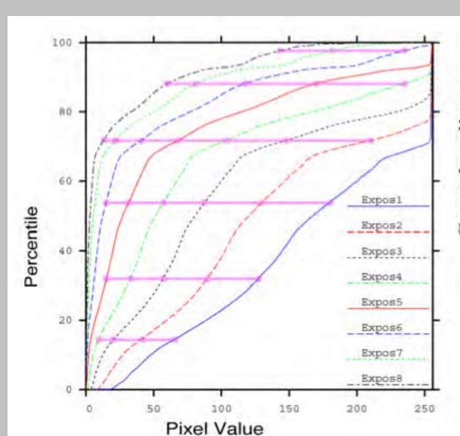

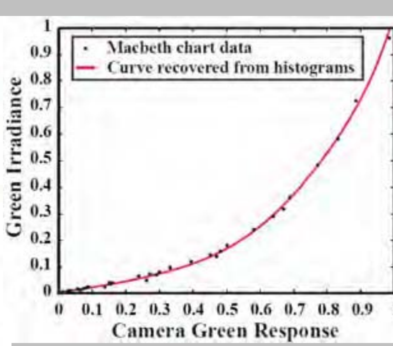

Robust to scene motion

What can be Known about the Radiometric Response from Images? ECCV 2002

Lin, Gu, Yamazaki, Shum. ***Radiometric Calibration from a Single Image***. CVPR 2004

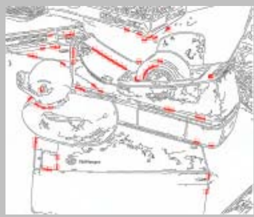

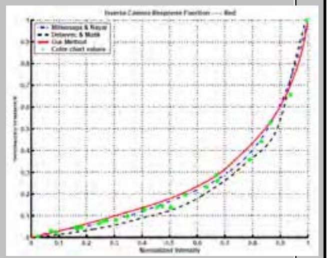

Assumes that edge regions are blends of two colors, with all intermediate values represented equally (in linear space)

Locates edge regions in image

Estimates response based on intermediate value distributions

Ways to vary exposure

- Shutter Speed (preferred)
- F/stop (aperture, iris)

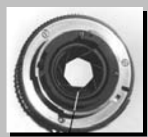

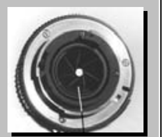

- Neutral Density (ND) Filters

Shutter Speed

Ranges: Canon D30: 30 to 1/4,000 sec.

Sony VX2000: ¼ to 1/10,000 sec.

Pros:

- Directly varies the exposure
- Usually accurate and repeatable

Issues:

- Digital: Noise in long exposures
- Film: Reciprocity failure at > ~ 5 sec.

Shutter Speed

Note: shutter times usually obey a power series – each "stop" is a factor of 2

¼, 1/8, 1/15, 1/30, 1/60, 1/125, 1/250, 1/500, 1/1000 sec

Usually really is:

¼, 1/8, 1/16, 1/32, 1/64, 1/128, 1/256, 1/512, 1/1024 sec

F/stop (aperture)

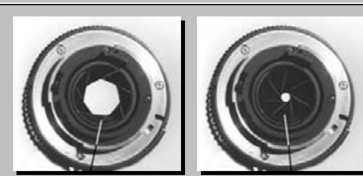

Ranges: Canon D30: f/2.8 to f/22
Sony VX2000: f/1.6 to f/11

Standard f-numbers: 2, 2.8, 4, 5.6, 8, 11, 16, 22

Exposure is proportional to the inverse square of the f-number:

f/22 is 1/64 the light of f/2.8 (6 stops)

Pros:

Can use aperture when you run out of shutter speed variation

F/stop (aperture)

Issues:

Changes depth of field
Less repeatable
Limited range of exposure variation
=> Not first choice for HDR

Radial Brightness Falloff (*Vignetting*)

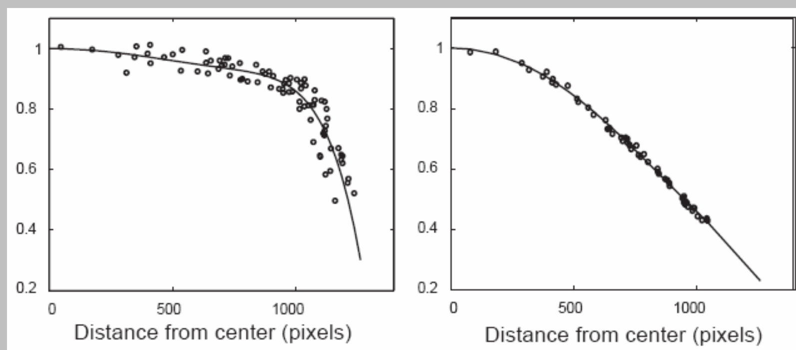

f/16 Sigma 8mm fisheye lens f/8

From Stumpfel et al, *Direct HDR Capture of the Sun and Sky*, Afrigraph 2004

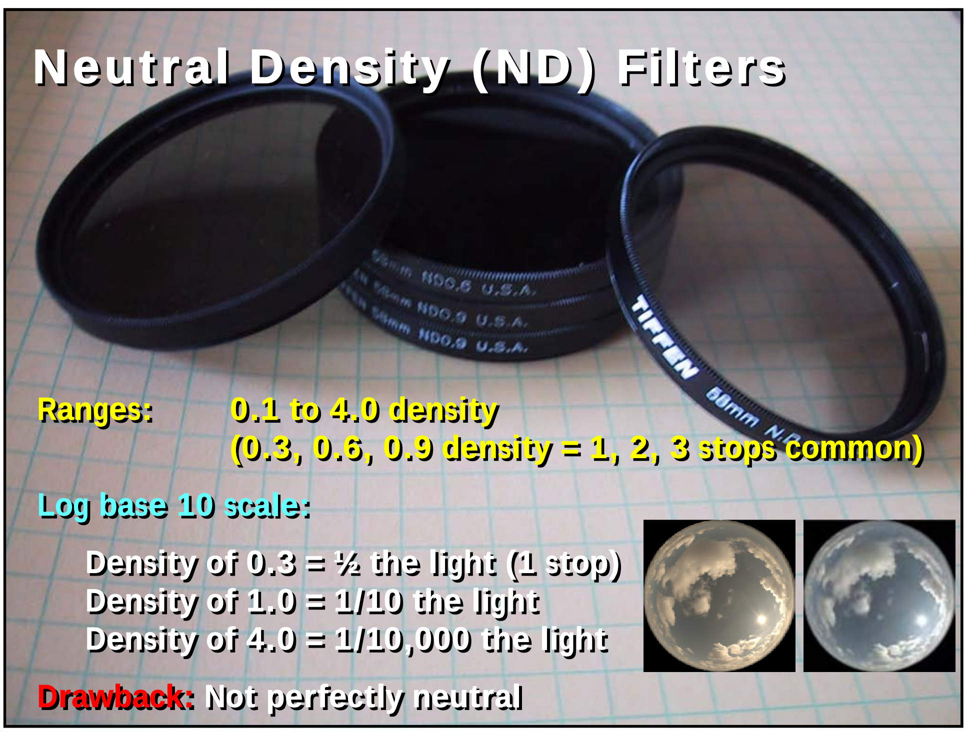

Gain / ISO / Film Speed

Range: ISO 100 to 1600
0dB to 18dB (3dB = factor of 2)

Problem:
Usually just adds noise to your image

+18dB Gain on Sony VX2000

Storing High Dynamic Range Images

- Portable FloatMap .pfm
- Floating-Point TIFF .tif
- RADIANCE .hdr, .pic
- LogLuv TIFF .tif
- OpenEXR .exr
- ...

- See Greg’s talk coming up…

Methods for taking light probes: omnidirectional HDR images

- Mirrored ball + camera
- Fisheye lens images
- Panoramic camera
- Image stitching

Sources of Mirrored Balls

- 2-inch chrome balls ~ $20 ea.
  - McMaster-Carr Supply Company
    www.mcmaster.com
- 6-12 inch large gazing balls
  - Baker's Lawn Ornaments
    www.bakerslawnorn.com
- Hollow Spheres, 2in – 4in
  - Dube Juggling Equipment
    www.dube.com
- FAQ on www.hdrshop.com

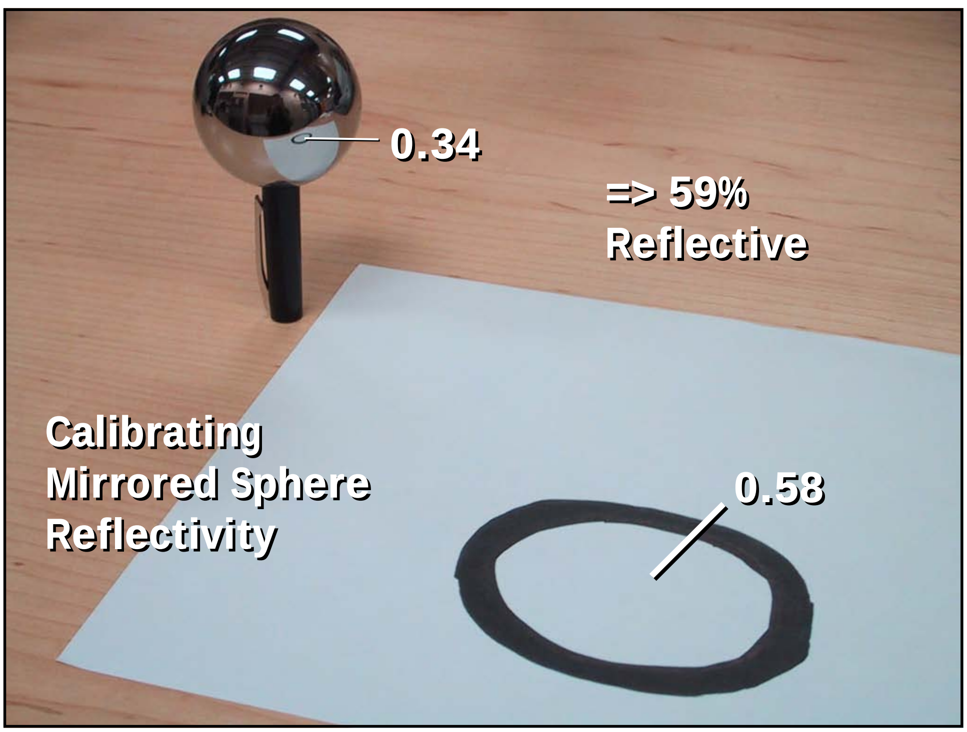

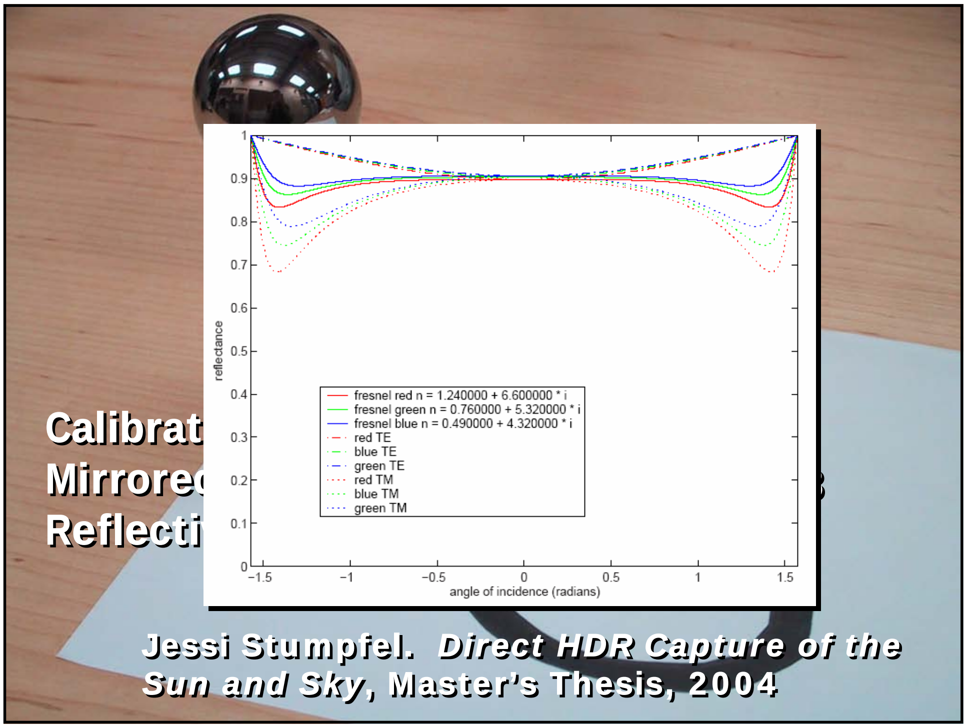
reflectance
angle of incidence (radians)
fresnel red n = 1.240000 + 6.600000 * i
fresnel green n = 0.760000 + 5.320000 * i
fresnel blue n = 0.490000 + 4.320000 * i
red TE
blue TE
green TE
red TM
blue TM
green TM
Jessi Stumpfel. *Direct HDR Capture of the Sun and Sky*, Master's Thesis, 2004


Fisheye Images
Canon

Fisheye Lens Radial Falloff

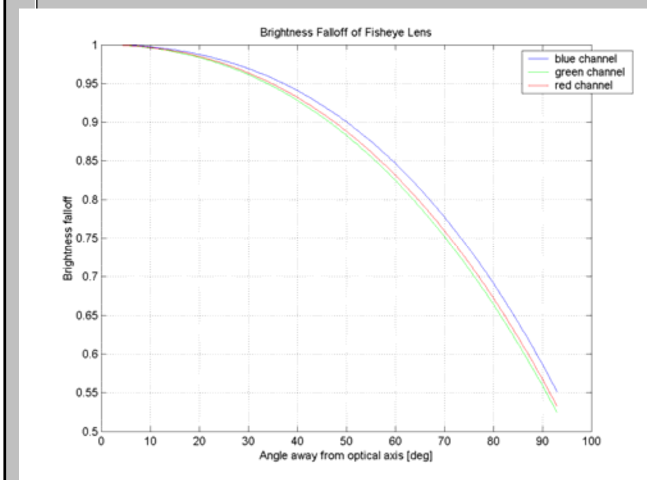

Scanning Panoramic Cameras (Panoscan, Spheron)

Pros:

- *very* high res (10K x 7K+)
- Full sphere in one scan – no stitching
- Good dynamic range, some are HDR

Issues:

- More expensive
- Scans take a while

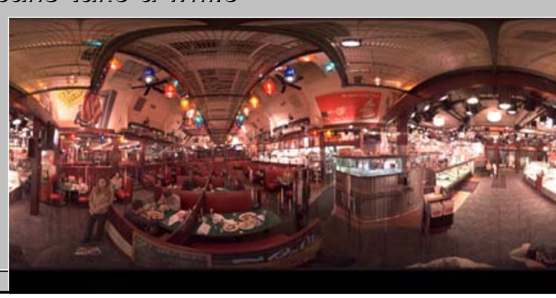

Tiled Photographs

Tiled Photographs – Nodal Acquisition Rig

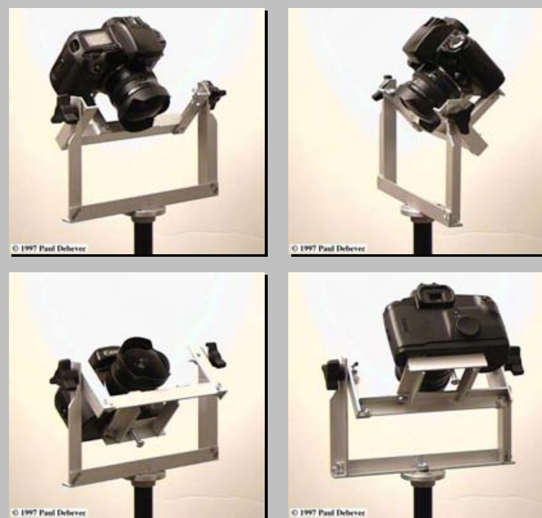

See also www.kaidan.com

Stitching HDRI with Realviz Stitcher

http://www.gregdowning.com/HDRI/stitched/

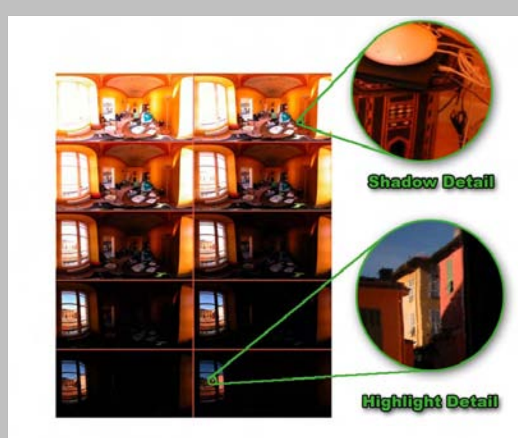

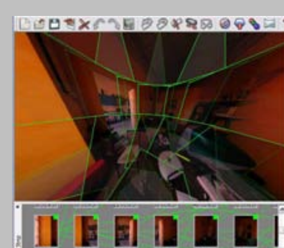

Types of Omnidirectional Images

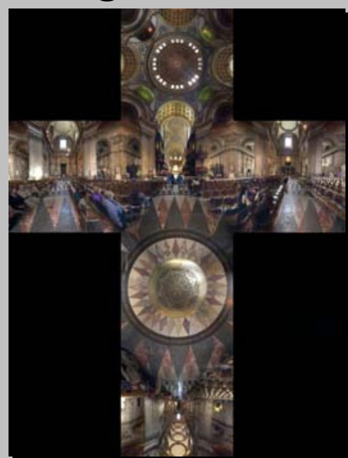

Cube Map

Latitude/Longitude

Types of Omnidirectional Images

Mirrored Ball

Angular Map

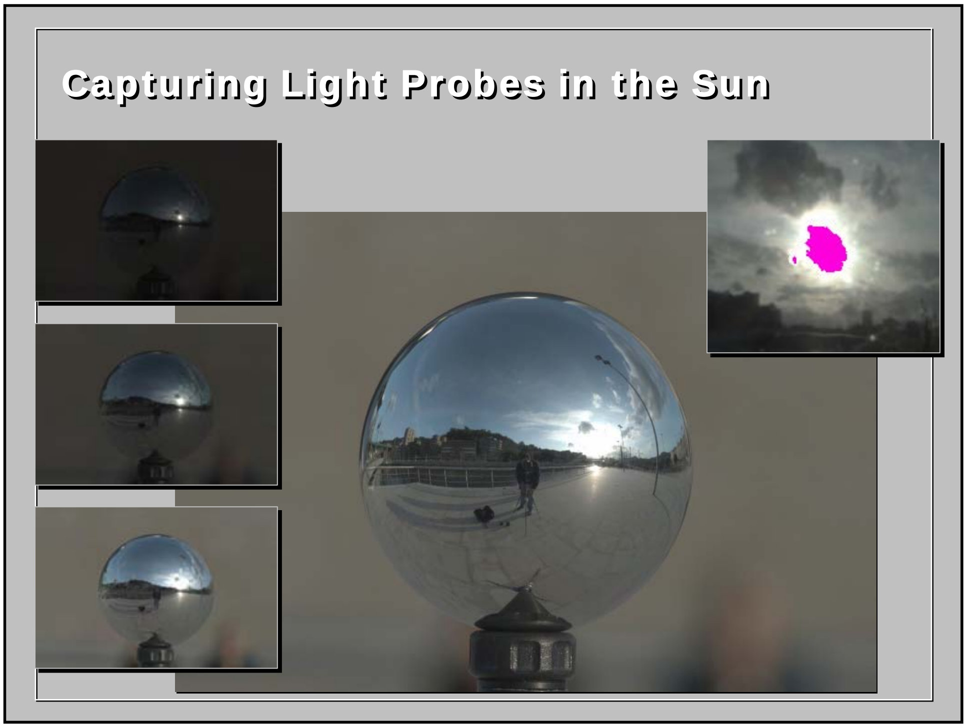
Capturing Light Probes in the Sun

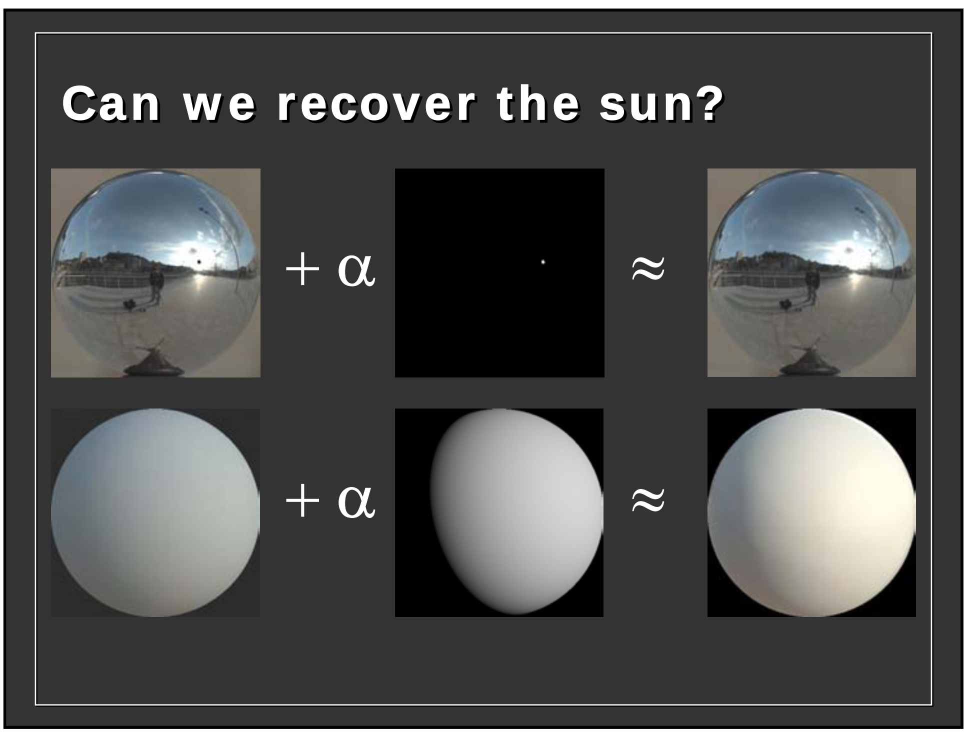
Can we recover the sun?
+ α
≈
+ α
≈

Photograph Diffuse Sphere

Solve for Sun Scaling Factor

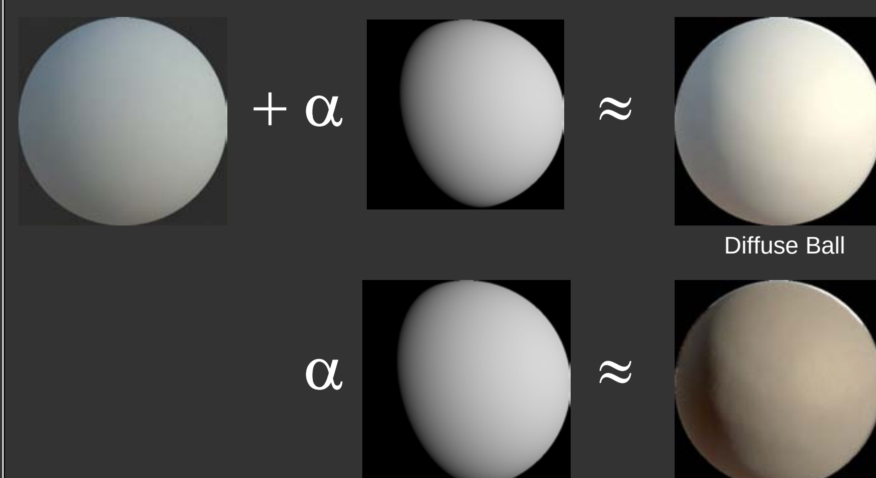

$\alpha = (1.166, 0.973, 0.701)$

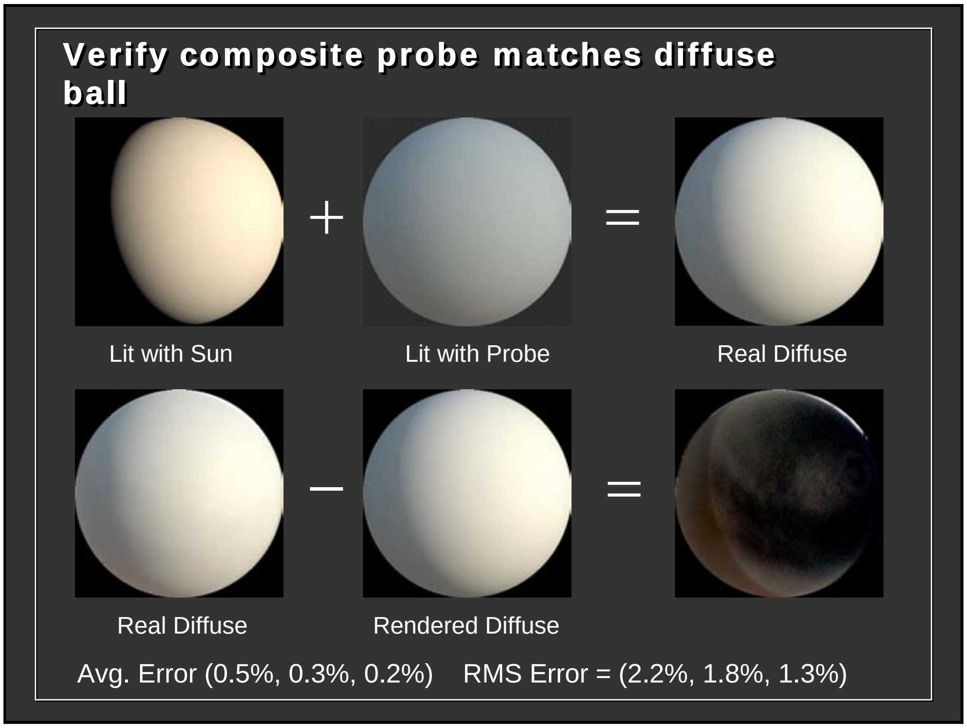
Verify composite probe matches diffuse ball
+
=
Lit with Sun
Lit with Probe
Real Diffuse
−
=
Real Diffuse
Rendered Diffuse
Avg. Error (0.5%, 0.3%, 0.2%)
RMS Error = (2.2%, 1.8%, 1.3%)

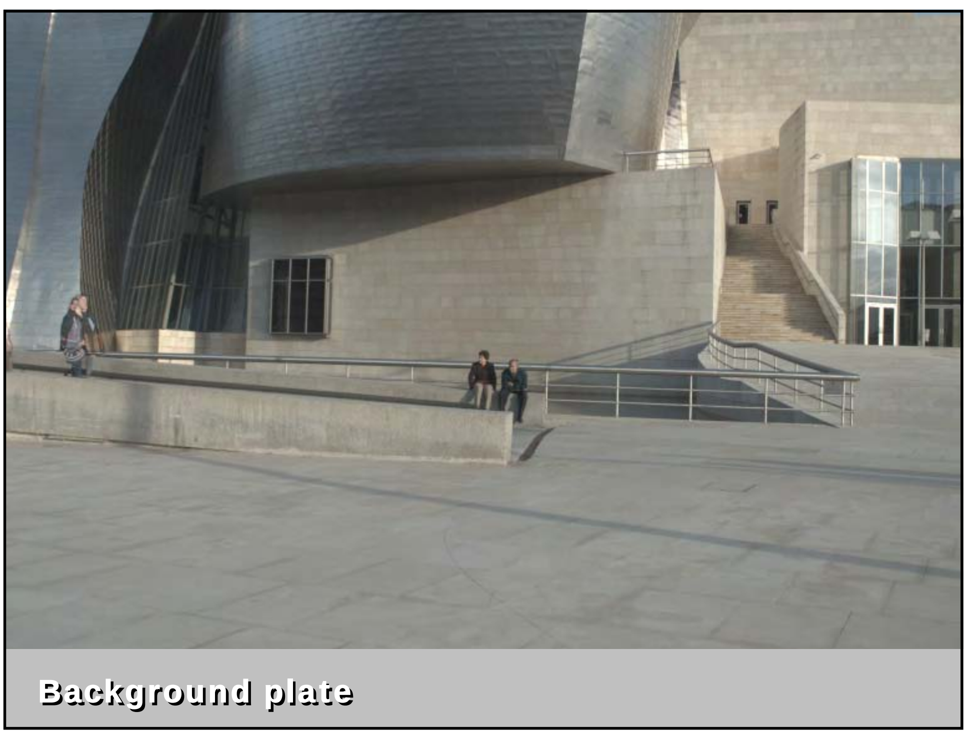
Background plate

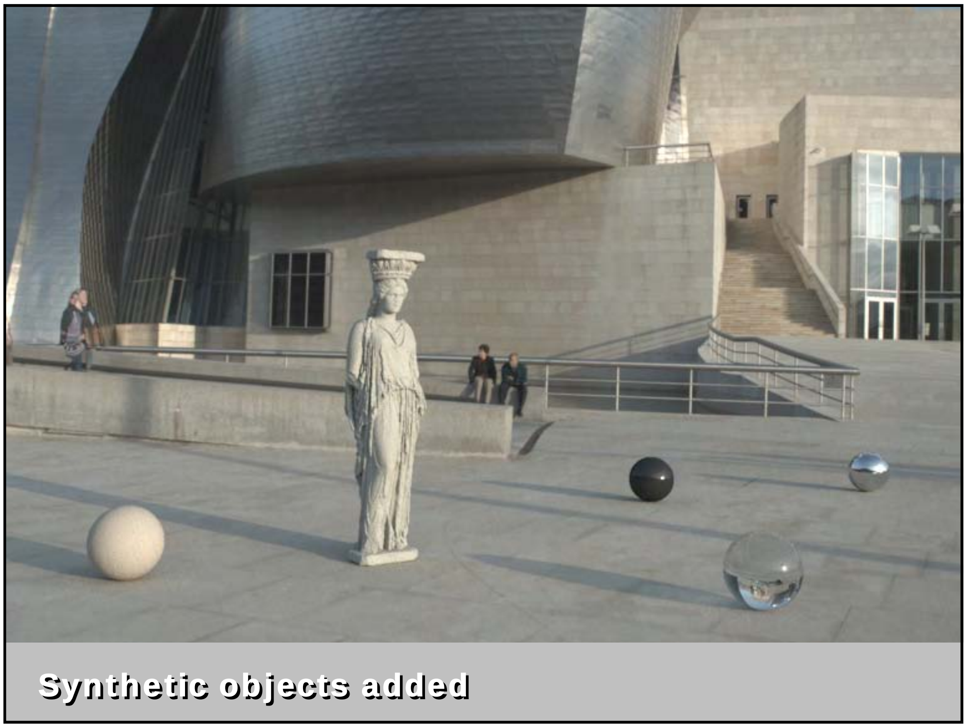
Synthetic objects added

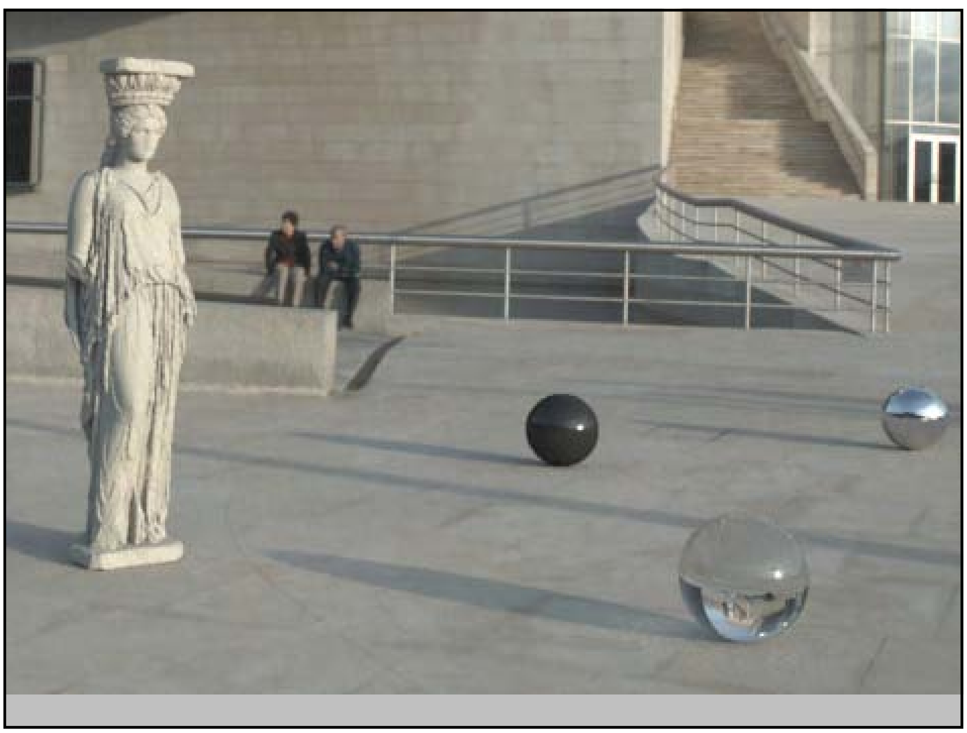

Direct HDR Capture of the Sun and Sky

- Use Sigma 8mm fisheye lens and Canon EOS 1Ds to cover entire sky
- Use 3.0 ND filter on lens back to cover full range of light

Stumpfel, Jones, Wenger, Tchou, Hawkins, and Debevec. “Direct HDR Capture of the Sun and Sky”. To appear in Afrigraph 2004.

Extreme HDR Image Series

1 sec
f/4

1/4 sec
f/4

1/30 sec
f/4

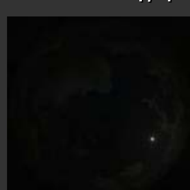

1/30 sec
f/16

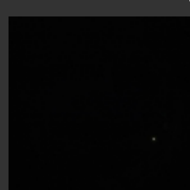

1/250 sec
f/16

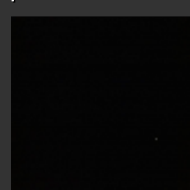

1/1000 sec
f/16

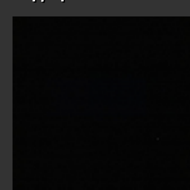

1/8000 sec
f/16

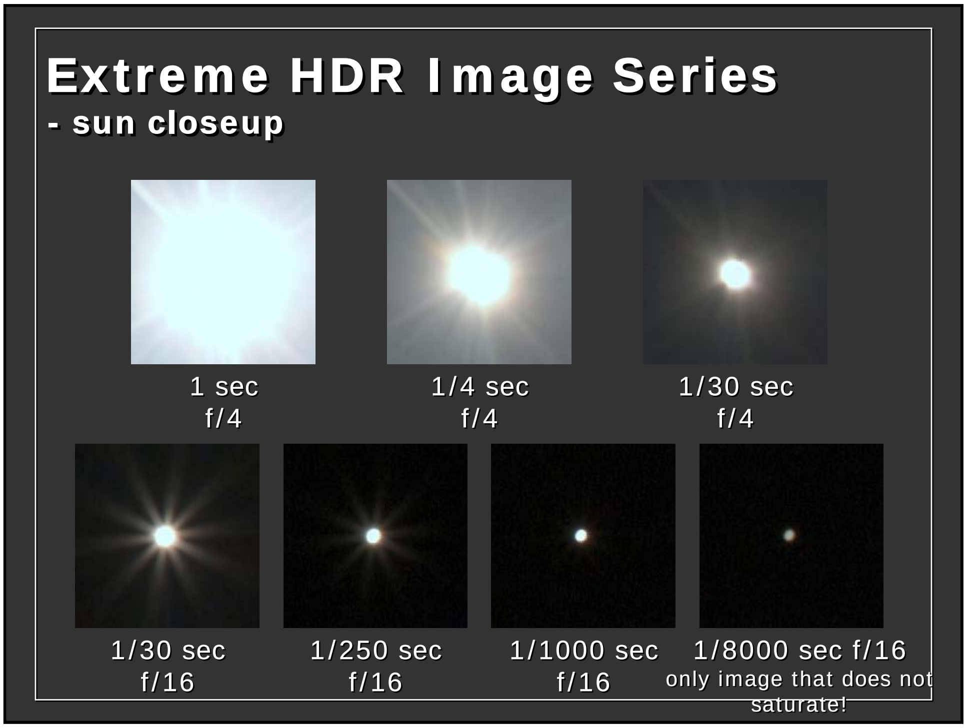
Extreme HDR Image Series
- sun closeup
1 sec
f/4
1/4 sec
f/4
1/30 sec
f/4
1/30 sec
f/16
1/250 sec
f/16
1/1000 sec
f/16
1/8000 sec f/16
only image that does not saturate!


Lit by sun and sky probe

SIGGRAPH2006

HDR Image Representation

Greg Ward

SIGGRAPH2006

HDRI Representation

Greg Ward

Anyhere Software

Dynamic Range

From Ferwerda et al, SIGGRAPH '96

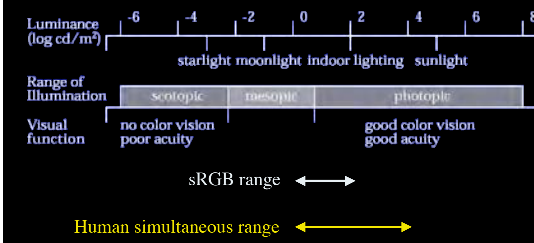

Value Encoding Methods

- Linear quantization
- Gamma function (e.g., CRT curve)
- Logarithmic encoding
- Floating point
- Perceptual

Linear vs. Gamma Encoding

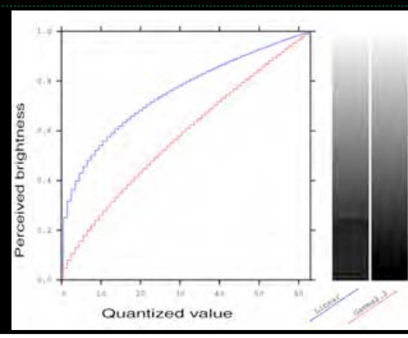

Logarithmic vs. Floating Point

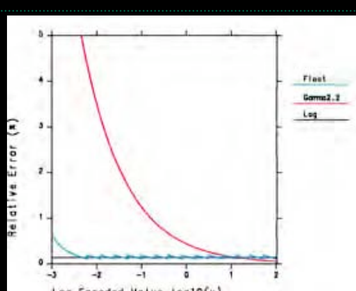

Fictitious 12-bit encoding

Perceptual Encoding

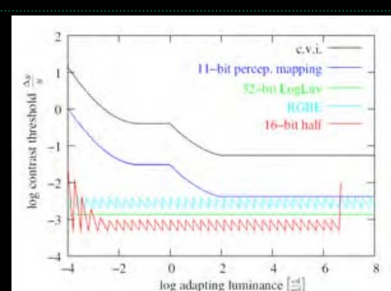

Mantiuk et al., SIGGRAPH 2004

CCIR-709 Color Space

- Human visible gamut is much larger than standard display's
- Saturated blues, greens, and purples are lost in sRGB
- Many HDR image formats also cover a larger color gamut

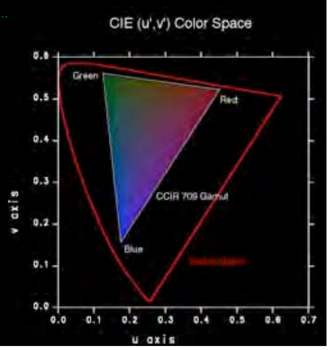

HDR Image & Video Formats

- Available high dynamic-range formats:
  - *Radiance* 32-bit RGBE and XYZE pictures
  - TIFF 48-bit integer and 96-bit float formats
  - SGI 24-bit and 32-bit LogLuv TIFF
  - ILM OpenEXR format
  - JPEG-HDR format
- Proposals and extensions:
  - HDR extensions to MPEG from MPI [Mantiuk et al. 2004]
  - HDR extensions to JPEG 2000 from UFL [Xu et al. 2005]
  - HDR texture compression (two papers at this conference)

Encoding Comparison Chart

Encoding	Bits / pixel	Dynamic Range	Quant. Step	Covers Gamut
sRGB	24	1:10^1.6	Variable	No
Radiance RGBE	32	1:10^76	1%	No
Radiance XYZE	"	"	"	Yes
LogLuv 24	24	1:10^4.8	1.1%	Yes
LogLuv 32	32	1:10^38	0.3%	Yes
OpenEXR	48	1:10^10.7	0.1%	Yes
JPEG-HDR	1-7	1:10^9.5	Variable	Can

Full Gamut Encodings

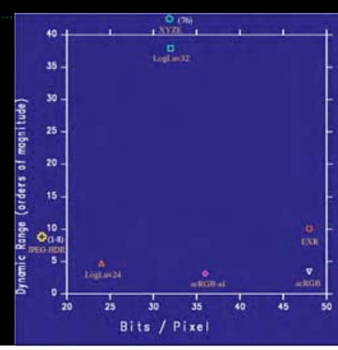

HDR Acid Test Image

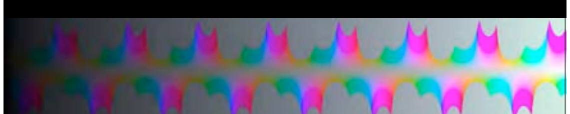

1:10⁸ dynamic range, covering visible gamut

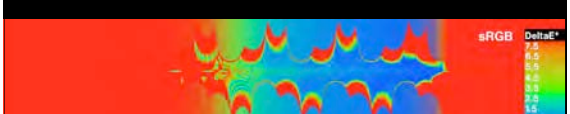

Visible error for 24-bit/pixel sRGB encoding

48-bit RGB TIFF Accuracy

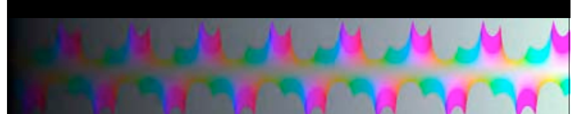

1:10⁸ dynamic range, covering visible gamut

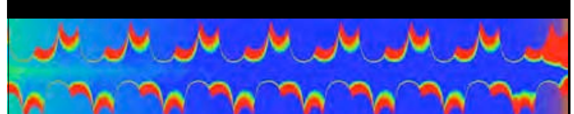

Visible error for 48-bit/pixel RGB encoding (γ = 2.2)

Radiance RGBE Accuracy

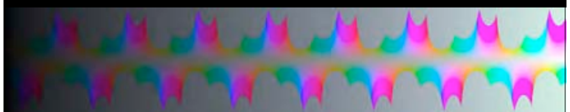

1:10^8 dynamic range, covering visible gamut

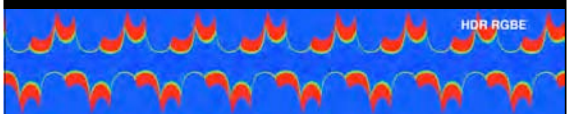

Visible error for 32-bit/pixel RGBE encoding

Radiance XYZE Accuracy

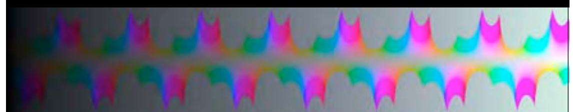

1:10^8 dynamic range, covering visible gamut

Visible error for 32-bit/pixel XYZE encoding

TIFF LogLuv24 Accuracy

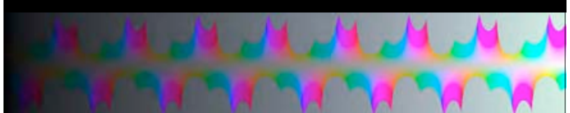

1:10^8 dynamic range, covering visible gamut

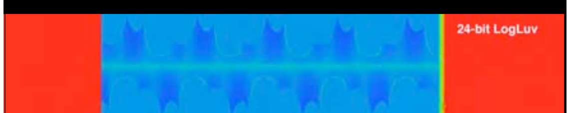

Visible error for 24-bit/pixel LogLuv encoding

TIFF LogLuv32 Accuracy

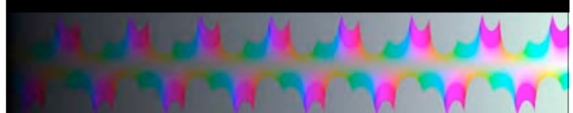

1:10^8 dynamic range, covering visible gamut

Visible error for 32-bit/pixel LogLuv encoding

OpenEXR Accuracy

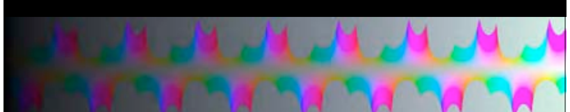

1:10^8 dynamic range, covering visible gamut

Visible error for 48-bit/pixel EXR `half` encoding

HDR Format Performance

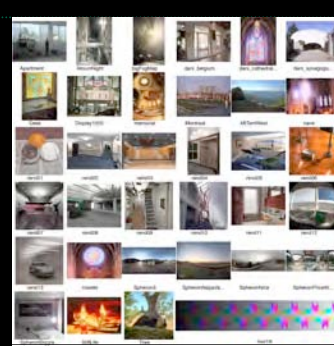

Test Set
(34 HDR images)

Measured i/o speed and image size

Read & Write Speed

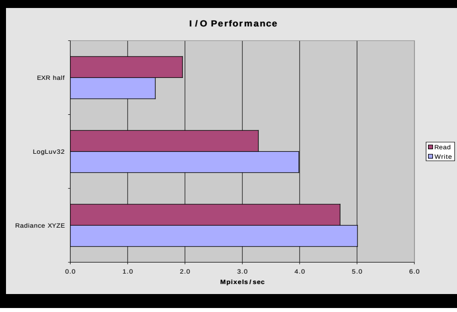

Total Image Size

Radiance RGBE and XYZE

- Simple format with free source code
- 8 bits each for 3 mantissas and 1 exponent
- 76 orders of magnitude in 1% steps
- Run-length encoding (20% avg. compr.)
- RGBE format does not cover visible gamut
- Dynamic range at expense of accuracy
- Color quantization not perceptually uniform

Radiance Format (.pic, .hdr)

(145, 215, 87, 149) =
(145, 215, 87) * 2^(149-128) =
(1190000, 1760000, 713000)

(145, 215, 87, 103) =
(145, 215, 87) * 2^(103-128) =
(0.00000432, 0.00000641, 0.00000259)

Ward, Greg. "Real Pixels," in Graphics Gems IV, James Arvo ed., Academic Press, 1994

IEEE 96-bit TIFF

- Most accurate representation
- Support (with compression) in Photoshop CS2
- Uncompressed files are enormous
  - 32-bit IEEE floats look like random bits

16-bit/sample TIFF (RGB48)

- Supported by Photoshop and TIFF library
- 16 bits each of log red, green, and blue
- 5.4 orders of magnitude in < 1% steps
- LZW lossless compression available
- Does not cover visible gamut
- Most applications think of max. value as “white”

SGI 24-bit LogLuv TIFF Codec

- Implemented in Leffler's TIFF library
- 10-bit LogL + 14-bit CIE (u',v') lookup
- 4.8 orders of magnitude in 1.1% steps
- Just covers visible gamut and range
- Amenable to tone-mapping as look-up
- Dynamic range is less than we would like
- No compression

24 -bit LogLuv Pixel

From Larson, CIC '98

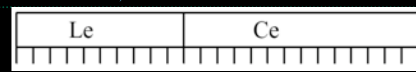

$$L_e = \lfloor 64(\log_2 L + 12) \rfloor$$

$$u' = \frac{4x}{-2x+12y+3}$$

$$v' = \frac{9y}{-2x+12y+3}$$

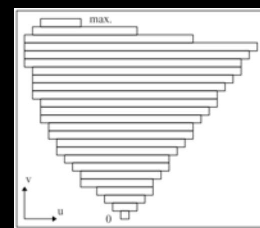

SGI 32-bit LogLuv TIFF Codec

- Implemented in Leffler's TIFF library
- 16-bit LogL + 8 bits each for CIE (u',v')
- 38 orders of magnitude in 0.3% steps
- Run-length encoding (30% avg. compr.)
- Allows negative luminance value
- Amenable to tone-mapping as look-up

32-bit LogLuv Pixel

From Larson, JGT '98

$$L_e = \lfloor 256(\log_2 L + 64) \rfloor$$

$$u_e = \lfloor 410u' \rfloor$$

$$v_e = \lfloor 410v' \rfloor$$

$$u' = \frac{4x}{-2x+12y+3}$$

$$v' = \frac{9y}{-2x+12y+3}$$

ILM OpenEXR Format

- 16-bit/primary floating point (sign-e5-m10)
- 9.6 orders of magnitude in 0.1% steps
- Additional order of magnitude near black
- Wavelet compression of about 40%
- Negative colors and full gamut RGB
- Alpha and multichannel support
- Open Source I/O library released Fall 2002
- Slow to read and write

ILM's OpenEXR (.exr)

6 bytes per pixel, 2 for each channel, compressed

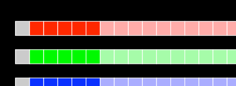

sign exponent mantissa

- Several lossless compression options, 2:1 typical
- Compatible with the "half" datatype in NVidia's Cg
- Supported natively on GeForce FX and Quadro FX

- Available at www.openexr.com

BrightSide Technology's JPEG-HDR Format

- Backwards-compatible JPEG extension for high dynamic range images
- Very compact: 1/10th size of other formats
- Naïve software displays tone-mapped sRGB
  - Different tone-mappings possible
- Desaturation can encompass visible gamut
- Lossy encoding so repeated read/write degrades
- Expensive (three pass) write process

File Size & HDR Adoption

- Compression can match size of JPEG images + 20%
- Rationale for "lossy" HDR:
  - Lossy encodings are all about perception
  - Lossy HDR supports display to the limits of human vision
  - Required for digital photography & web applications
  - Mantiuk et al.'s MPEG-4 extension (SIGGRAPH 2004)
  - Xu et al.'s JPEG-2000 extension (CG+A 2005)
  - Two 2006 papers on 8-bit/pixel HDR texture compression
- What if HDR format was backwards-compatible?
  - JPEG-HDR & new MPI technique (SIGGRAPH 2006)

JPEG-HDR Format

Ward & Simmons 2004 & 2005

1. Tone-map HDR input into 24-bit sRGB
2. Write as *output-referred* JPEG
3. Record restorative information as metadata (supplement)

- Naïve applications see tone-mapped image
- HDR applications use supplement to recover *scene-referred* original
- Similar to Kodak's ERI (Spaulding et al. 2003)

JPEG-HDR Encoding Process

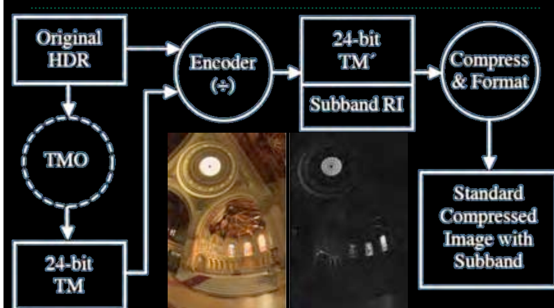

RI(x,y) = Y(HDR(x,y)) / Y(TM(x,y))

Decoding Process

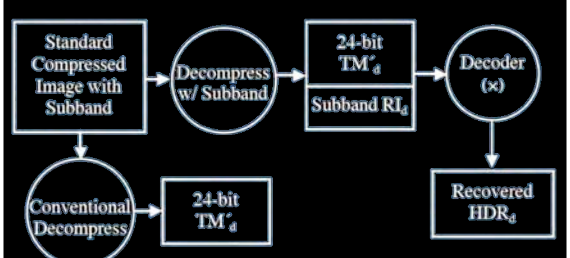

Compressed JPEG-HDR size: 1-7 bits/pixel (between 1/3 & 1/20 size of other HDR formats)

JPEG-HDR Software Availability

- Implemented as extension to Tom Lane's public JPEG library (www.ijg.org)
  - Tools & libraries included on HDRI book's DVD
  - Contact BrightSide Technologies for licensing info. at www.brightsidetech.com/software
- JPEG-HDR included in Photosphere
  - Handy export function for batch conversion and webpage creation
  - Download from www.anyhere.com

HDRI Encoding Conclusions

- Sufficient still formats to meet most needs:
  - Radiance RGBE for legacy systems
  - TIFF for greatest encoding variety
  - OpenEXR for good accuracy and support
  - JPEG-HDR for space efficiency
- HDR texture formats are in the works
- HDR video formats are coming soon

HDR Capture Refinements

- Automatic exposure alignment
- “Ghost” removal
- Lens flare removal
- Implementing HDR in still & video cameras

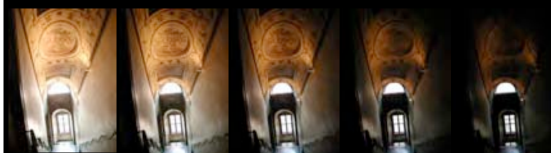

LDR Exposure Registration

[Ward 2003, *Journal of Graphics Tools*, 8(2)]

The *median threshold bitmap* (MTB) allows us to quickly compare and align different images, because it is constant with respect to exposure for any camera with a monotonic response function

The same is not true for an edge map, which changes with exposure even with careful normalization and approximate response curves

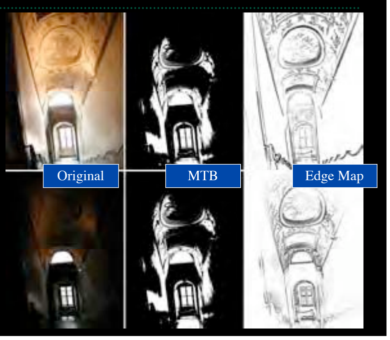

Image Pyramid Alignment

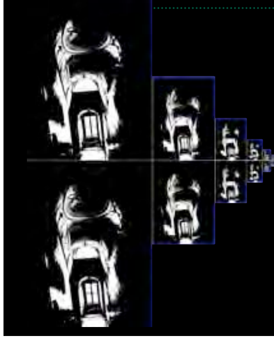

Grayscale images are scaled down repeatedly to create an image pyramid, which is then converted into MTBs for comparison

The smallest images are aligned first within a ±1 pixel distance, which corresponds to a ±32 pixel distance in the original

This becomes the MSB in the offset, which is shifted and used as the starting point for the next higher resolution alignment, and so on to the top

Alignment Results

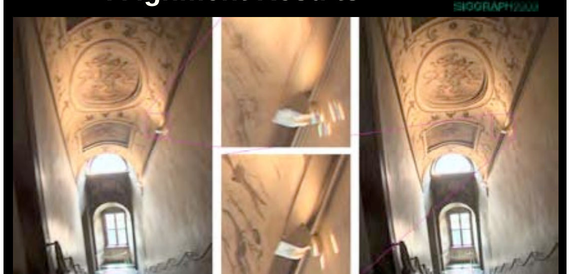

5 unaligned exposures | Close-up detail | MTB alignment

Time: About .2 second/exposure for 3 MPixel image

Automatic “Ghost” Removal

Before | After

Object Movement

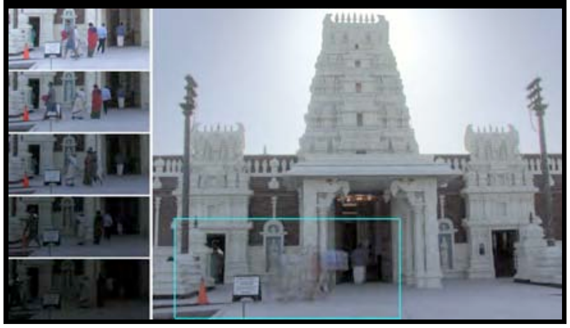

Variance-based Detection

Region Masking

Best Exposure in Each Region

Lens Flare Removal


HDR capture of perfect hole cut in aluminum foil

- From left image, we may directly measure the lens Point Spread Function (PSF)
- PSF is a function of focal length and aperture, so comprehensive measurement is impractical

More Usual Input

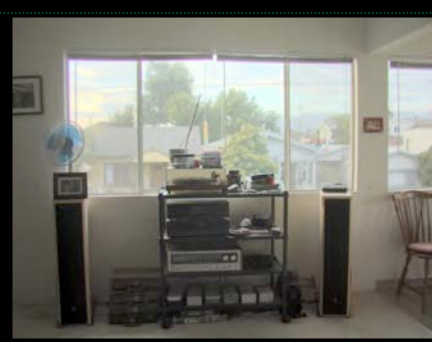

Estimate PSF from HDR Capture

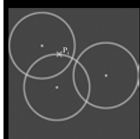

- For each “hot” pixel in input:
  - Find “coldest” relative pixel from at each radius
  - Consider overlapping hot pixel contributions
- PSF is minimal, monotonically decreasing function measured relative to cold pixels
  - Computed by fitting 3rd degree polynomial over all identified cold pixels:

“Cold Pixel” $$P_i = \sum_j P_j \cdot \left(C_0 + \frac{C_1}{r_{ij}} + \frac{C_2}{r_{ij}^{\,2}} + \frac{C_3}{r_{ij}^{\,3}} \right)$$

Fitted vs. Measured PSF

PSF estimate (apt. capture fit vs. tin foil spot)

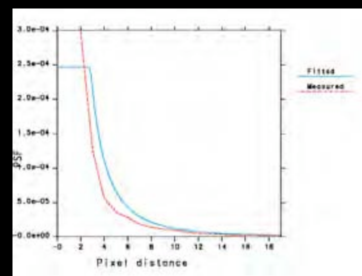

Apply: Simulate Flare & Subtract

Before After Difference

Photosphere Demo

- HDRI Browsing & Cataloging Application
  - Also builds HDRI’s from bracketed exposures
- Available from www.anyhere.com
  - Mac OS X app., Linux command-line tool

Launch Photosphere

Fast Forward to HDR Cameras

- Leverage CMOS Sensor Technology
  - Fuji has sensor with dual-sensitivity pixel
  - SMaL Camera has log sensor
  - Pixim sensor has local pixel exposure
- Alter camera/sensor design
  - Multi-image capture using modified scanout
  - Multiple sensors
  - Spatially varying filters for video mosaicing
  - Sensors with assorted pixels
  - Adaptive dynamic range system

Two-Exposure HDR

- Compared Results to 5 exp. on 12 Scenes
  - Two exposures was usually sufficient
  - Less noise averaging but otherwise comparable
- Camera Implementation Reduces Artifacts
  - No alignment issues on short exposures
  - Longer exposure akin to “slow flash”
- Marginal Manufacturing Cost: $0.00

Interline CCD Scanout

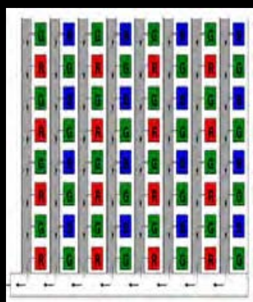

Old Program:

Electronic shutter holds each exposure during scanout

Preview/movie uses electronic shutter, while still capture relies on mechanical shutter

New Program:

Instead, shift pixels under electronic shutter with 1/16th of mechanical exposure still remaining

After scanning out long exposure, shift and scan out short exposure

Result: two exposures separated by 4 f-stops

Sample Results

Assorted Pixels

Nayar and Mitsunaga, IEEE CVPR 2000
Nayar and Narasimhan, ECCV 2002

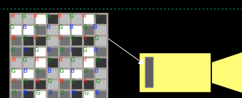

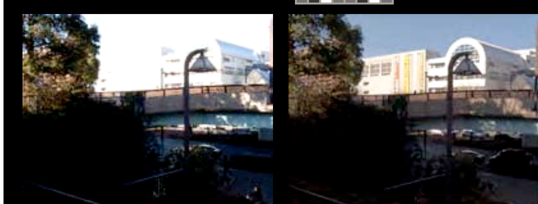

Digital Still Camera

Camera with Assorted Pixels

Adaptive Dynamic Range

Nayar and Branzoi, ICCV 2003

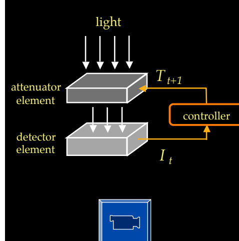

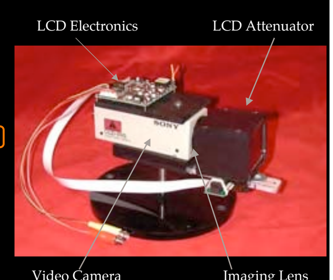

Conclusions

- HDR capture is still under active development
- HDR video is challenging but has many potential benefits
- Movie industry is an early adopter
- Home entertainment market will soon follow

HDR Video and Applications

Karol Myszkowski

High Dynamic Range Video

Rafal Mantiuk, Grzegorz Krawczyk, **Karol Myszkowski**
MPI Informatik, Germany

Overview

- HDR video compression
- Real-time HDR playback with perceptual effects
- HDR Visible Differences Predictor
- HDR video environment maps

HDR Video Compression

SIGGRAPH2006

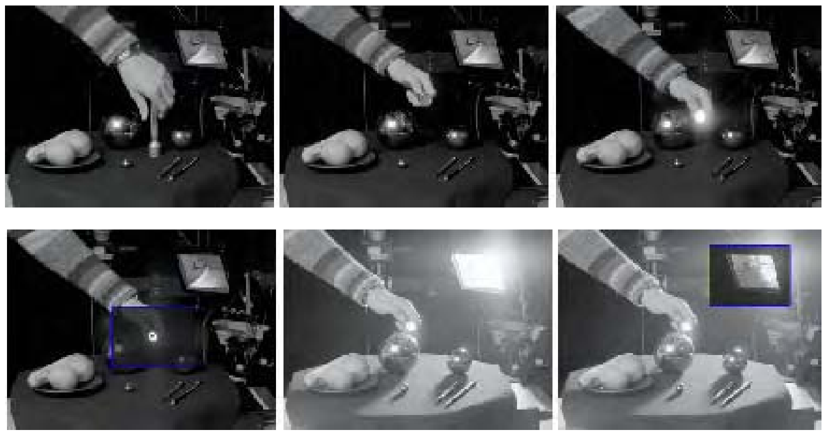

http://www.mpi-inf.mpg.de/resources/hdr/compression

HDR Video Compression

SIGGRAPH2006

- Color space for HDR pixel encoding
- 12-bit HDR MPEG
- 8-bit Backward Compatible HDR MPEG
- Demo

Standard *HDR-ready* Image and Video Formats

- Basic requirement
  - Capability of storing larger number of bits
- Formats capable of HDR
  - JPEG-2000
    - ➔Up to **16**-bit integers per component
    - ➔HDR extension [Xu et al. IEEE CG&A 2005]
  - MPEG-4 (IOS/IEC 14496-2 or ISO/IEC 14496-10)
    - ➔Up to **12**-bit integers per component
    - ➔HDR extension [Mantiuk et al. ACM Siggraph 2004]
- Needed special color space for HDR pixel encoding

Color Space for HDR pixels

- Full visible range of luminance and color gamut
- Perceptually uniform
- Only positive integer numbers
- Quantization below the detection threshold
- Minimum correlation between color channels
- Direct transformation: color space <--> XYZ

Luminance and Luma

SIGGRAPH2006

$$\underbrace{l}_{\text{luma}} \rightarrow \underbrace{y}_{\text{luminance}} : y(l)$$

From Wikipedia:

Luma is a new word proposed by the NTSC in 1953 (..) While luminance is the weighted sum of the linear RGB components of a color video signal, proportional to intensity, luma is the weighted sum of the non-linear R'G'B' components **after gamma correction** has been applied ...

$$y(l_o + 0.5) - y(l_o) \approx 0.5 \cdot \frac{dy}{dl}$$

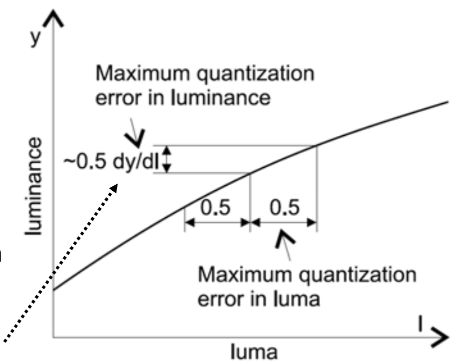

Luminance ←→ Luma

SIGGRAPH2006

Maximum quantization error < the detection threshold (JND)

$$0.5 \cdot \frac{dy(l)}{dl} < t(y_{\text{adapt}})$$

Adaptation to a single pixel: $y_{\text{adapt}} = y$

Inequality to equality:

$$\frac{dy(l)}{dl} = 2 \cdot \frac{t(y)}{k} \qquad \text{JND scaled by } k > 1$$

Boundary conditions:

$$y(0) = 10^{-5}\left[\frac{\text{cd}}{\text{m}^2}\right], \quad y(l_{\max}) = 10^{10}\left[\frac{\text{cd}}{\text{m}^2}\right] \quad \text{for } l_{\max} = 2^{\text{bits}} - 1$$

Contrast Detection Models: $t(y_{adapt})$

- Threshold versus intensity function (t.v.i.)
  - Measured for a fixed pattern on an uniform background
    - ➔Ferwerda’s t.v.i. (computer graphics)
    - ➔Blackwell’s t.v.i. (CIE 19/2.1 standard)

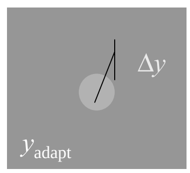

Contrast Detection Models: $t(y_{adapt})$

- Contrast Sensitivity Function (CSF)
  - Measured for sinusoidal patterns or Gabor patches of various spatial frequencies and orientations:
    - ➔Barten’s CSF model (DICOM’s grayscale standard display function)

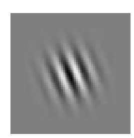

    - ➔CSF model used in Daly’s Visible Differences Predictor
  - for each y_{adapt} take the maximum sensitivity

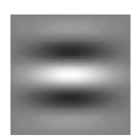

Contrast Detection Models: $t(y_{adapt})$

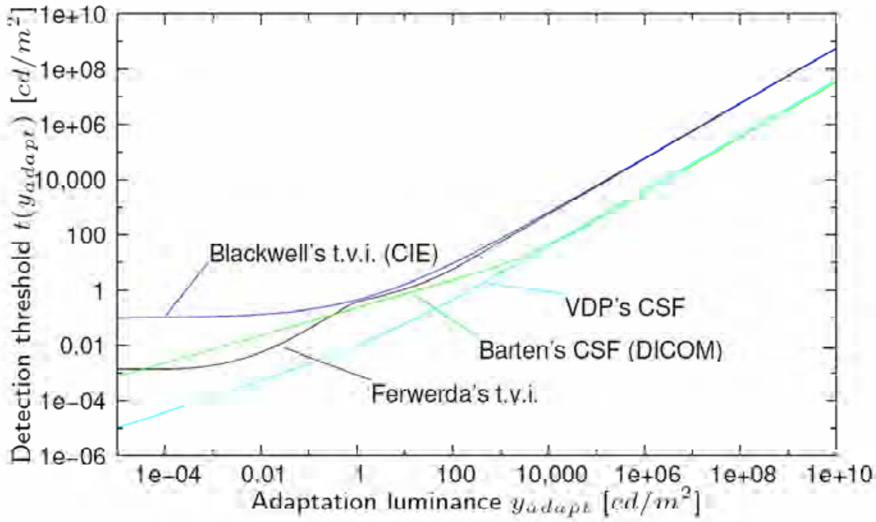

Luminance to Luma Mapping

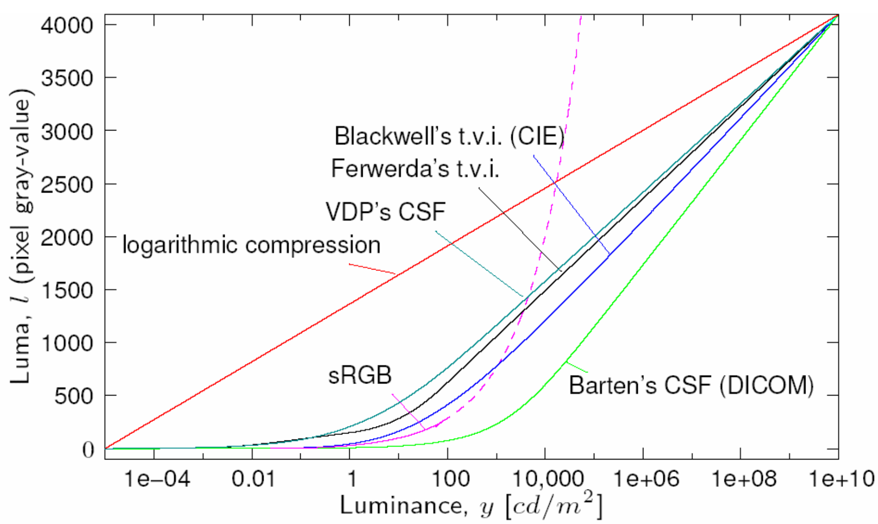

Conversion: Luminance → Luma

$$l(y) = \begin{cases} a \cdot y & \text{if } y < y_l \\ b \cdot y^c + d & \text{if } y_l \leq y < y_h \\ e \cdot \log(y) + f & \text{if } y \geq y_h \end{cases}$$

Model	a	b	c	d	e	f	y_l	y_h
CIE *t.v.i.*	17.554	826.81	0.10013	-884.17	209.16	-731.3	5.6046	10469
VDP's CSF	769.18	449.12	0.16999	-232.25	181.70	-90.16	0.061843	164.10

Conversion: Luma → Luminance

$$y(l) = \begin{cases} a' \cdot l & \text{if } l < l_l \\ b' \cdot (l + d')^{c'} & \text{if } l_l \leq l < l_h \\ e' \cdot \exp(f' \cdot l) & \text{if } l \geq l_h \end{cases}$$

Model	a'	b'	c'	d'	e'	f'	l_l	l_h
CIE *t.v.i.*	0.0570	7.3014e-30	9.987	884.17	32.994	0.00478	98.381	1204.7
VDP's CSF	0.0013	2.4969e-16	5.883	232.25	1.6425	0.00550	47.568	836.59

Chrominance and Chroma

- CIE 1976 Uniform Chromacity Scales (u' v')

$$u' = \frac{4X}{X + 15Y + 3Z} \qquad v' = \frac{9Y}{X + 15Y + 3Z}$$

$$u_{8bit} = u' \cdot 410 \qquad v_{8bit} = v' \cdot 410$$

- Just noticeable differences (JND) [Hunt 1995, p. 154]

$$\Delta u', \Delta v' \approx 0.002 \quad (\text{or } \Delta u_{8bit}, \Delta v_{8bit} \approx 0.82)$$

- 8-bit encoding gives sufficient precision for complex images
  - For smooth patterns we could see contouring artifacts for blue and purple colors for the highest luminance levels
- $L^*u^*v^*$ - not suitable for compression
  - too much correlation between luma and chroma

HDR Color Space: Summary

- Color space for HDR image and video encoding
  - 12-bit luma, 2 × 8-bit chroma
  - Derived from contrast detection data
  - Meets all requirements: full color gamut and luminance range; perceptually uniform; no countouring; positive integer numbers; channels decorrelated; conversion formulas from/to XYZ
- Simple extension for the existing formats

HDR Video Encoding [Mantiuk et al. Siggraph 2004]

- An extension to MPEG-4 (ISO/IEC 14496-2)
- $l\,u'v'$: 11 bits for luma, 2×8 bits for chroma

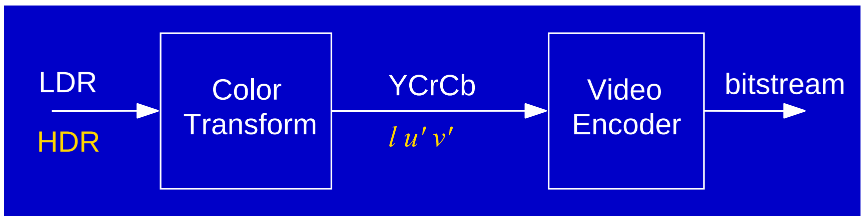

- White: MPEG
- Orange: HDR Encoder

l: Ferwerda's t.v.i. for luma encoding

→ Similar to Capacity function [Ashikhmin, EGSR 2002]

Luminance Quantization Comparison

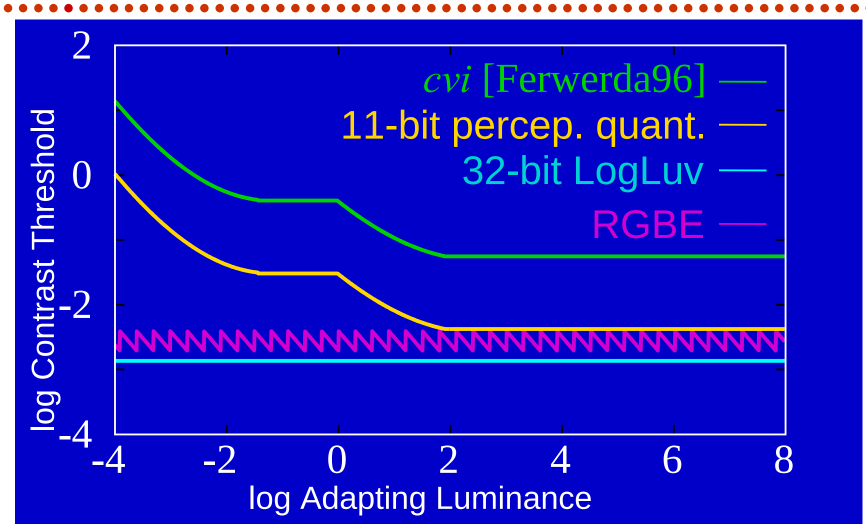

Edge Coding: Motivation

SIGGRAPH2006

- HDR video can contain sharp contrast edges
  - Light sources, shadows
- DCT coding of sharp contrast may cause high frequency artifacts

DCT coding

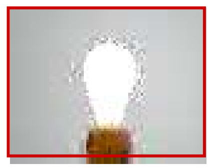

Edge coding

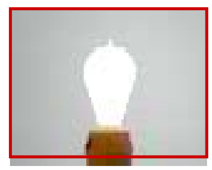

Edge Coding

SIGGRAPH2006

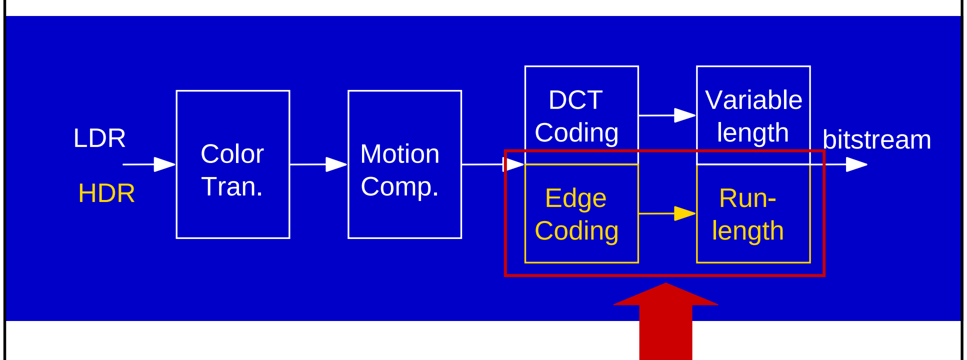

- White: MPEG
- Orange: HDR Encoder

Edge Coding: Solution

- Solution: Encode sharp edges in spatial domain, the rest in frequency domain

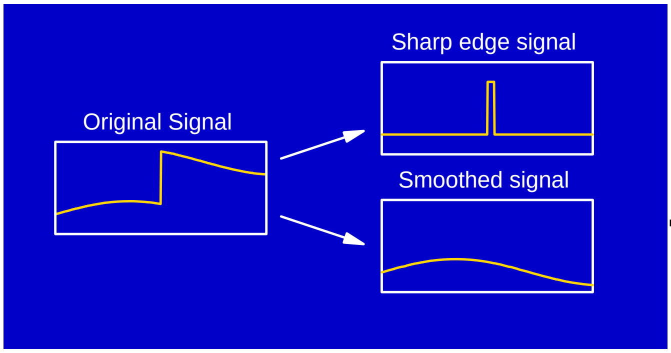

➔Run-length encoding

➔DCT encoding

Edge Coding: Algorithm

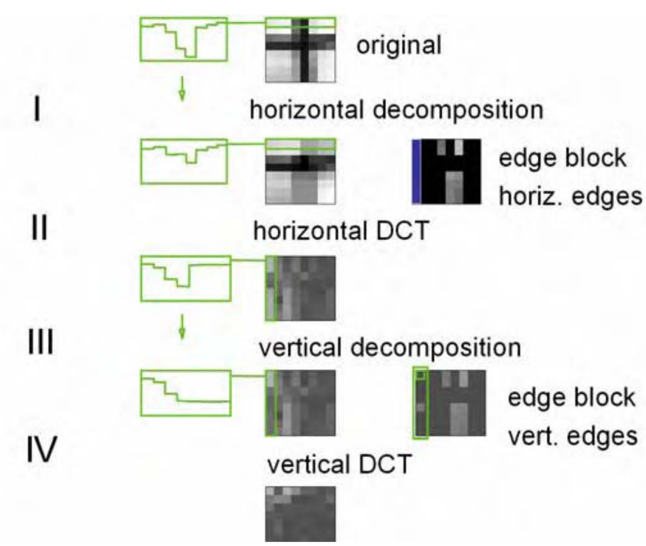

Backward Compatible HDR MPEG [Mantiuk et al. Siggraph 2006]

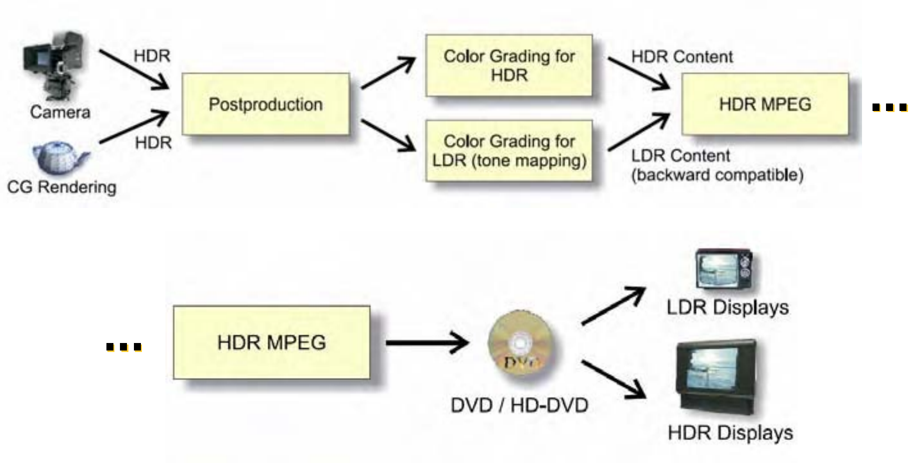

Backward Compatible HDR MPEG: Overview

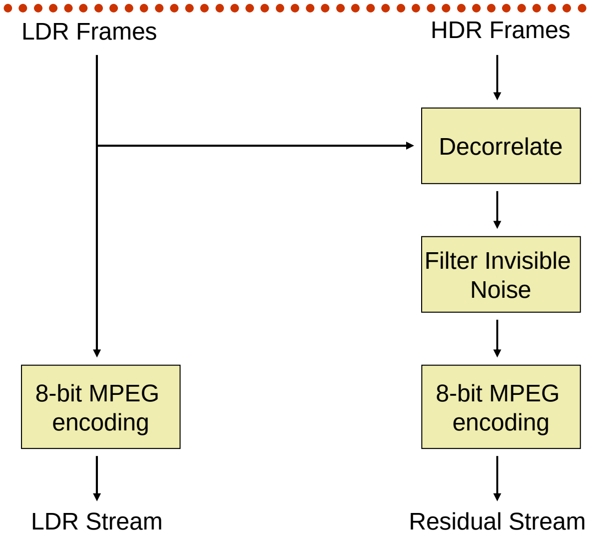

Backward Compatible HDR MPEG

SIGGRAPH2006

- Standard chipset can be used for 8-bit MPEG-4 decoding
- HDR and LDR video appearance not affected by the compression scheme
  - Arbitrary tone mapping can be used
- Backward compatibility of LDR stream with existing DVD players
- Modest size of residual stream
  - Perception-based HDR-enabled filtering of invisible details

12-bit (HDRV) vs. 8-bit Backward Compatible HDR MPEG

SIGGRAPH2006

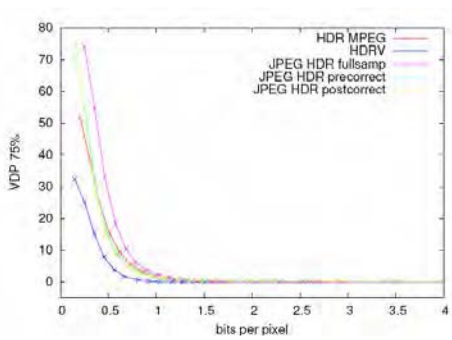

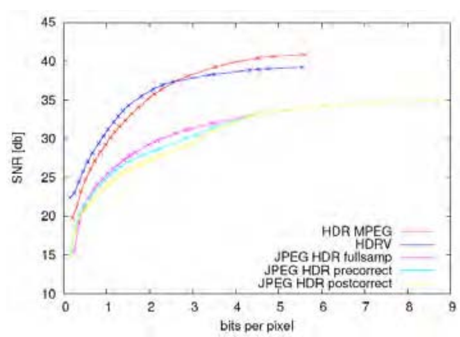

Real-Time HDR Playback with Perceptual Effects

http://www.mpi-inf.mpg.de/resources/hdr/peffects

HDR Video Player

- Tone mapping adjusted to display device
- Inspecting various luminance ranges with a linear luminance mapping
- Physically based post-processing effects: blooming, motion blur, night vision
  - Require HDR information
  - Computed on-the-fly by graphics hardware
  - Can be scripted through annotations in video stream

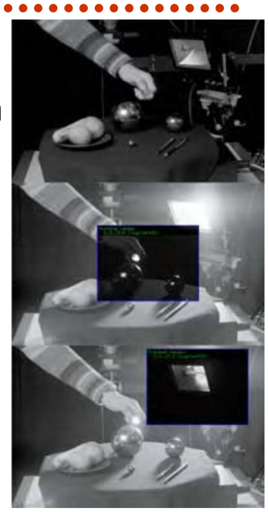

Real-time Tone Mapping for LDR/HDR Displays

SIGGRAPH2006

HDR Video Stream → HDR Video Player → LDR/HDR Display

tone mapping customized for a given display

Logarithmic mapping

Photographic (local) mapping

Night Vision

SIGGRAPH2006

- Human Vision operates in three distinct adaptation conditions:

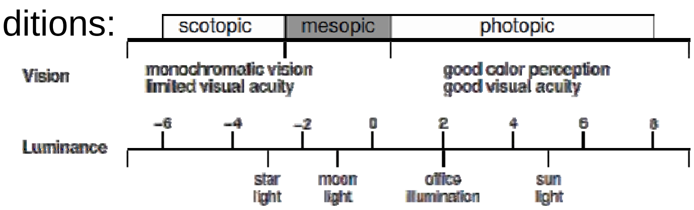

Temporal Luminance Adaptation

Light adaptation

- Compensates changes in illumination
- Simulated by smoothing adapting luminance in tone mapping equation
- Different speed of adaptation for photopic and scotopic vision

Tunnel entrance: dark adaptation

Visual Acuity

- Perception of spatial details is limited with decreasing luminance level (Y)
- Details can be removed using convolution with Gaussian Kernel
- Highest resolvable spatial frequency [Shaley 1937, Ward 1997]:

$$RF(Y) = 17.25 \cdot \arctan(1.4 \log_{10} Y + 0.35) + 25.72$$

Veiling Luminance (Glare)

- Decrease of contrast and visibility due to light scattering in the optical system of the eye
- Described by the optical transfer function [Deeley 1991]:

$$OTF(\rho, d(\bar{Y})) = \exp\left(-\frac{\rho}{20.9 - 2.1 \cdot d}^{1.3-0.07 \cdot d}\right)$$

ρ spatial frequency, d pupil aperture

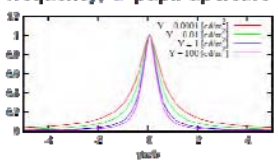

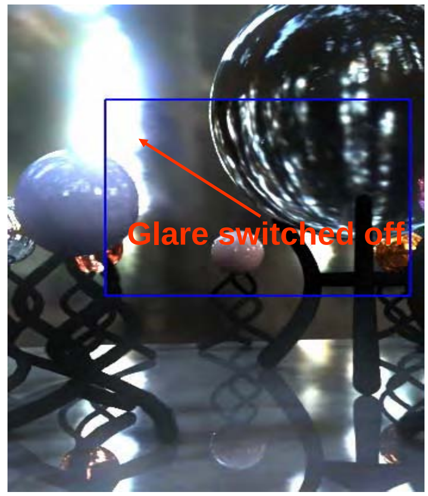

Motion Blur

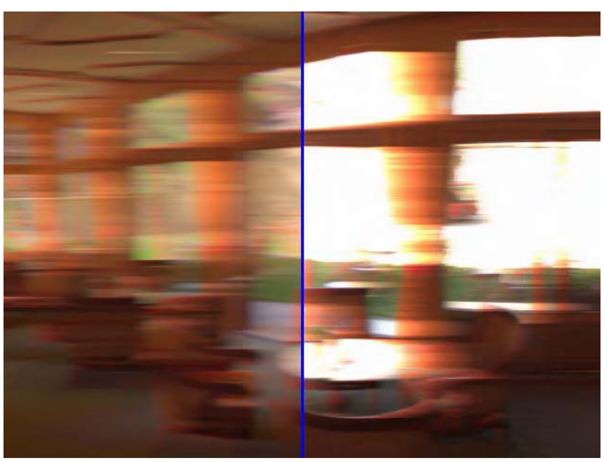

LDR HDR

HDR
Visible Difference Predictor

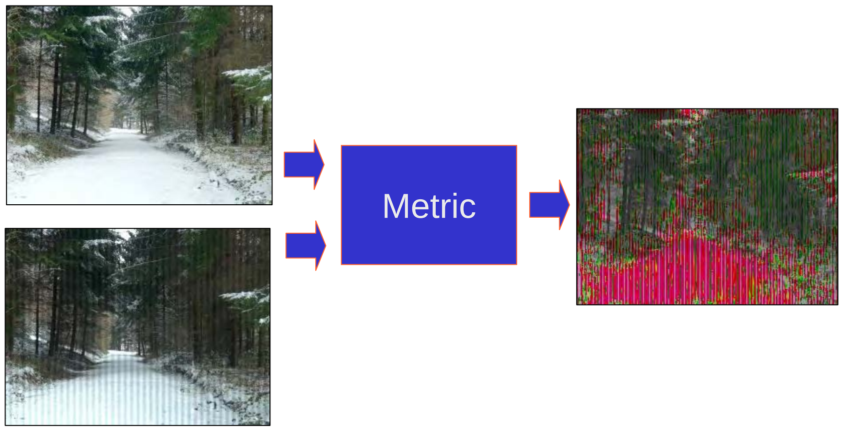

http://www.mpi-inf.mpg.de/resources/hdr/vdp

HDR VDP Overview

- Relevant characteristics of the Human Visual System
- LDR VDP
- HDR VDP extensions

Subjective vs. Objective Methods

SIGGRAPH2006

- Subjective methods involving human subjects
  - Very costly
- Simple objective metrics e.g. RMS Error
  - Not reliable
- Basic characteristics of the Human Visual System (HVS) must be modeled to improve difference prediction reliability:
  - + Luminance adaptation
  - + Contrast sensitivity
  - + Visual masking

Photoreceptor Response

SIGGRAPH2006

$$\text{Michelis} - \text{Menten Equation}: R(y) = \frac{y^n}{y^n + \sigma^n_{y_{adapt}}} R_{\max} \qquad n \approx 0.74$$

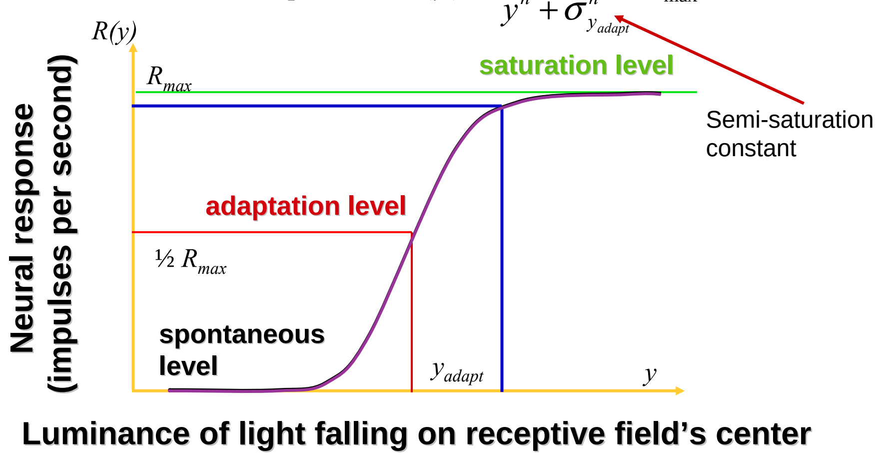

Photoreceptor Response

$$\text{Michelis} - \text{Menten Equation}: R(y) = \frac{y^n}{y^n + \sigma^n_{y_{adapt}}} R_{\max} \qquad n \approx 0.74$$

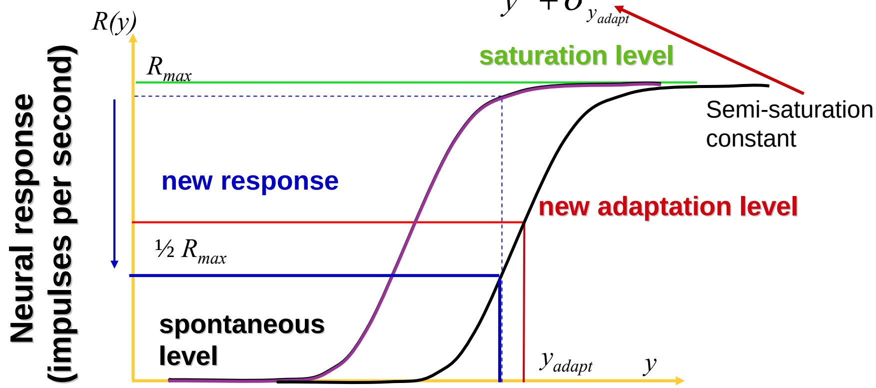

Luminance of light falling on receptive field's center

Contrast Sensitivity Function

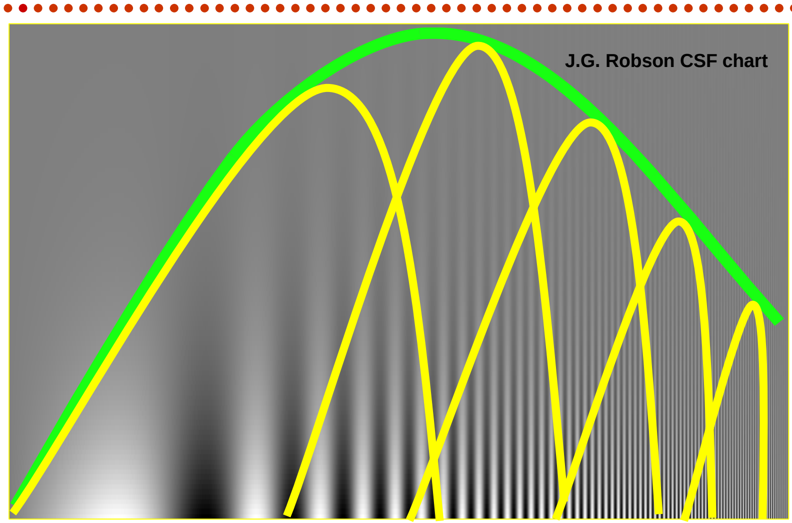

Visual Masking

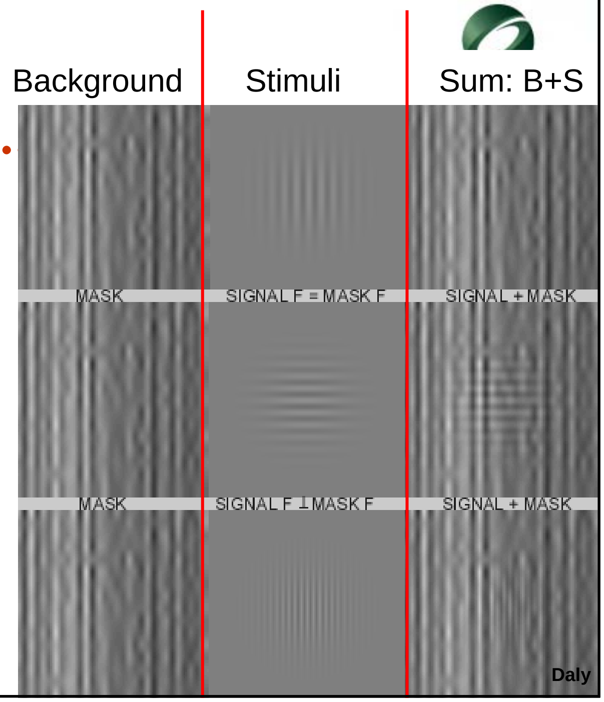

- Strong masking: similar spatial frequencies
- Weak masking: different orientations
- Weak masking: different spatial frequencies

Typical HVS Model

SIGGRAPH2006

Detection of perceivable differences between images depends on the following characteristics of the HVS:

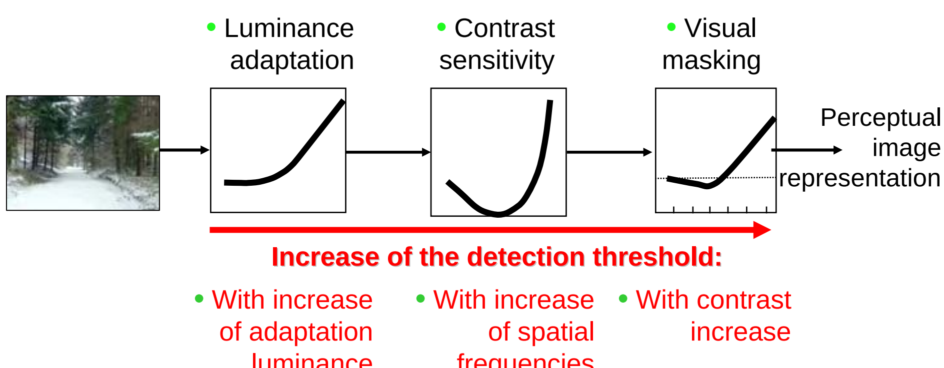

Pattanaik

Perceivable Differences Predictor [Daly SPIE 1992]

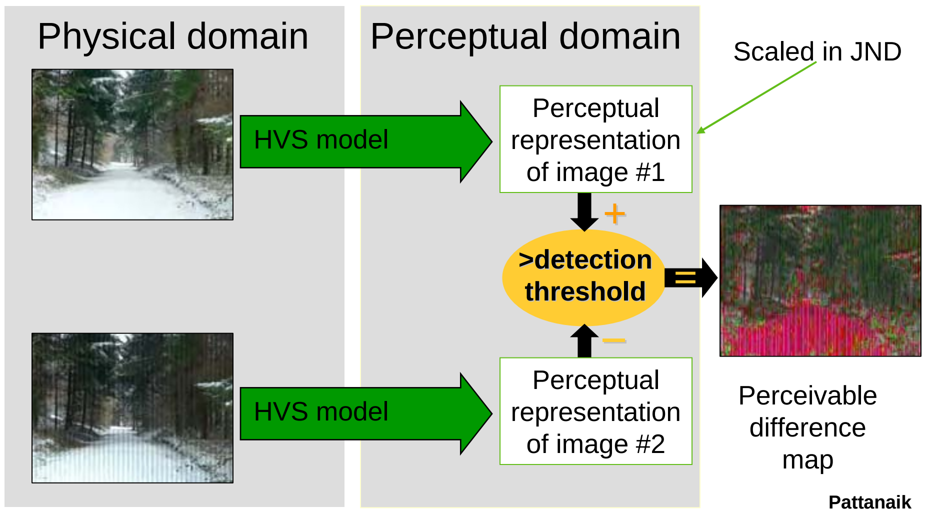

Pattanaik

HDR VDP – Extensions [Mantiuk et al. SPIE 2005]

- Optical Transfer Function
  - Glare effect important for high contrasts
- JND scaled contrast
  - Contrast measure that does not depend on adaptation luminance
- Locally adaptive CSF
  - Instead of global adaptation for the whole image

➢ <u>Assumption:</u> The eye can adapt to luminance of very small patches (to a single pixel)

Optical Transfer Function

- Modulation Transfer function that simulates optics of the human eye [Deeley et al. 1991]

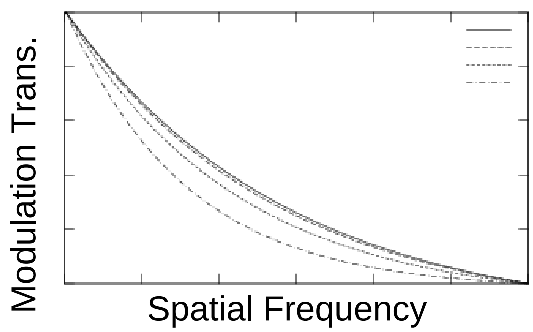

Optical Transfer Function

- OTF significantly reduces contrast in HDR scenes

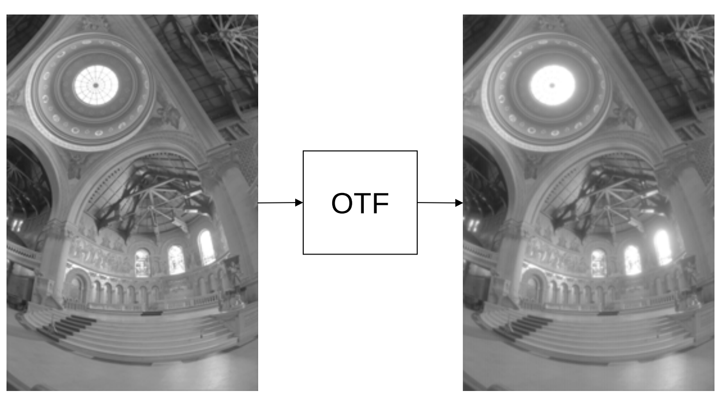

Amplitude Nonlinearity

- Introduced contrast measure independent of adaptation luminance
- Scaled in JND units
  - ➔ our HDR luma space *l*

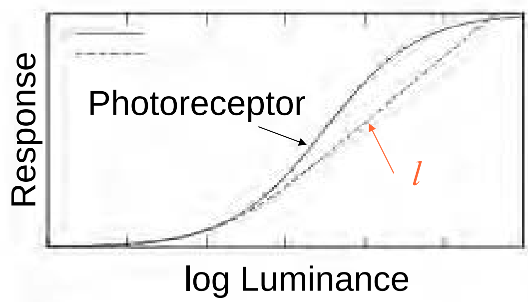

Local Adaptation for CSF

- The shape of CSF varies with adaptation luminance

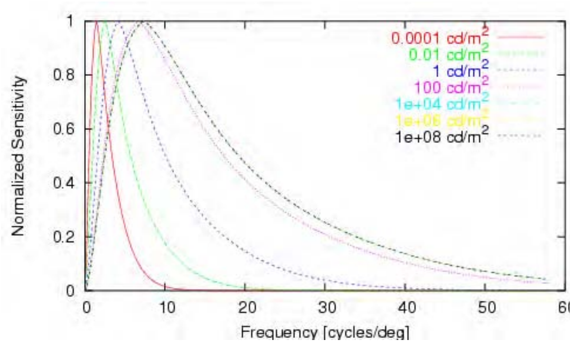

- Adaptation luminance can vary significantly across an image

Local Adaptation for CSF

1. Use different CSF for filtering each part of an image
2. Interpolate filtered images depending on adaptation luminance [Durand03]

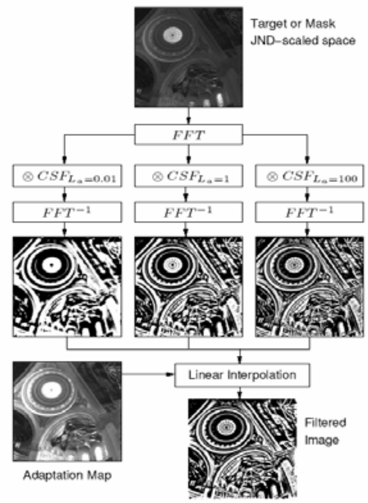

Calibration

- An experiment conducted on an HDR display to find subjective difference probability map
- HDR VDP parameters optimized to fit closely subjective data

Calibration: Experiment

SIGGRAPH2006

- Subjects were to mark visible differences using rectangular blocks

- Results averaged across subjects

  ➔fuzzy detection probability map

Calibration: Data Fitting

SIGGRAPH2006

- HDR VDP response converted to format of the subjective data

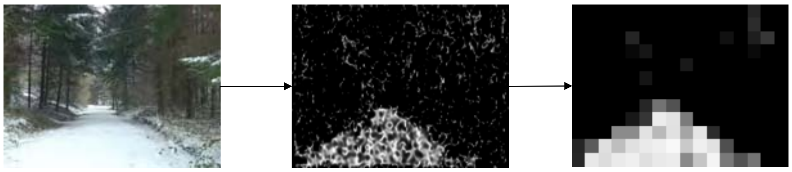

Distorted Image   VDP Response   Integrated Resp.

- Found the best fit for peak threshold contrast and masking function slope

HDR VDP: Summary

- Objective visual difference metric
  - Not limited to existing display technology
  - Predicts changing sensitivity in bright and dark regions of an image
  - Small performance overhead
- Applications
  - HDR JPEG and MPEG compression
  - New display technology (HDR display)
  - Assessment of visibility for varying luminance

HDR Video Environment Maps

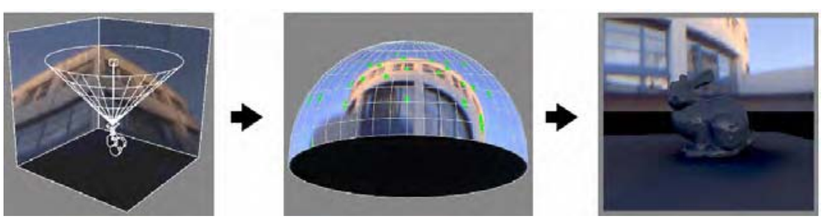

HDR Illumination Acquisition — Light Sources Computation — Rendering (GPU) (+Compositing)

http://www.mpi-inf.mpg.de/resources/hdr/vem

Illumination Acquisition

SIGGRAPH2006

- Source of data: HDR camera
  - 12-bit logarithmic response
  - 8 orders of magnitude dynamic range
- Fish-eye lens
- Resolution 640x480
- Photometric calibration necessary

Video Environment Map (VEM) [Havran et al. EGSR 2005]

SIGGRAPH2006

- HDR image on hemisphere defined by meridians and parallels
- Extension to time domain
- Luminance corresponds to the probability density function on the hemisphere (grid points)

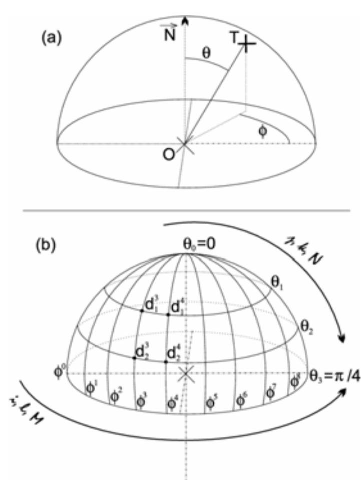

VEM Decomposition into Directional Light Sources

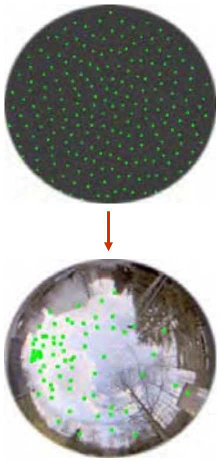

Importance Sampling Properties

- Arbitrary number of directional light sources
- The same energy for each light source
- Progressive light source sequence
- Small memory requirements
- Real-time performance
- Dependence on the surface normal
- Energy and position of light sources processed with FIR filters

Rendering System

- Full pipeline from HDR acquisition to rendering
- Real-time computation of light sources with many properties useful for rendering
- Performance 5-10 FPS for 320×240 pixels on NVidia GeForce 6800GT
- Principles also useful for CPU rendering systems
- Dynamic geometry supported
- No interreflections

Snapshots & Video

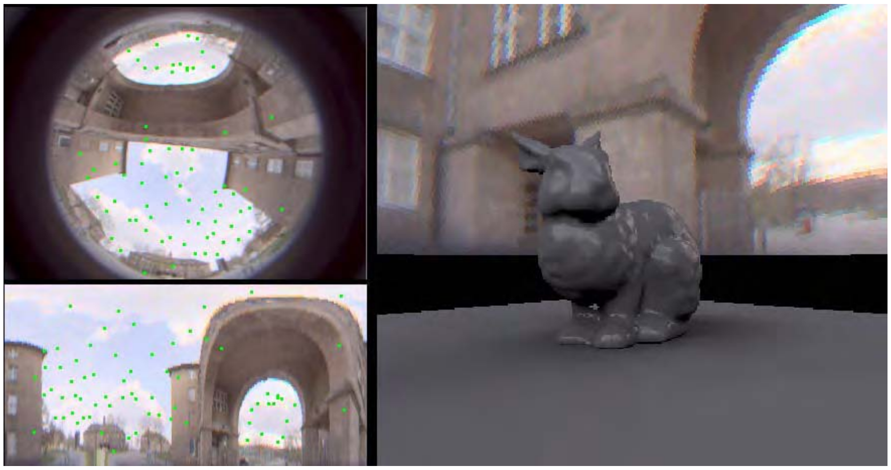

Mixed Reality Applications

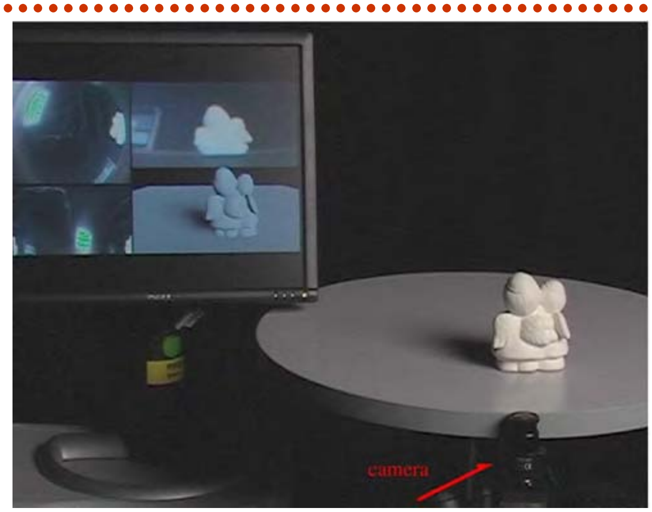

Special Case: Dynamic Lighting and Static Geometry

- CAVE system for car interior design and ergonomic studies
- Real-world driving scenarios
  – Tunnels, sunset, city driving, highways, etc
- Two HDR video cameras with fisheye lenses that capture:
  – Environment map
  – Windshield image

Precomputed Radiance Transfer (PRT) Rendering

- Car model contains approx. 500K triangles
- Visualization frame rate ~10fps

Acknowledgements

- We would like to thank Tom Annen, Scott Daly, Kirill Dmitriev, and Vlastimil Havran for their respective contributions to the presented work.
- We are grateful to Helge Seetzen and Greg Ward for their friendly advice and BrightSide Technologies for making the whole family of their HDR displays available to us.
- We are also grateful to Bernd Hoefflinger and IMS CHIPS for providing us with HDR cameras.

SIGGRAPH2006

Tone Reproduction

Erik Reinhard

Tone reproduction

Match range of image intensities to those of the display device.

Possible approaches:

- Linear scaling
- Image formation
- Human visual system

Linear scaling

Linear

Linear + cropping

Image formation

Assumption: pixel intensities are the result of light reflecting off surfaces and hitting an image sensor.

A much simplified model: $\boxed{L = i * r}$

- L: pixel intensity
- i: illuminance
- r: surface reflectance

Surface reflectance

Surfaces reflect a fraction of between 0.005 and 1.0 of incident light (Stockham 1972)

Surface reflectance does not by itself create high dynamic range images.

Illuminance

The amount of light incident upon a surface could have any range of values (as long as these values are positive)

$L = i * r$: If L has a high dynamic range, then this is because i has a high dynamic range

Intensity vs. Density

In photography images are often represented as densities, i.e. in the logarithmic domain:

$$D = \log(L) = \log(i) + \log(r)$$

Note that D could be negative

The notation given here will be used throughout this section of the course. D is a 'density image'

Frequency domain analysis

One more observation about reflectance and illuminance:

- Reflectance tends to exhibit high spatial frequencies
- Illuminance is slowly varying over surfaces

Oppenheim 1968

Given a set of pixel intensities, we wish to attenuate the illuminance component while keeping the reflectance component

Solution: Take the Fourier transform of a density image and attenuate the low frequencies. Then exponentiate to form a displayable image

Oppenheim 1968

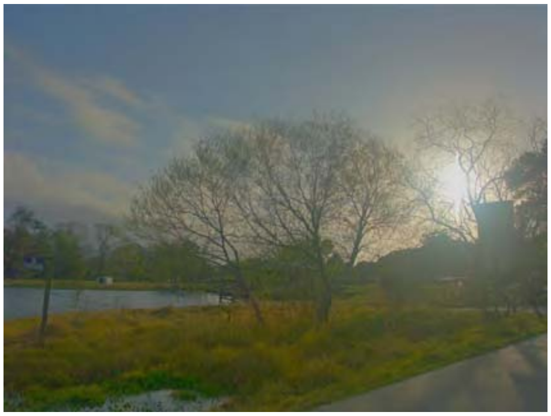

Durand and Dorsey 2002

Split image into base- and detail layers:

- Use edge-preserving smoothing operator to filter density image and call result ‘base layer’
- Divide density image by base layer and call result ‘detail layer’

Durand and Dorsey 2002

- Note that base layer is high dynamic range and could be viewed as the illuminance component
- The detail layer is low dynamic range and could be viewed as the reflectance component

Durand and Dorsey 2002

Input image

Smoothed result

Durand and Dorsey 2002

Related techniques

- Choudhury and Tumblin (2003): Trilateral filter
  - A refinement of the edge-preserving smoothing operator used by Durand and Dorsey
- Pattanaik and Yee (2002):

  Gain control operator
  - Different kernel shapes

Alternatively…

SIGGRAPH2006

Illuminance is slowly varying over a surface, whereas reflectance produces sharp discontinuities

In other words: large image gradients are caused by reflectance edges

Horn 1974

SIGGRAPH2006

Model of lightness (i.e. a model of the perception of surface reflectance):

- Differentiate a density image
- Threshold density gradients (keep large gradients)
- Integrate image (numerical process)
- Exponentiate

Horn 1974

Original Processed

Image formation

This model of image formation breaks down

- if specular highlights are visible in the image
- if light sources are directly visible
- if image depicts fluorescent materials

In fact, we might want to attenuate large gradients rather than preserve them!

Fattal et al 2002

Gradient domain compression:

- Differentiate density image
- Attenuate large gradients
- Integrate
- Exponentiate

Fattal et al 2002

Human visual system

Sumanta already explained the principles by which the HVS deals with high dynamic range scenery

Here, we build on those insights to classify HVS-based tone reproduction operators

Global vs. Local

We may apply compressive functions directly in image space

- Global: use the same function for each image pixel (use a global adaptation value)
- Local: modulate compressive function by pixel neighborhood (use a local adaptation value)

Global operators

Since each pixel is treated independently:

- Efficient and fast
- Suitable for medium dynamic range imagery

Global operators

Only a small number of different functional shapes are known:

- Linear scaling
- Logarithmic scaling
- Histogram-based compression
- Sigmoidal functions (discussed by Sumant)

Linear scaling

SIGGRAPH2006

- Already shown an ad-hoc linear scaling function at the start of this section
- Ward 1994 and Ferwerda et al 1996: Linear scaling based on TVI functions
  - Bring the log average luminance to a sensible display level
  - Better than ad-hoc scaling, not suitable for very high dynamic range images

Ward 1994 (linear scaling)

SIGGRAPH2006

Ferwerda et al 1996

SIGGRAPH2006

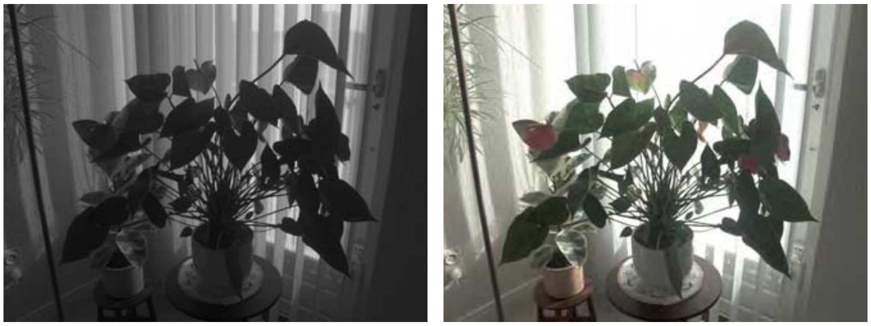

Applied different prescaling factors

Logarithmic scaling

SIGGRAPH2006

Simple version:

$$L_d = \frac{\log(L_w + 1)}{\log(L_{w,\max} + 1)}$$

Drago et al 2003:

$$L_d = \frac{\log_b(L_w + 1)}{\log_b(L_{w,\max} + 1)}$$

b is a function of the pixel intensity and is modulated by a user parameter to steer overall appearance

Logarithmic compression

Drago et al 2003

Histogram adjustment

Ward et al (1997) use the shape of the image's histogram

- Compute histogram
- Compute cumulate histogram
- Result is a monotonically increasing function
- Reshape this function to avoid slopes greater than 1
- Remap luminances according to this function

Ward et al 1997

Sigmoids

S-shaped compression function used in a large number of tone reproduction operators.

Discussed in detail by Sumanta in previous section. Will omit discussion here.

Global operators

Recap:

- Easy to implement
- Fast
- Some provide very reasonable amounts of compression and plausible results are frequently obtained
- Useful for medium to high dynamic range images

Local operators

SIGGRAPH2006

Sometimes the mismatch between HDR image and display device is very large

More compression afforded by local operators

Often at the cost of artifacts

Local operators

SIGGRAPH2006

Two approaches:

- $L_d(x, y) = s(x, y) L_w(x, y)$
- Replace global adaptation value with a per-pixel local adaptation value

Local operators

- Chiu et al, Rahman et al, iCAM model:

$$L_d(x,y) = s(x,y)L_w(x,y)$$
$$s(x,y) = f\left(\frac{1}{L_w^{\mathrm{LPF}}(x,y)}\right)$$

i.e. divide the image by a low-pass filtered version of itself

Chiu et al 1993

Rahman et al

Johnson and Fairchild (iCAM)

Local operators

When dividing by a low-pass filtered image,

we need two things to minimize artifacts:

– A very large kernel size
– Reduce the weight of the LPF image

Global to local

- Replace global average with local average within global operators
- Sigmoids! General shape:

$$L_d(x,y) = \frac{L_w(x,y)^n}{L_w^{\mathrm{LPF}}(x,y)^n + L_w(x,y)^n}$$

Local adaptation

There is a catch, though:

- The size of the low-pass filter kernel is important.
- Too large and haloing artifacts will occur
- Too small and compression will not be better than global operators

Low-pass filter kernel size

- For each pixel the size of the LPF kernel should be such that it does not overlap with sharp discontinuities
- But we still want to average over the largest spatial area for which the above is true (which may be different for each pixel)
- In practice often a small kernel size!

How to compute?

SIGGRAPH2006

- Multi-scale techniques using difference of Gaussians and scale selection (Reinhard 2002, Ashikhmin 2002)
- Edge-preserving smoothing operator (bilateral filter, mean shift algorithm, LCIS…)

Reinhard et al 2002

SIGGRAPH2006

Ashikhmin 2002

Global vs. local

Sigmoid with global adaptation

Sigmoid with local adaptation

Informal comparison

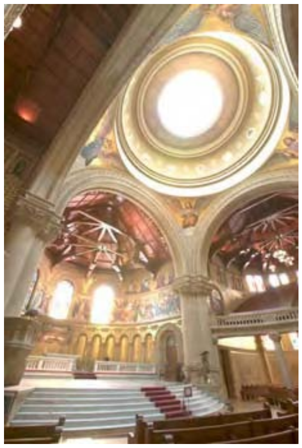

Linear scaling

Linear with clamping

Informal comparison

Bilateral filter

Gradient compression

Informal comparison

SIGGRAPH2006

Histogram adjustment

Gradient compression

Informal comparison

SIGGRAPH2006

Photographic operator

Photoreceptor-based

Informal comparison

Li et al, SIGGRAPH 2005

Conclusions

Only a few fundamentally different approaches to tone reproduction exist

- Based on image formation
  - Frequency domain
  - Gradient domain
- Based on the human visual system
  - Global operators
  - Local operators

Conclusions

SIGGRAPH2006

Trade-offs exist between:

- Amount of compression
- Presence of artifacts
- Computation time

Validation studies are an important tool to figure out which operator is most suited to which task

SIGGRAPH2006

HDR Image-Based Lighting

Paul Debevec

Image-Based Lighting

Paul Debevec
University of Southern California
Institute for Creative Technologies
Graphics Laboratory

SIGGRAPH 2005 Course #29
High-Dynamic-Range Imaging and Image-Based Lighting
Half-Day, Tuesday, 2 August, 8:30 am - 12:15 pm

www.debevec.org

"A McIntosh in the Kitchen" by Marc Jacquier (2004)

Image courtesy of Nick Bertke

Image courtesy of Nick Bertke


Nick Bertke
pogo767@gmail.com


Canon Powershot G2
f/8, 1 sec


Canon Powershot G2
f/8, 1/2 sec


Canon Powershot G2
f/8, 1/5 sec


Canon Powershot G2
f / 8, 1 / 13 sec


Canon Powershot G2
f / 8, 1 / 30 sec


Canon Powershot G2
f/8, 1/80 sec


Canon Powershot G2
f/8, 1/200 sec


Nick Bertke


Nick Bertke
pogo767@gmail.com


Nick Bertke
pogo767@gmail.com

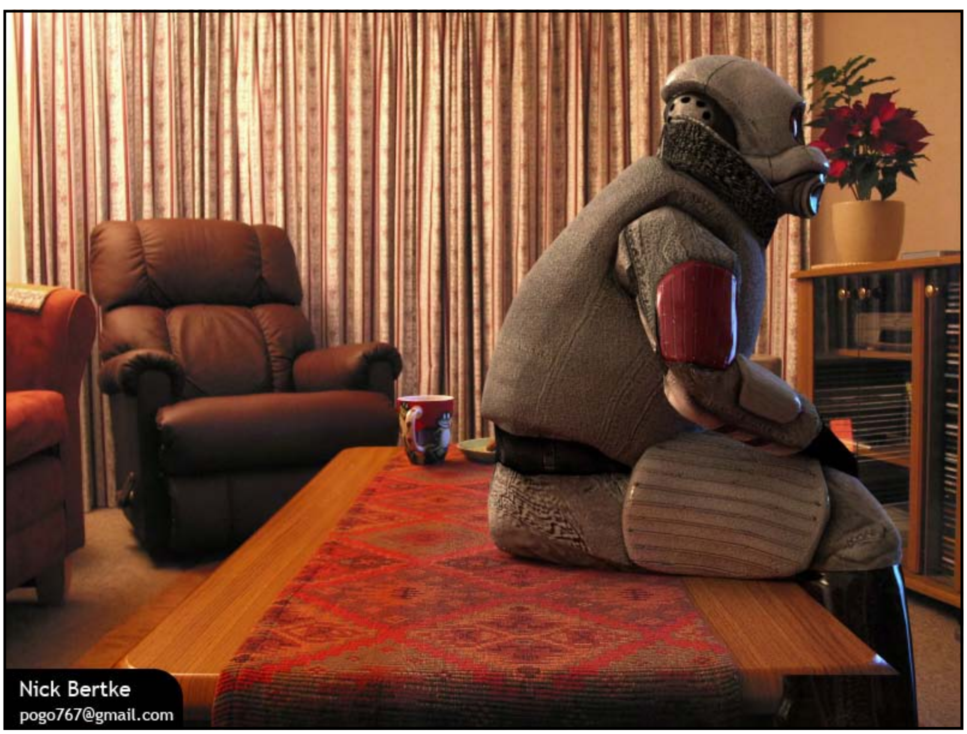
Nick Bertke
pogo767@gmail.com


Rendering with Natural Light
SIGGRAPH 98 Electronic Theater

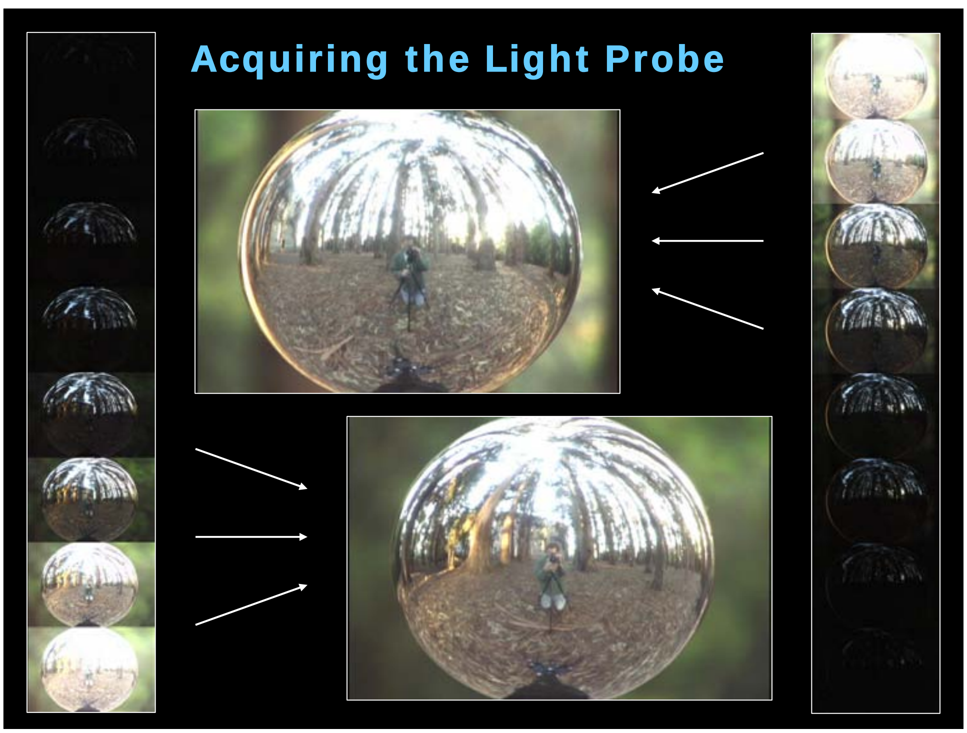
Acquiring the Light Probe

Assembling the Light Probe

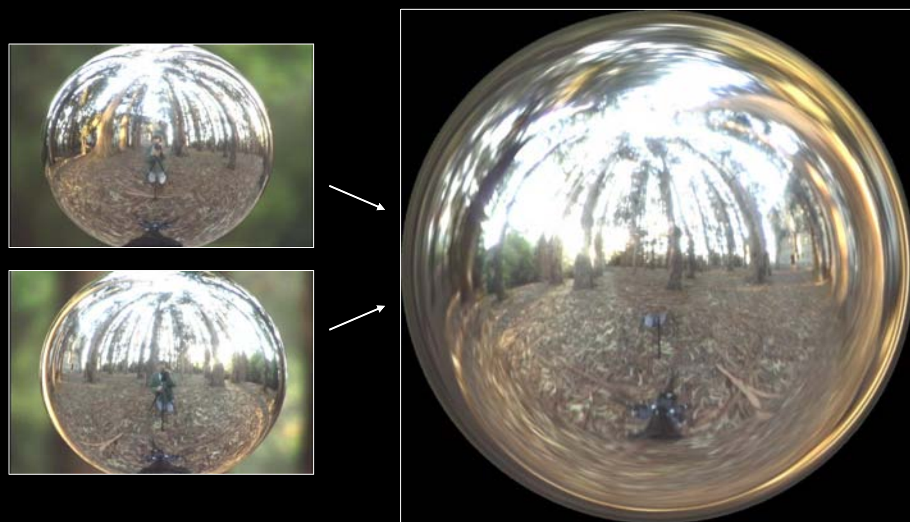

See HDRShop Tutorial #5, www.hdrshop.com

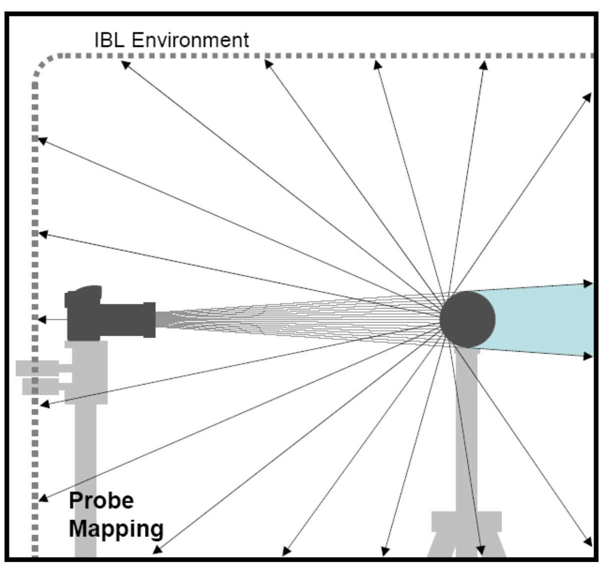


RNL Probe mapped onto interior of large cube

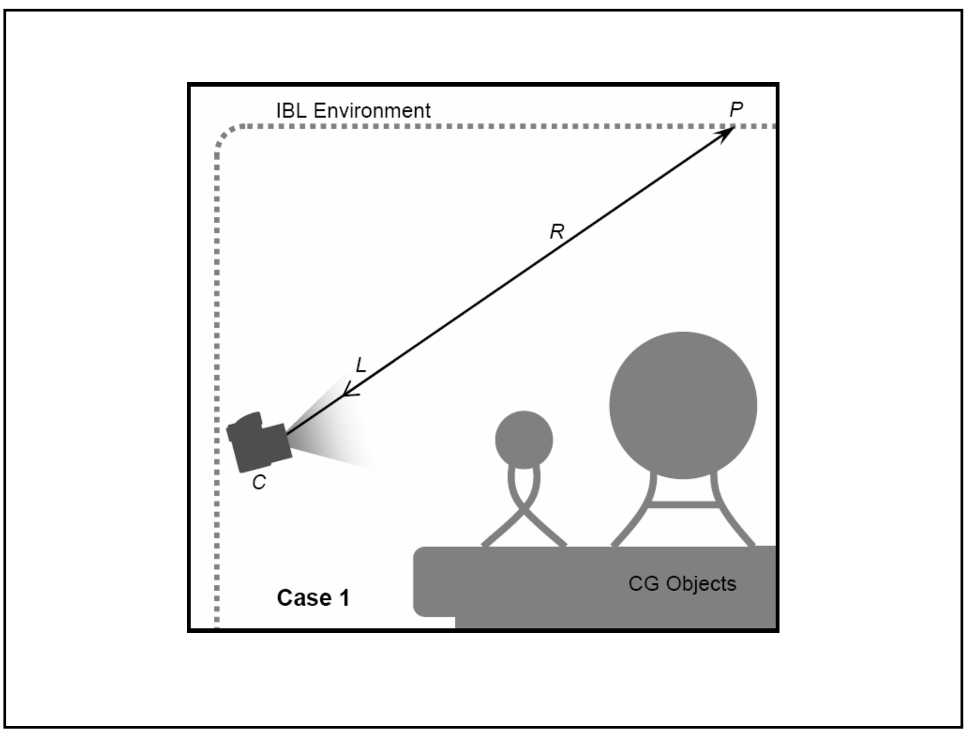
IBL Environment
P
R
L
C
Case 1
CG Objects

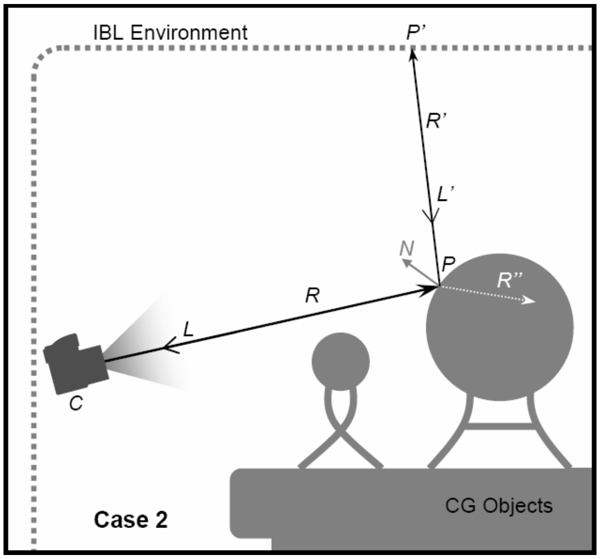
IBL Environment
P'
R'
L'
N
P
R"
R
L
C
Case 2
CG Objects

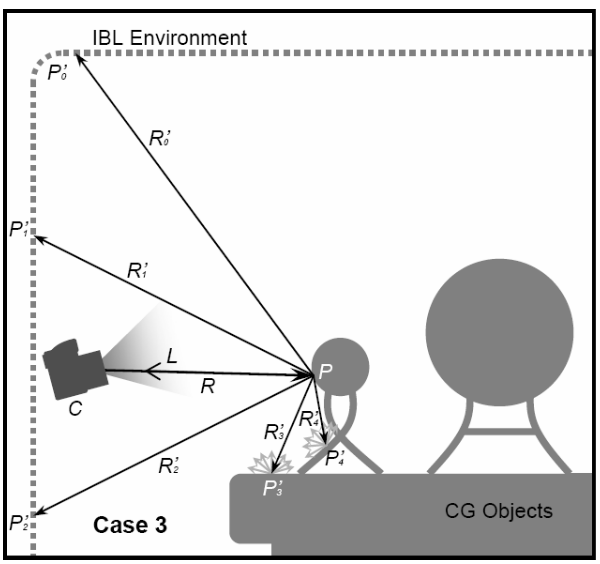
IBL Environment
P'_0
R'_0
P'_1
R'_1
L
R
P
C
R'_3
R'_4
P'_4
R'_2
P'_3
P'_2
Case 3
CG Objects

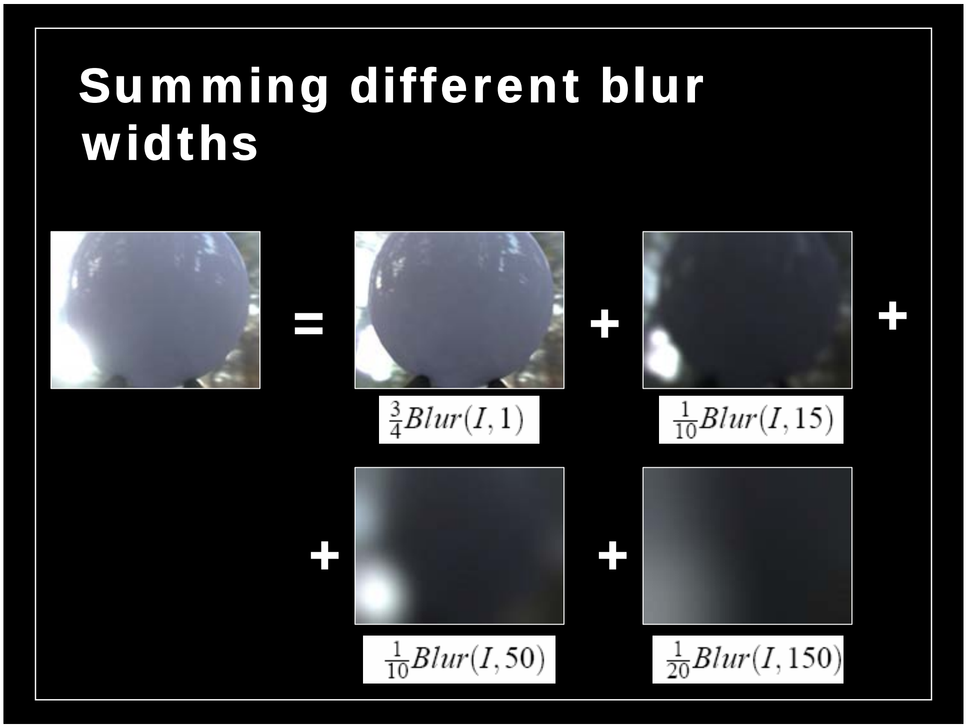
Summing different blur widths
=
+
+
+
+
$\frac{3}{4}Blur(I,1)$
$\frac{1}{10}Blur(I,15)$
$\frac{1}{10}Blur(I,50)$
$\frac{1}{20}Blur(I,150)$

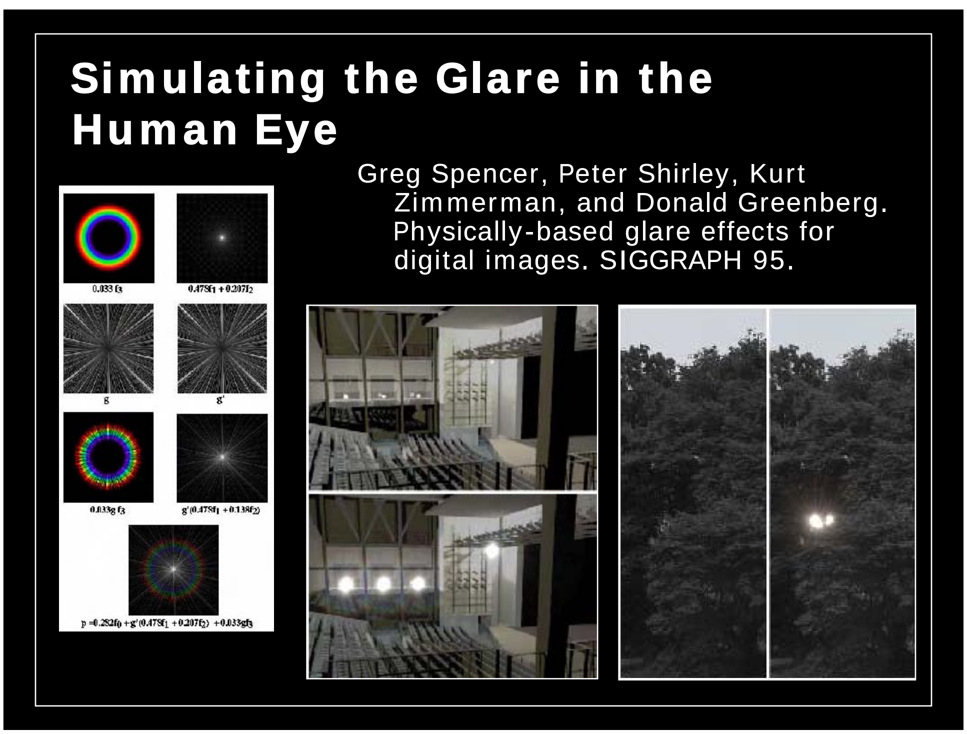
Simulating the Glare in the Human Eye
Greg Spencer, Peter Shirley, Kurt Zimmerman, and Donald Greenberg. Physically-based glare effects for digital images. SIGGRAPH 95.

The Sampling Problem

HDRI Sky Probe from http://www.ict.usc.edu/graphics/skyprobes/

Lit by sun and sky

Rendered in Marcos Fajardo's "Arnold" renderer

9 samples per pixel, 17 min.

Rendered in Marcos Fajardo's "Arnold" renderer

9 samples per pixel, 17 min.

Rendered in Marcos Fajardo's "Arnold" renderer

16 samples per pixel, 46 min.

Rendered in Marcos Fajardo's "Arnold" renderer

100 samples per pixel, 189 min.

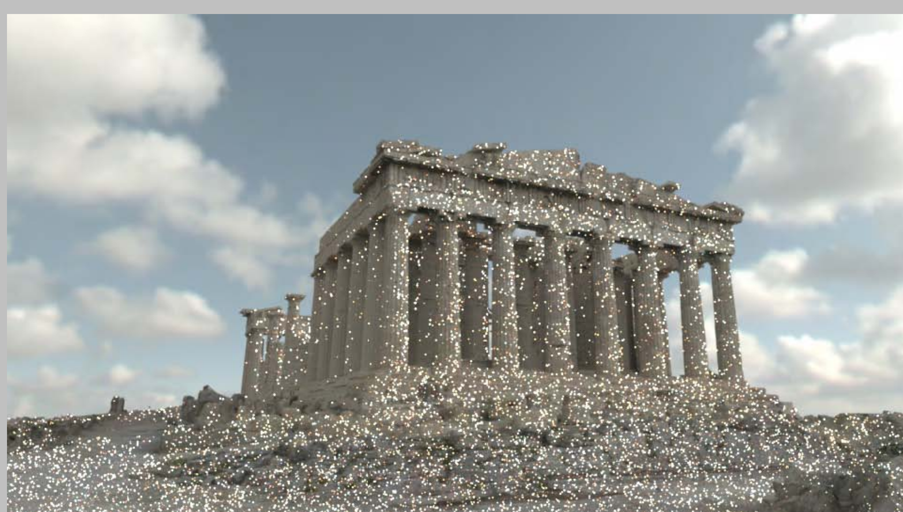

Rendered in Marcos Fajardo's "Arnold" renderer

A sunlit sample point

Rendered in Marcos Fajardo's "Arnold" renderer

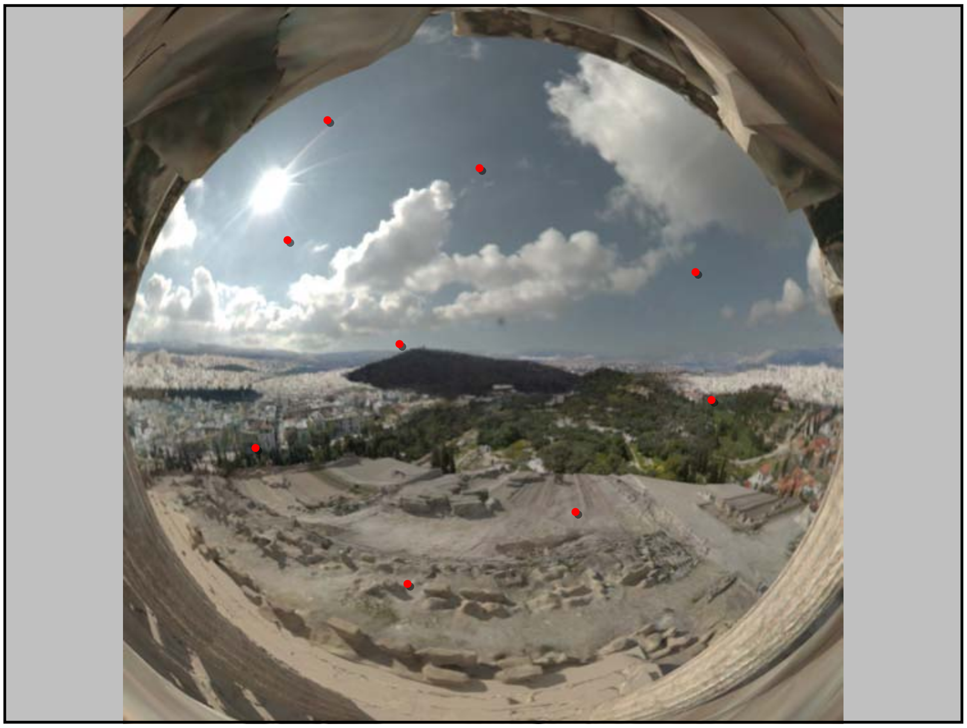


A shadowed sample point

HDRI Sky Probe

Clipped Sky + Sun Source

+

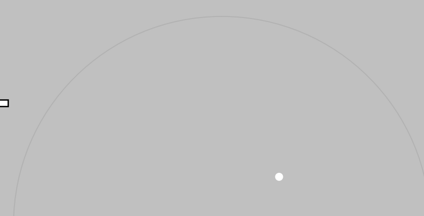

```
texture_map _HDRI_probe_map_clip_
"probe_09-55_clipped.hdr" {
swrap periodic
 filter bilinear
}

shader _HDRI_clip sky_HDRI {
 HDRI_map "_HDRI_probe_map_clip_"
 multiplier 1.000000
}
```

```
light sun directional 1 {
 direction -0.711743 -0.580805 -0.395078
 angle 0.532300
 color 10960000 10280000 866000
}
```

Lit by sky only, 17 min.

Lit by sun only, 21 min.


Lit by sun and sky, 25 mi

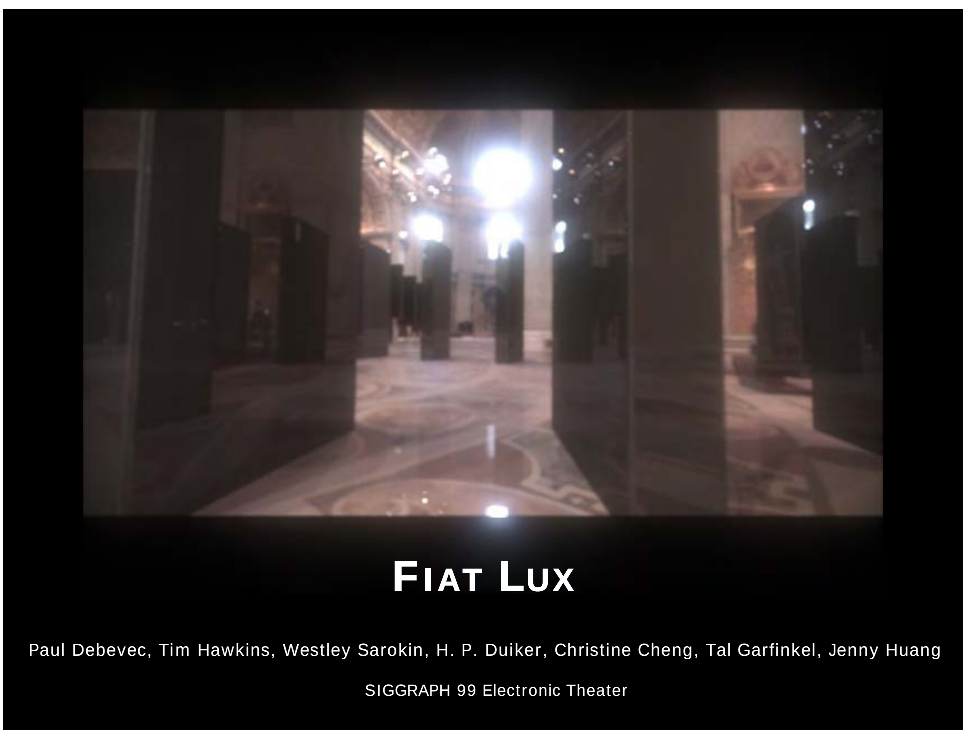
FIAT LUX
Paul Debevec, Tim Hawkins, Westley Sarokin, H. P. Duiker, Christine Cheng, Tal Garfinkel, Jenny Huang
SIGGRAPH 99 Electronic Theater

Assembled Panorama

Identified Light Sources

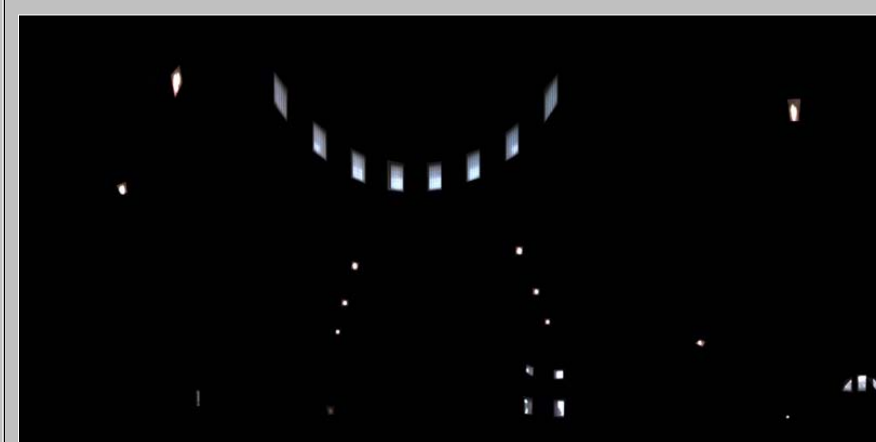

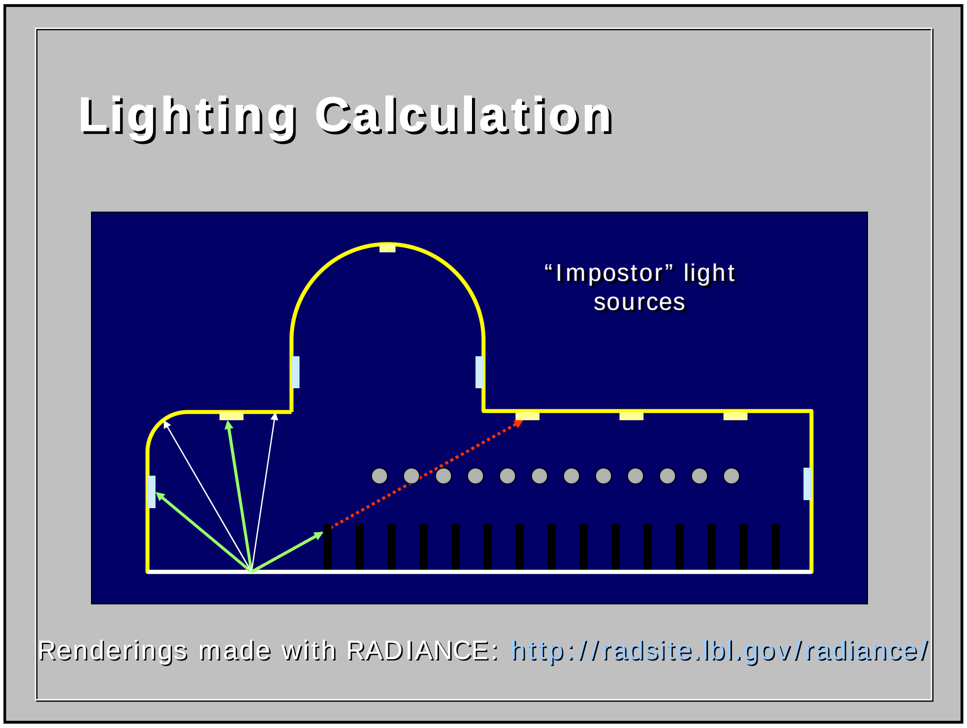
Lighting Calculation
"Impostor" light
sources
Renderings made with RADIANCE: http://radsite.lbl.gov/radiance/


Rendering Light Probes as
Light Sources
1999

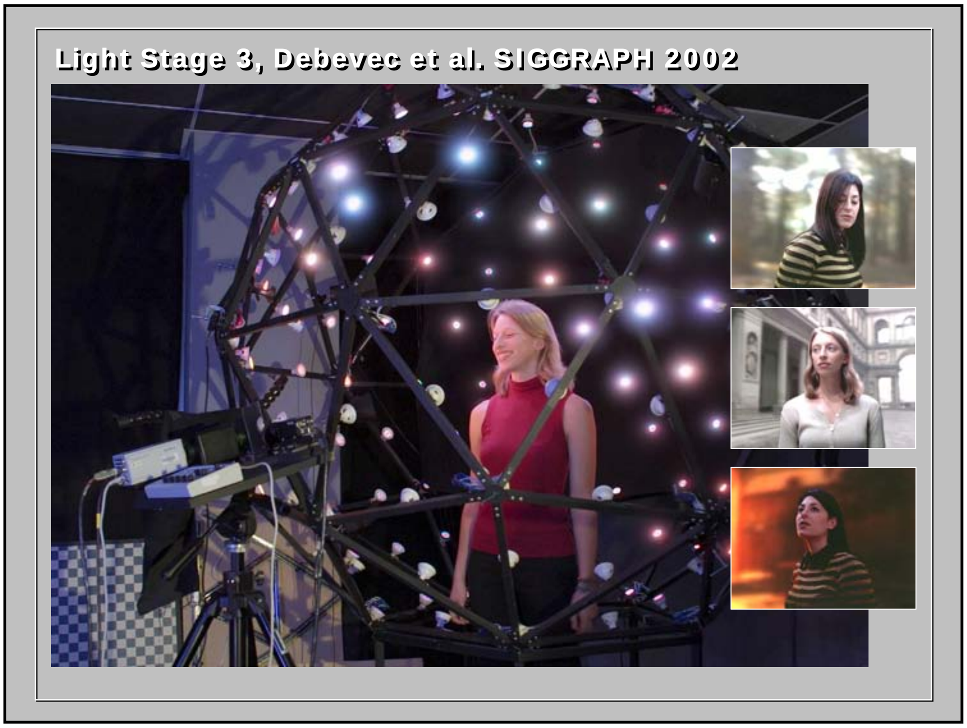
Light Stage 3, Debevec et al. SIGGRAPH 2002

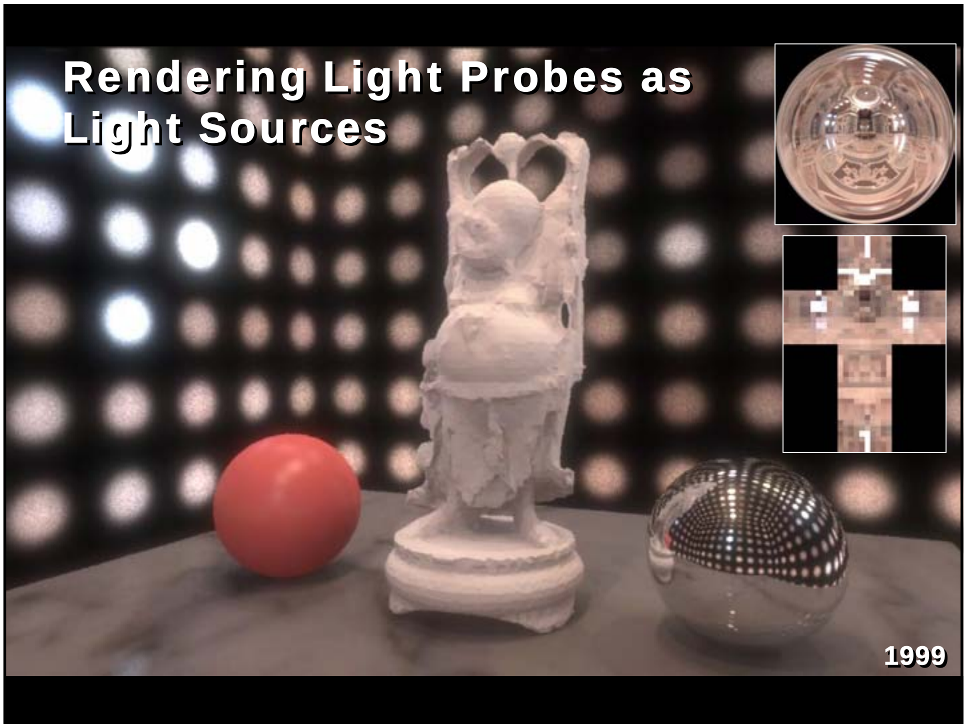
Rendering Light Probes as Light Sources
1999

Cohen and Debevec
"LightGen" HDR Shop Plugin, 2001
www.debevec.org/HDRShop/

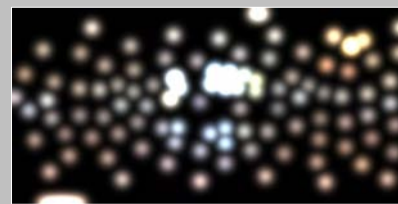

Kollig and Keller. *Efficient illumination by high dynamic range images*. Eurographics Symposium on Rendering 2003.

Agarwal, Ramamoorthi, Belongie, and Jensen, *Structured Importance Sampling of Environment Maps*, SIGGRAPH 2003

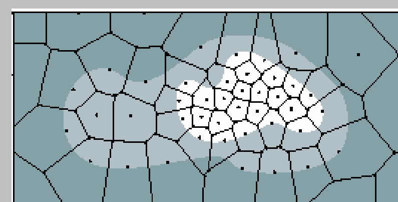

Ostromoukhov, Donohue, and Jodoin
Fast Hierarchical Importance Sampling with Blue Noise Properties
Proc. SIGGRAPH 2004

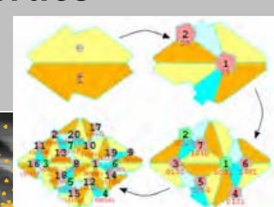

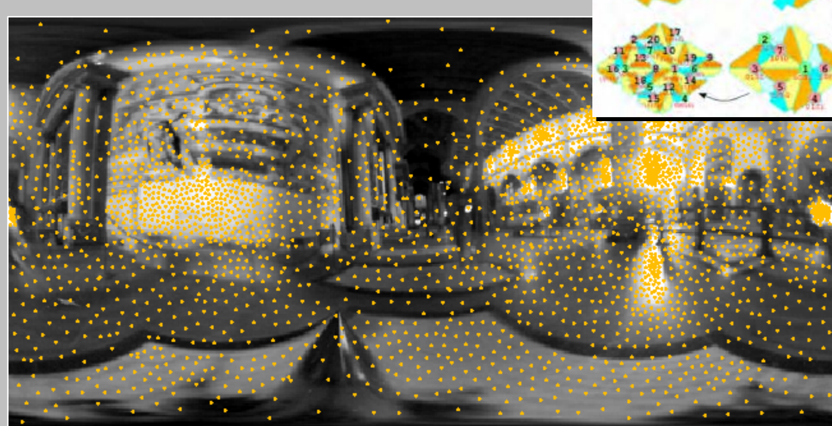

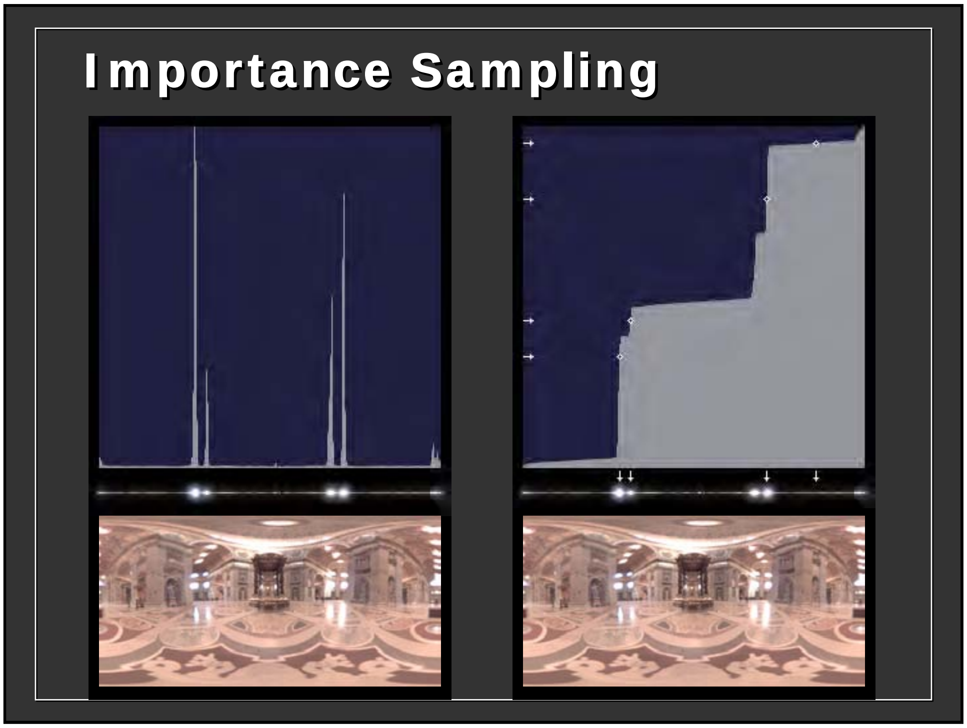
Importance Sampling

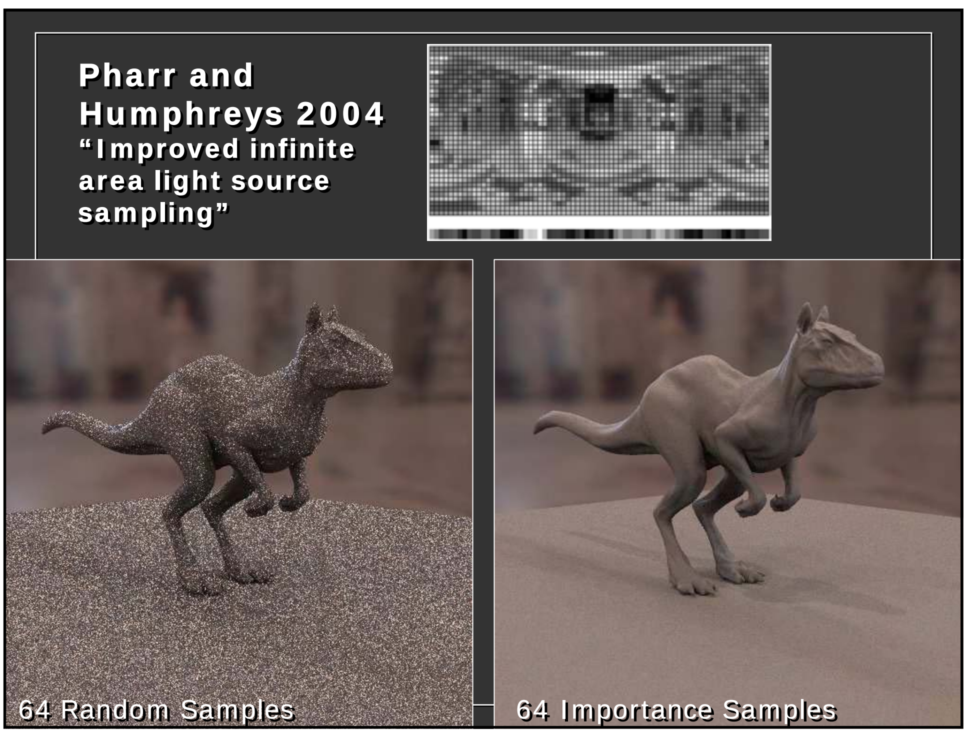
Pharr and Humphreys 2004
"Improved infinite area light source sampling"
64 Random Samples
64 Importance Samples

A Median Cut Algorithm for Light Source Placement

Inspired by Heckbert "Color Image Quantization for Frame Buffer Display", SIGGRAPH 82

1. Add the entire light probe image to the region list as a single region
2. For each region, subdivide along the longest dimension such that its light energy is divided evenly
3. If the number of iterations is less than *n*, return to step 2.

Paul Debevec, SIGGRAPH 2005 Poster

Median Cut Algorithm

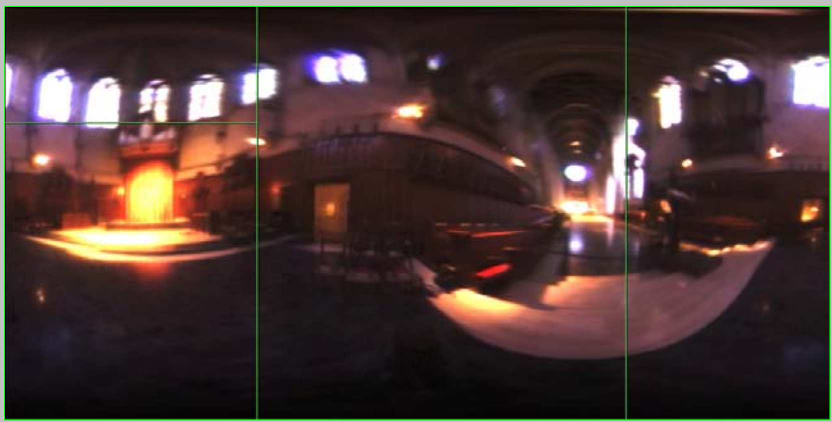

4 regions

Median Cut Algorithm

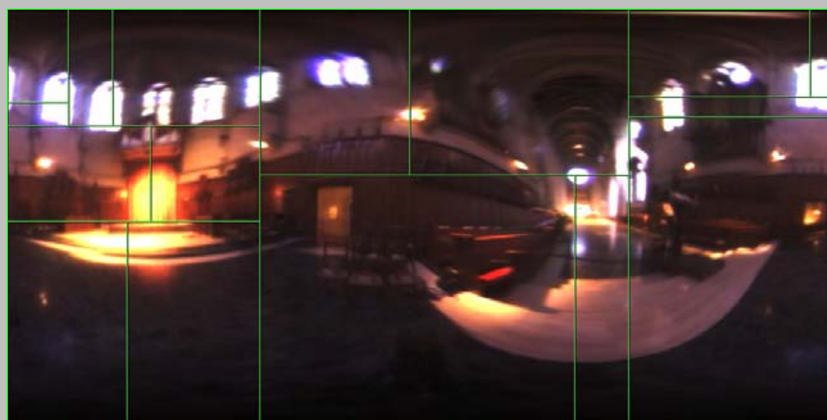

16 regions

Median Cut Algorithm

64 regions

Median Cut Algorithm

256 regions

Final Light Sources

256 lights

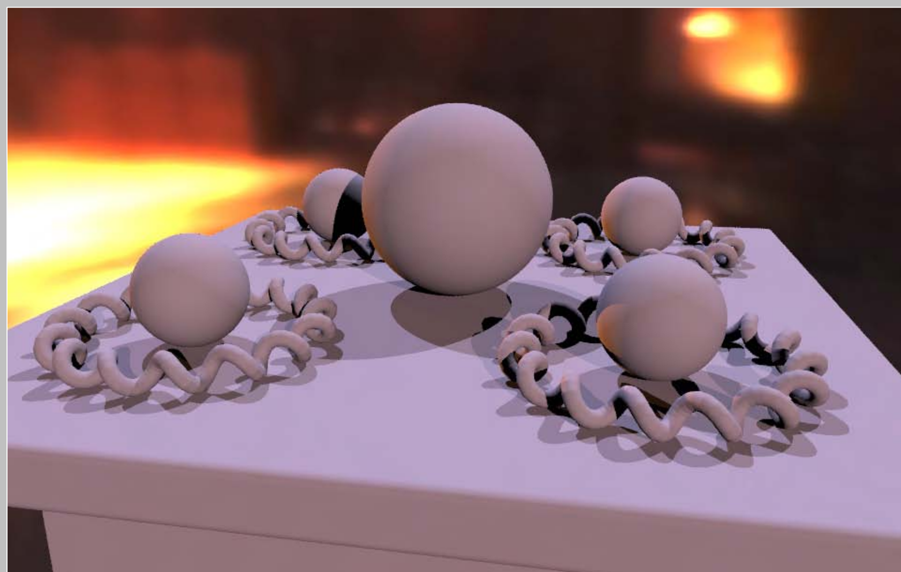

4 lights

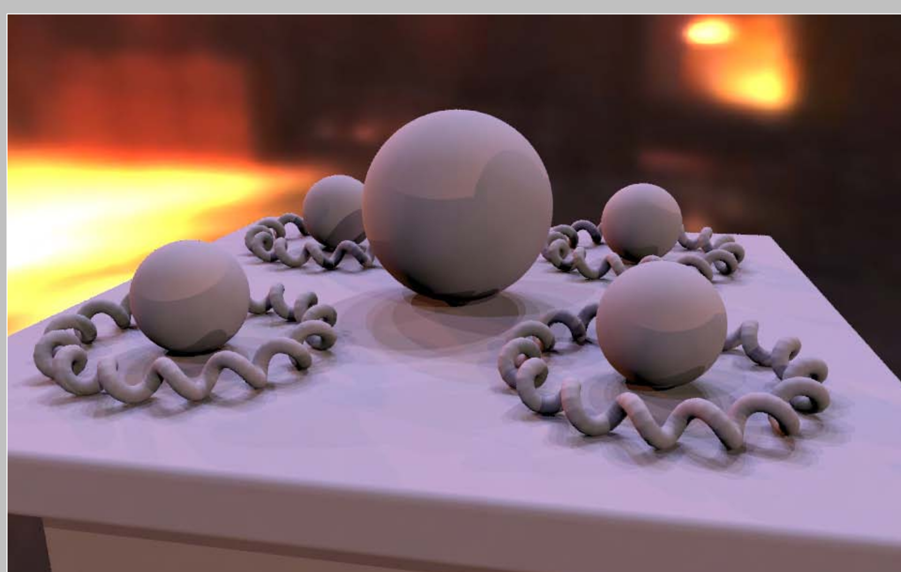

16 lights

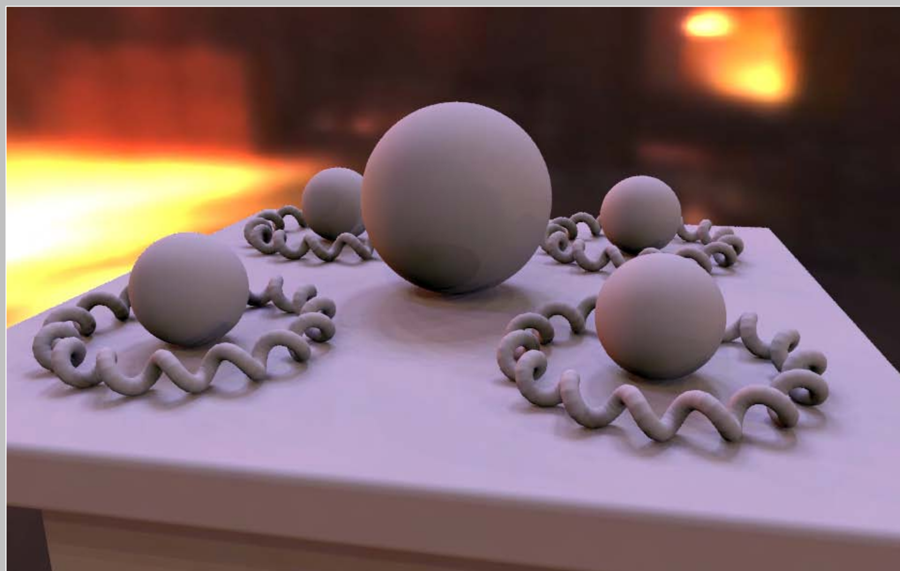

64 lights

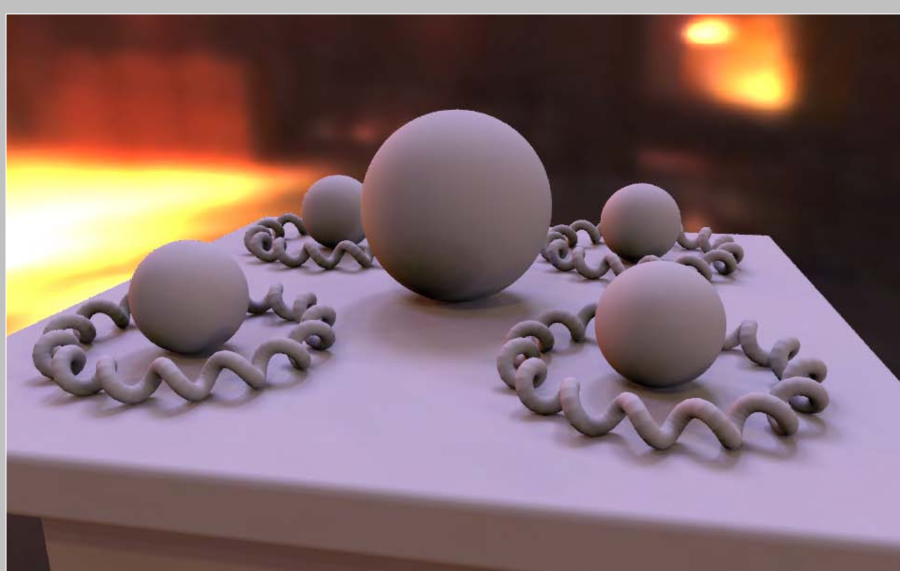

256 lights

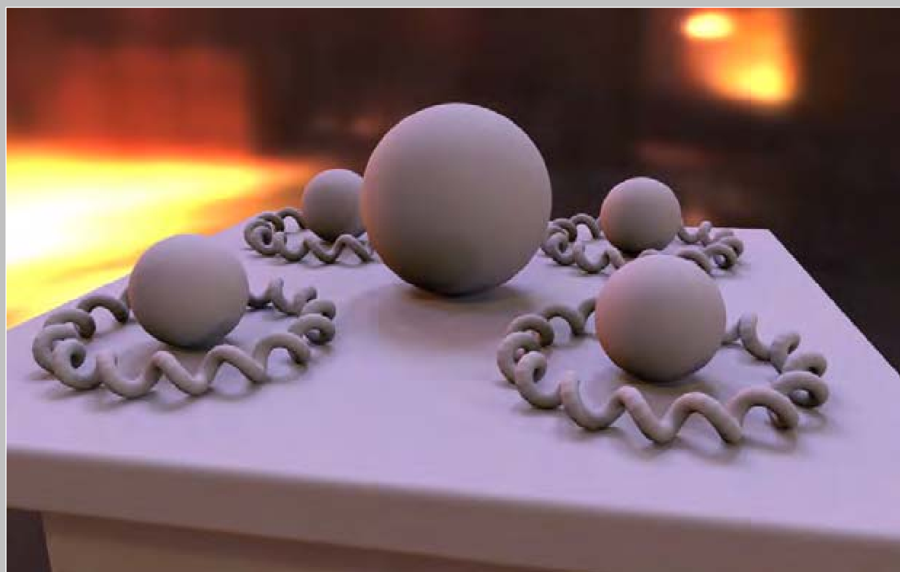

4096 random samples

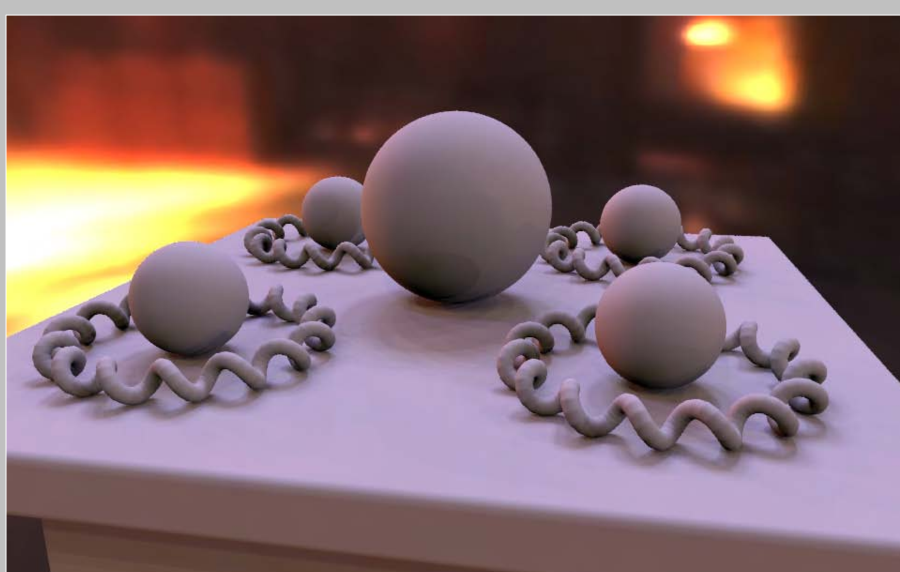

64 lights

Implementation Details

- Perform energy split decisions on a grayscale version of the image
- Can speed up region sums using summed area tables (Frank Crow, *Summed Area Tables for Texture Mapping*, SIGGRAPH 84)
- Must take into account area stretching near poles => multiply intensity and area widths by cos ϕ

Ambient Occlusion

Landis, McGaugh, and Koch; Landis et al SIGGRAPH 2002

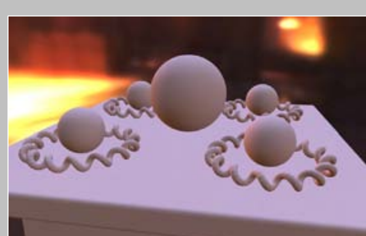

(a) diffuse env mapping

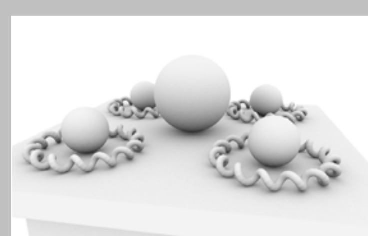

(b) ambient occlusion map

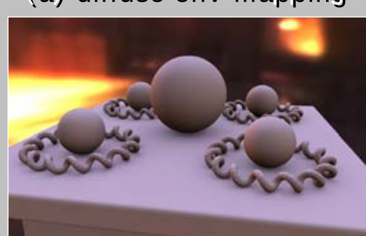

(a) * (b)

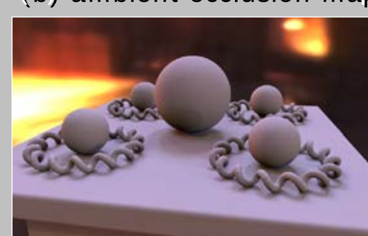

Full IBL

Ambient Occlusion

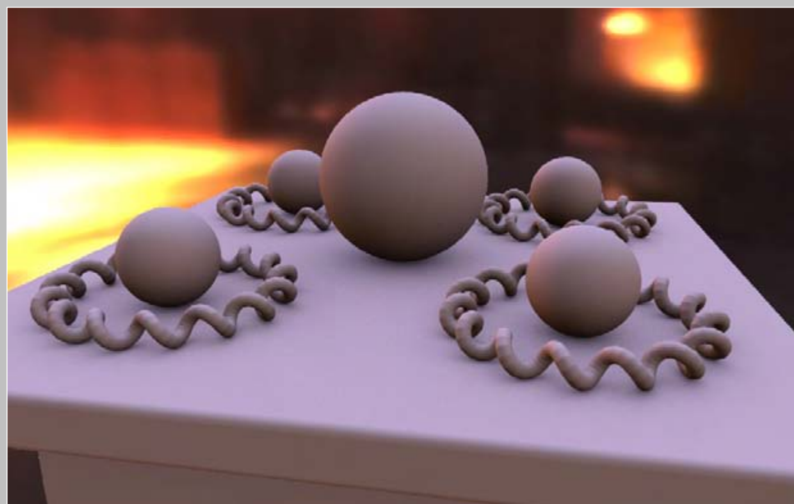

Ambient Occlusion Approximation

Ambient Occlusion

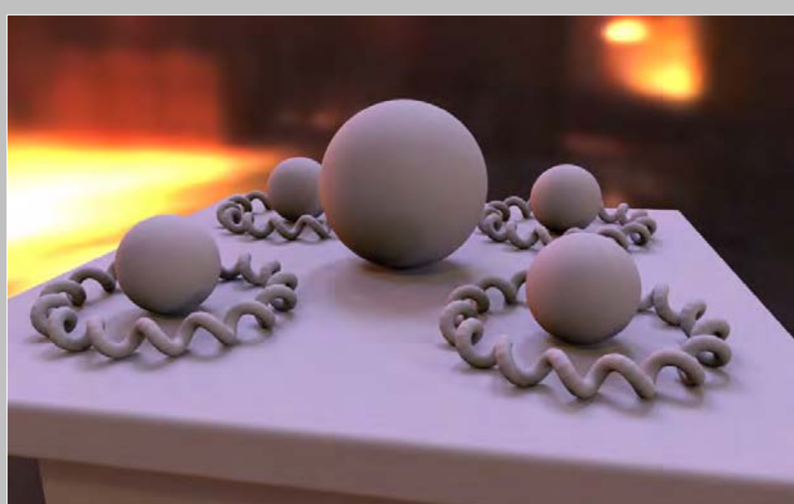

Full IBL

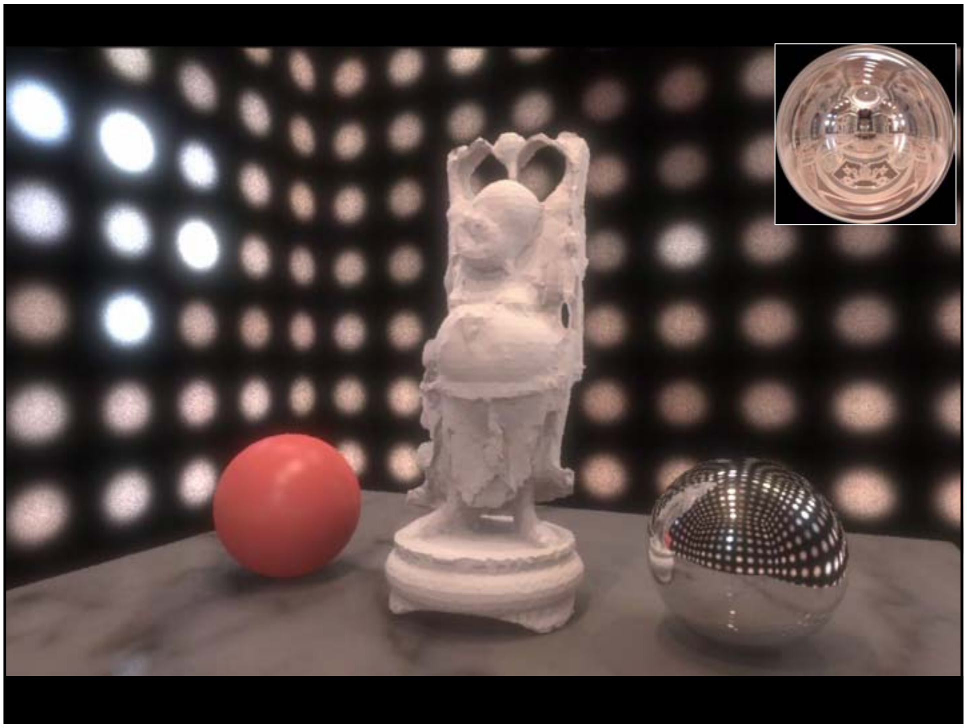

Sampling according to the BRDF

75 rays, uniform sampling

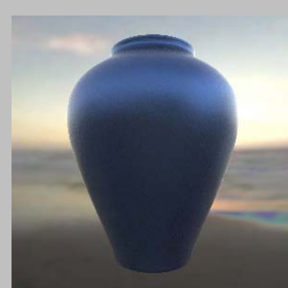

75 rays, factored BRDF sampling

From LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHI, R. 2004. Efficient BRDF importance sampling using a factored representation. *SIGGRAPH 2004*

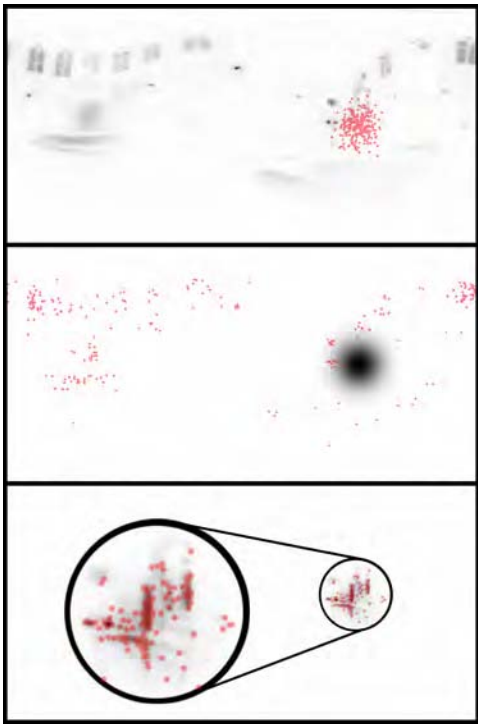

David Burke, Abhijeet Ghosh and Wolfgang Heidrich. *Bidirectional Importance Sampling for Direct Illumination*. EGSR2005.

Uses *Structured Importance Resampling (SIR)*

See Also:

Talbot et al. Importance Resampling for Global Illumination, EGSR2005

Lawrence et al, *Adaptive Numerical Cumulative Distribution Functions for Efficient Importance Sampling.* EGSR2005.

Ghosh and Heidrich. *Correlated Visibility Sampling for Direct Illumination.* SIGGRAPH 2005 Sketch.

Real-Time IBL with Spherical Harmonics

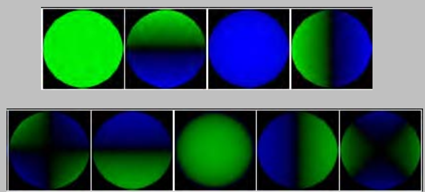

Frequency Space Environment Map Rendering
Ravi Ramamoorthi, Pat Hanrahan, SIGGRAPH 2002

Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments
Peter-Pike Sloan, Jan Kautz, John Snyder, SIGGRAPH 2002

Ng, Ramamoorthi and Hanrahan

All-frequency shadows using non-linear wavelet lighting approximation

SIGGRAPH 2003

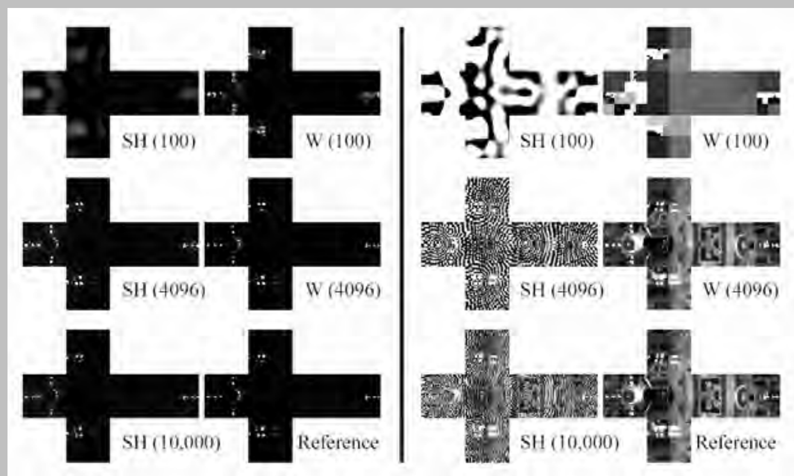

Ng, Ramamoorthi and Hanrahan

Triple Product Wavelet Integrals for All-Frequency Relighting

SIGGRAPH 2004

- **Lighting**
- **Visibility**
- **BRDF**

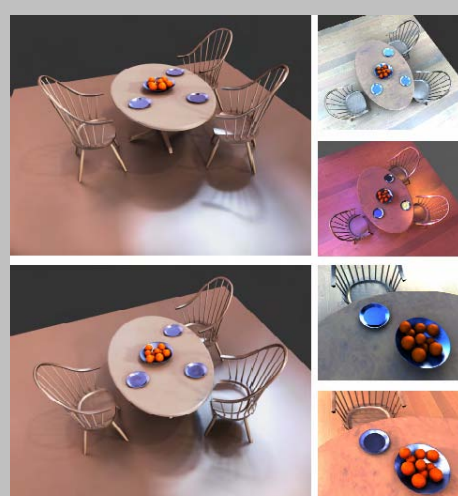

Wavelet Importance Sampling

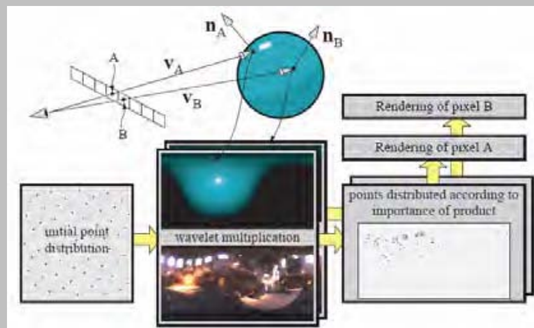

Clarberg, Jarosz, Akenine-Möller, and Jensen. Wavelet Importance Sampling: Efficiently Evaluating Products of Complex Functions. SIGGRAPH 2005.

Summary

- IBL lights objects with images of light from the real world
- Techniques such as light source identification, light source constellations, ambient occlusion, and importance sampling allow for efficient and useful approximations
- Sampling according to both the environment and the reflectance properties is the most efficient

Thanks!

- Nick Bertke, Marc Jacquier
- Greg Ward, Marcos Fajardo
- Henrik Wann Jensen, UCSD
- Hiro Matsuguma, Naomi Dennis, Hadi Ogawa, Toppan Printing
- UC Berkeley: Jitendra Malik, Larry Rowe, Westley Sarokin, HP Duiker
- ICT Graphics Lab: Chris Tchou, Brian Emerson, Marc Brownlow, Tim Hawkins, Andreas Wenger, Andrew Gardner, Andrew Jones, Jonas Unger, Frederik Gorranson, John Lai, Tom Pereira, Laurie Swanson
- ICT Sponsors: USC Office of the Provost, RDECOM, TOPPAN Printing Co, Ltd.
- Bill Swartout, David Wertheimer, Dell Lunceford, USC ICT
- Randy Hall, Max Nikias, USC
- Course Co-Organizer and speaker: Erik Reinhard, Greg Ward

gl.ict.usc.edu www.debevec.org

HDR Display Devices

Helge Seetzen

BRIGHTSIDE™

HDR Display Overview

August 2006

BRIGHTSIDE™

Imaging Pipeline – Low dynamic range rules!

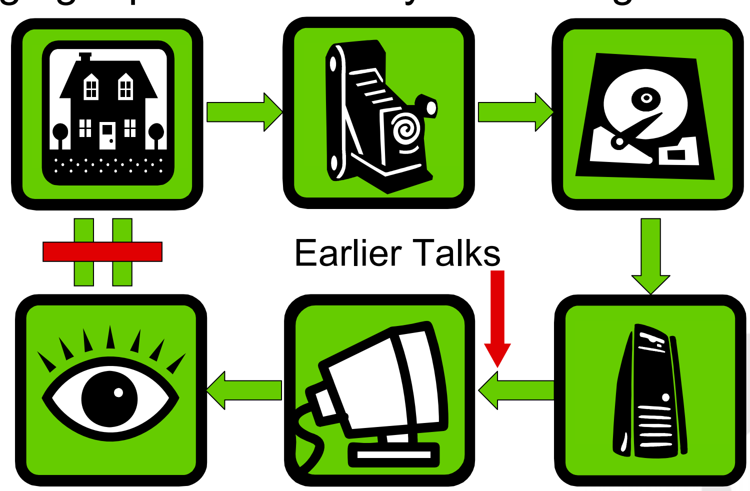

BRIGHTSIDE™

What are we missing?

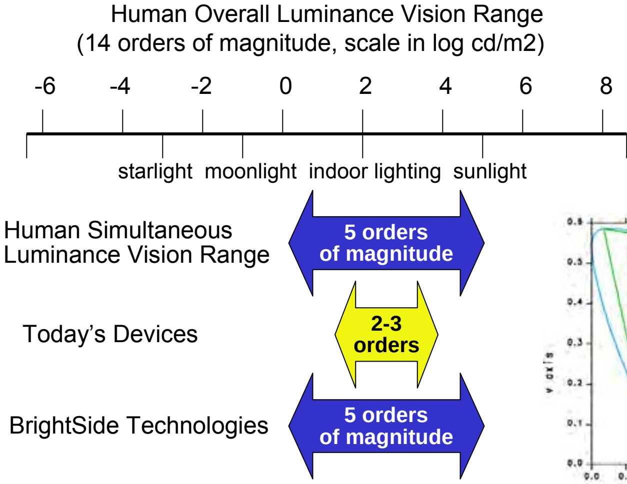

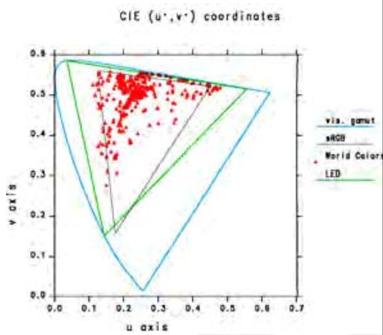

BRIGHTSIDE™

Designing Commercial HDR Devices

Requirements:

1. High Dynamic Range
2. Compatibility
   - New devices need to function in 8-bit environment and still deliver significant benefit
   - New devices need to be usable in stand-alone mode
3. Cost
   - Ideally no extra cost compared to 8-bit devices
   - If extra cost is necessary then in line with benefit

BRIGHTSIDE™

Limits of Conventional Output Devices

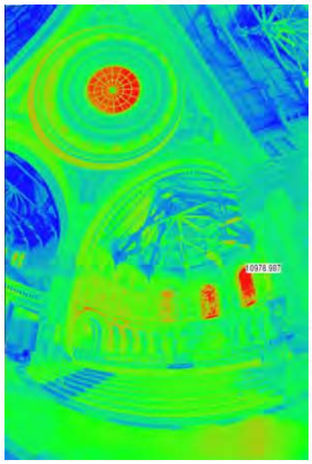

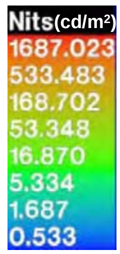

Stanford Memorial Church
Courtesy Paul Debevec

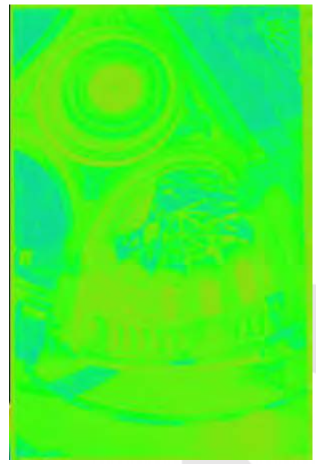

BRIGHTSIDE™

Display Technology – Conventional LCD’s

- LCD backlight is provided by CCFL tubes (fluorescent light)
- Light is spread evenly behind LCD panel and does not vary with image content
- Image control is limited to 8 bit single to Red, Green and Blue colour channels (255 steps of control)

Display Technology – HDR Display

- LCD backlight is provided by an array of LED's
- Each LED is controlled with 8 bit (255 step) signal
  - Brightness is adjusted to level demanded by source image
- LCD panel provides additional 8 bits of brightness control
- LED and LCD panel combine optically to deliver 16 bit performance
- LED's provide greater brightness

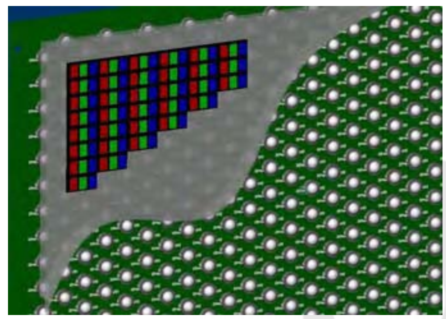

Display Technology – HDR Display

High resolution colour LCD | High Dynamic Range Display | Low resolution Individually Modulated LED array

- Dual modulation
- Low / high resolution and correction
- Veiling luminance
- Implementation in display and projection

BRIGHTSIDE™

Image Processing Algorithm

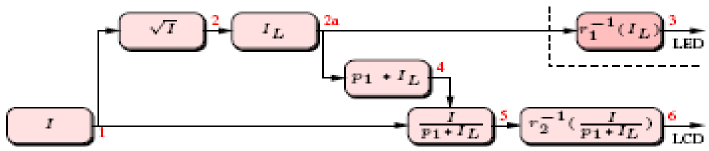

BRIGHTSIDE™

Image Processing Algorithm

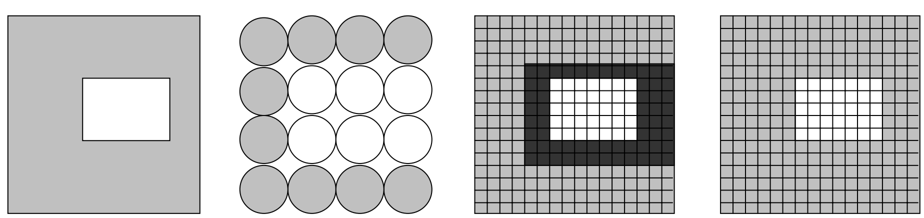

HDR Image **LED array** **LCD with correction** **Output image**

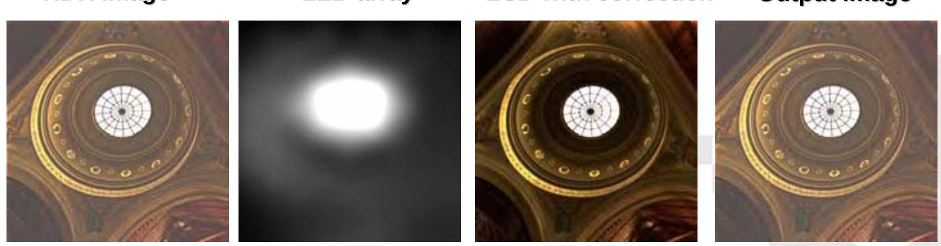

BRIGHTSIDE™

Image Processing Algorithm - Easy

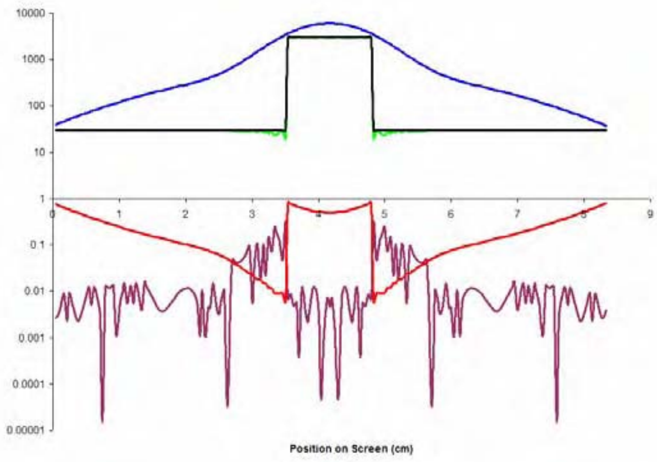

Receive Image

Drive LED

Divide Image by LED light field to obtain LCD values

Output Luminance is the product of LED light field and LCD transmission (modest error)

4/29/2006 ©2005 Brightside Technologies Inc. Proprietary & Confidential 11

BRIGHTSIDE™

Image Processing Algorithm - Hard

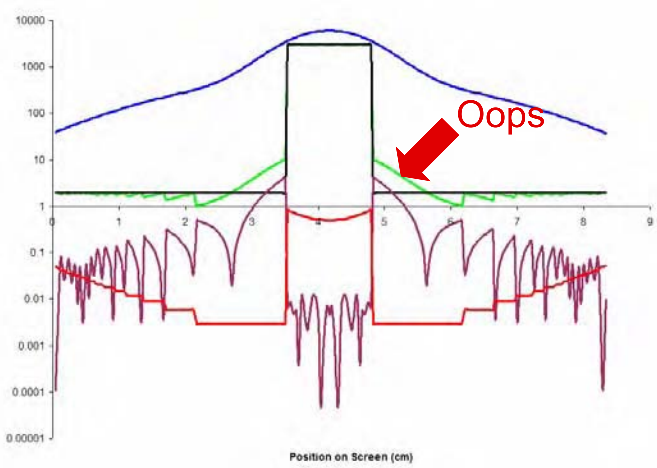

Receive Image

Drive LED

Divide Image by LED light field to obtain LCD values

Output Luminance is the product of LED light field and LCD transmission (Problematic error)

4/29/2006 ©2005 Brightside Technologies Inc. Proprietary & Confidential 12

BRIGHTSIDE™

Veiling Luminance

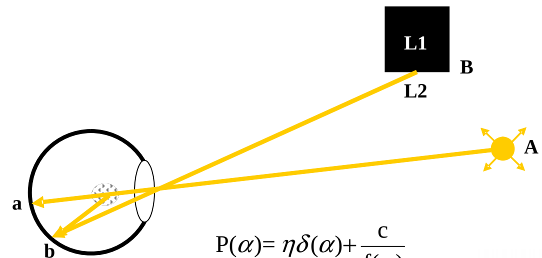

$$P(\alpha) = \eta\delta(\alpha) + \frac{c}{f(\alpha)}$$

BRIGHTSIDE™

Image Processing Algorithm - Hard

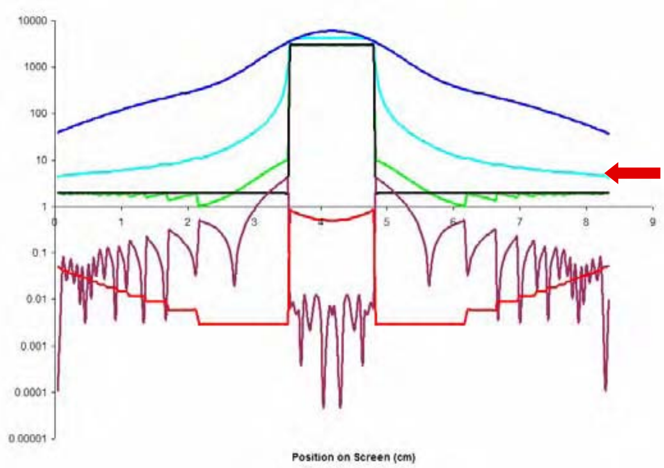

Veiling Luminance masks imperfection

BRIGHTSIDE™

Display Technology – Review

- Compatibility
  - Based on commercially available components (LCD, LED)
  - Legacy support through Reverse Tone Mapping and Saturation Extension
  - Small number of LEDs allows encoding of LED data in conventional video signal
- Cost
  - LED cost money (less every day)
  - Significant power reduction (~25% of comparable constant backlight LCD on average)

BRIGHTSIDE™

Display Technology – Results

Seeing is Believing!

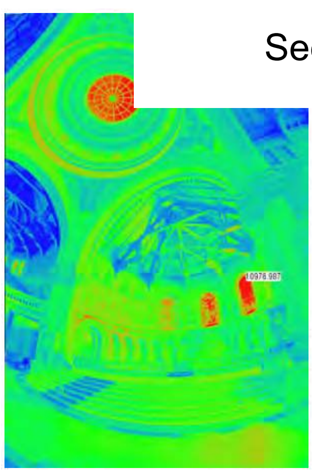

Stanford Memorial Church
Courtesy Paul Debevec

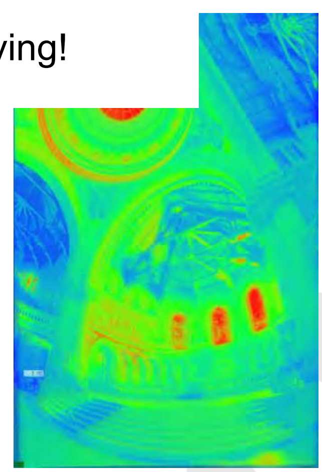

Projection Technology – Concept

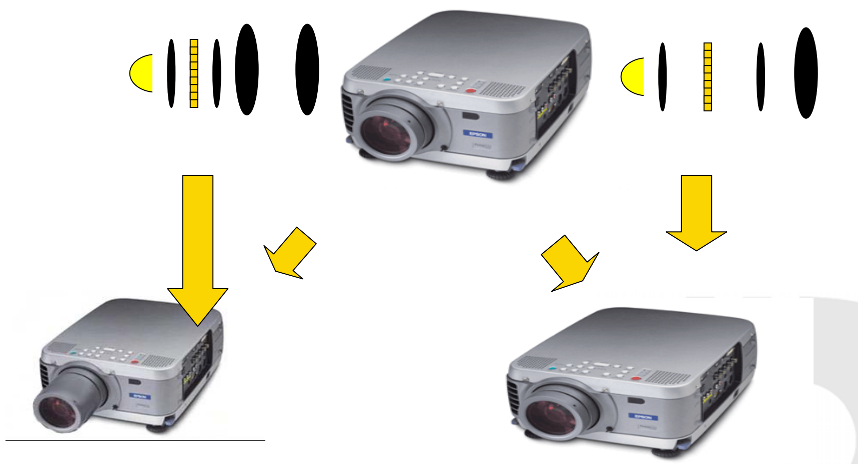

HDR Projection - LCoS Implementation

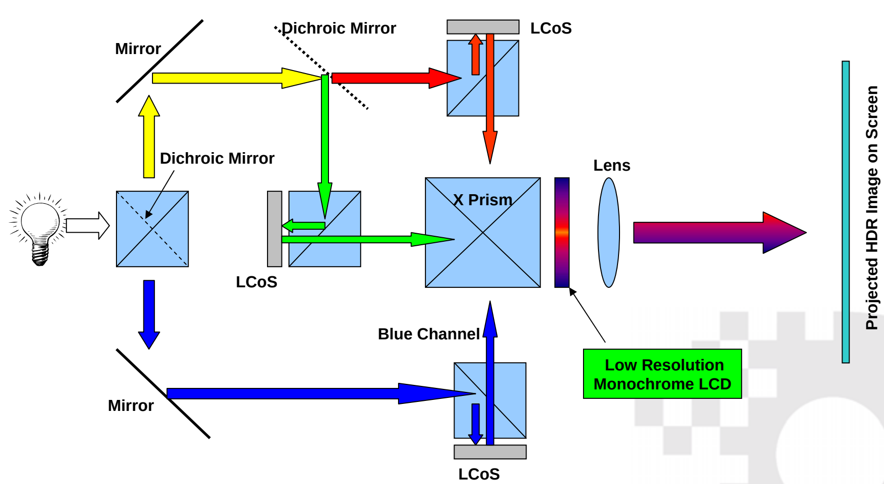

BRIGHTSIDE™

HDR Projection – 3 LCD Implementation

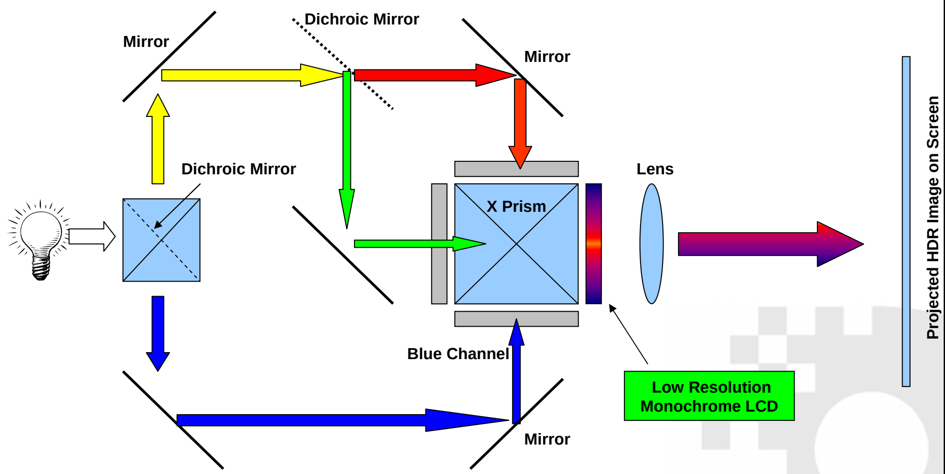

BRIGHTSIDE™

HDR Projection – DLP Implementation

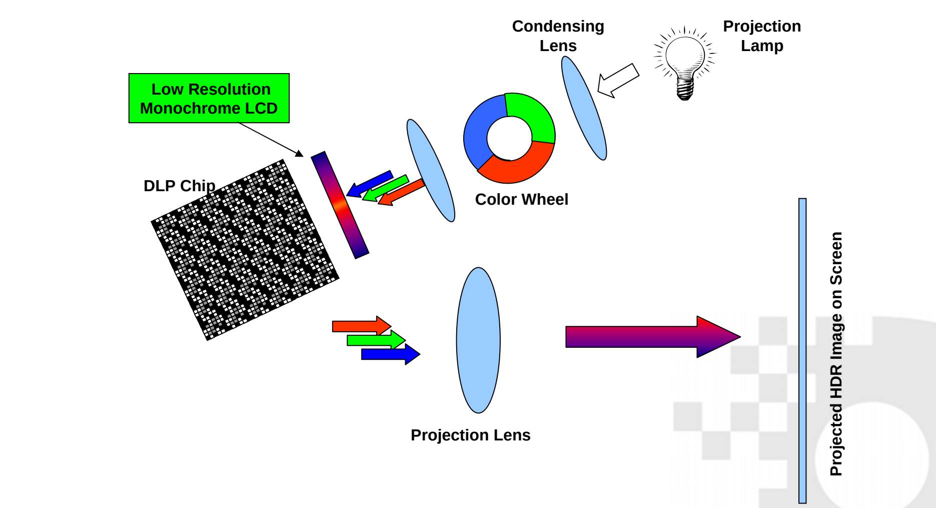

BRIGHTSIDE™

Technology Overview – A word on Contrast

- High Dynamic Range = Contrast + Amplitude Resolution

- Contrast depends on
  - Measurement technique
  - Ambient environment
  - Device Limitations

BRIGHTSIDE™

Technology Overview – A word on Contrast

Contrast Measurement Techniques

ANSI Contrast

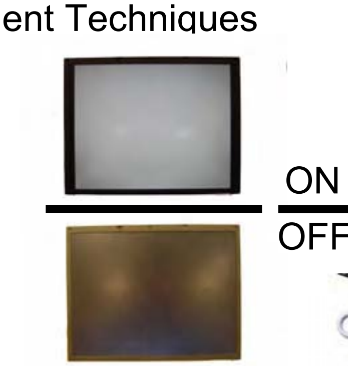

On/Off Contrast

BRIGHTSIDE™

Technology Overview – A word on Contrast

Ambient Environment

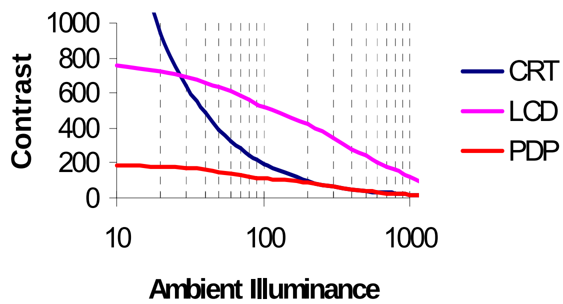

Source: Sharp Corp

BRIGHTSIDE™

Technology Overview – A word on Contrast

Device Limitations

- CRT: high front surface reflection
- LCD: Overdrive might prevent black
- Plasma: Size limit on white area

BRIGHTSIDE™

Technology Overview – Novel Concepts

Organic Light Emitting Diodes

- Good contrast
- Limited brightness
- Lifetime issues
- Optimists predict ~5years to compete with current LCD

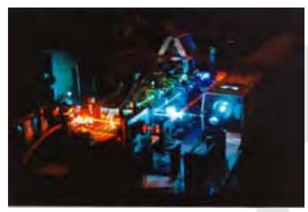

Laser Projection

- Very high contrast
- Green & blue laser not available
- Optimists predict ~10-15 years

BRIGHTSIDE™

We have HDR Display – How to feed them?

- Floating point GPUs are becoming standard
- Most CG and Games done in HDR today
- TV remains low dynamic range and is unlikely to change soon
- Reverse Tone Mapping and Saturation Extension provide quality gain even for legacy content

BRIGHTSIDE™

LDR to HDR Signal Extension

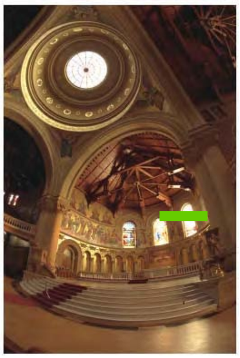

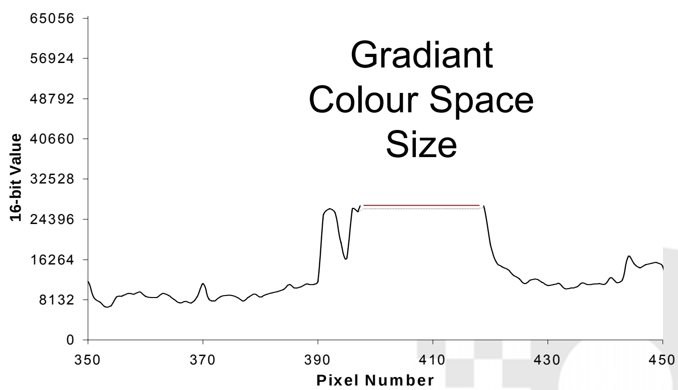

BRIGHTSIDE™

LDR to HDR Signal Extension

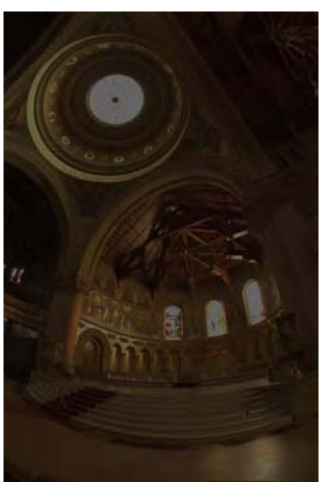

Compatibility
- only needed for legacy content
- calculation supported in 8-bit processor

Cost
- very fast (1-2 operations per pixel)

BRIGHTSIDE™

LDR to HDR Signal Extension

BRIGHTSIDE™

Summary

- HDR Displays are available
- HDR Projectors are coming
- HDR Output Devices can provide benefit to legacy content
- HDR Software / Input Devices are leading the way

HDR at Industrial Light + Magic

Drew Hess

HDR in Film Production

Drew Hess
<dhess@ilm.com>
Industrial Light + Magic

SIGGRAPH 2006 Course #5
High Dynamic Range Imaging: Theory and Applications
Full Day, Sunday, 30 July, 8:30AM - 5:30PM

Why is HDR important?

Motion Blur (bad)

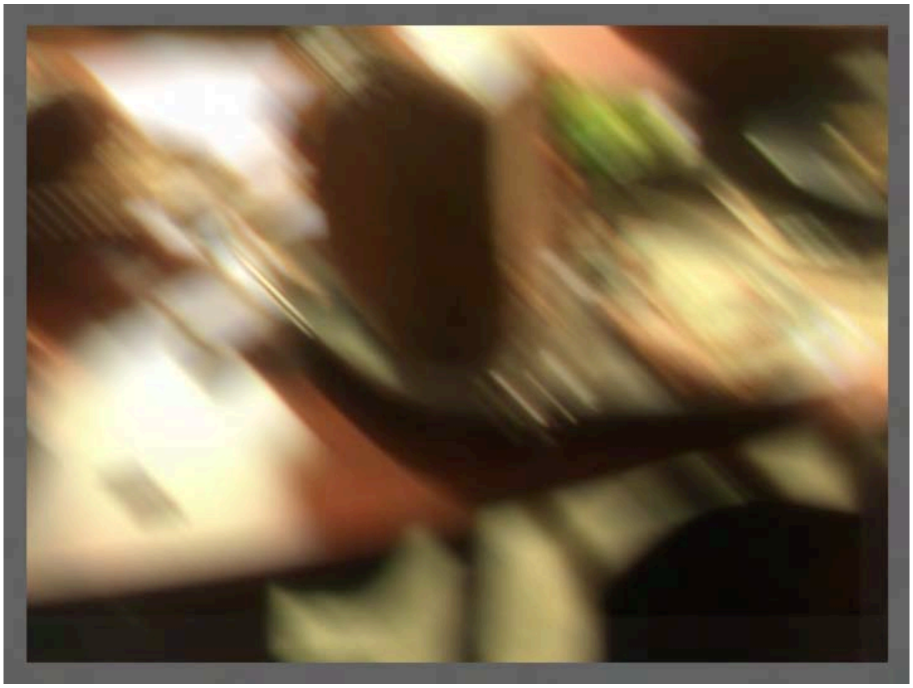

Motion Blur (good)

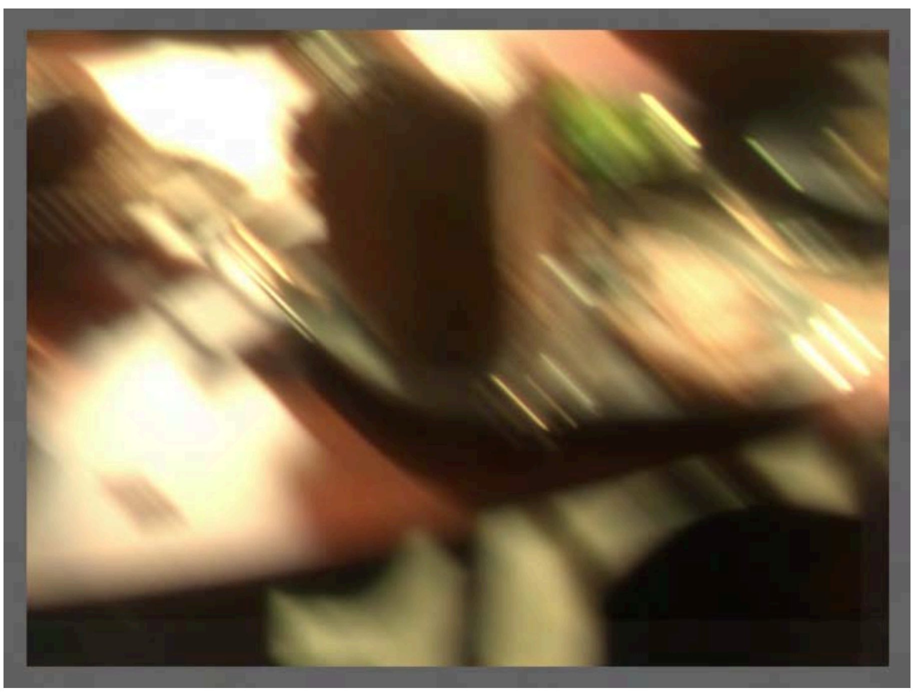

Film Characteristic Curve

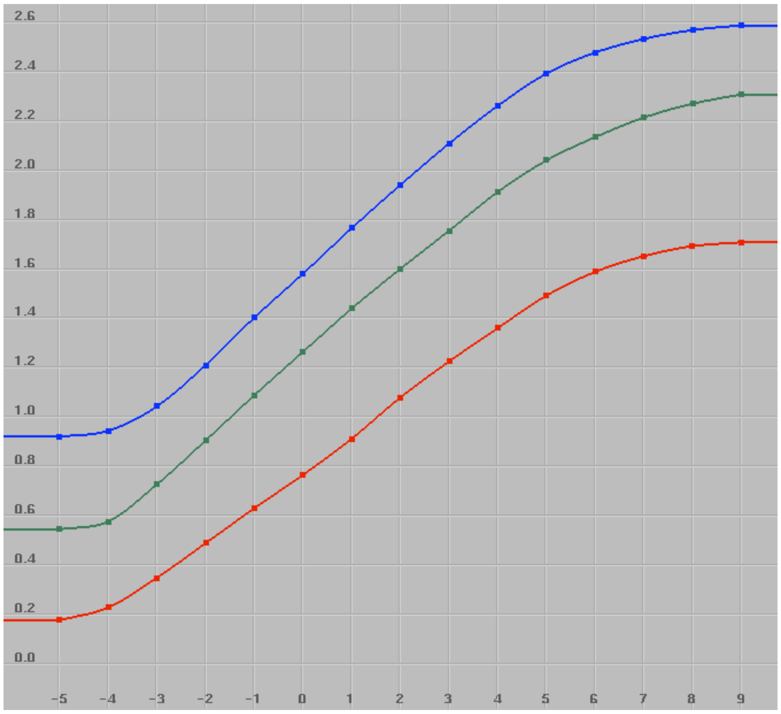

Film Characteristic Curve

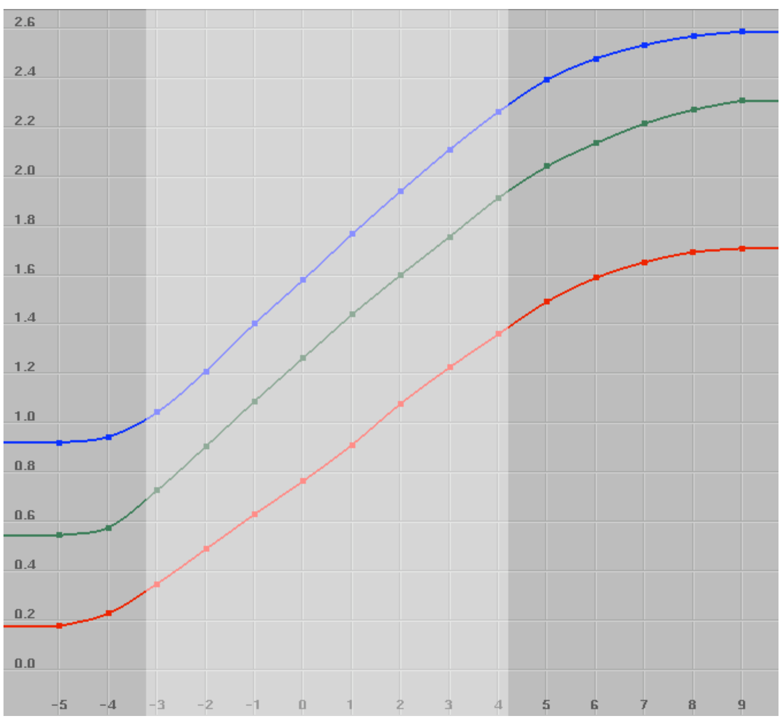

SpheroCamHDR

-8 stops

SpheroCamHDR

-4 stops

SpheroCamHDR

neutral

SpheroCamHDR

+4 stops

CG image

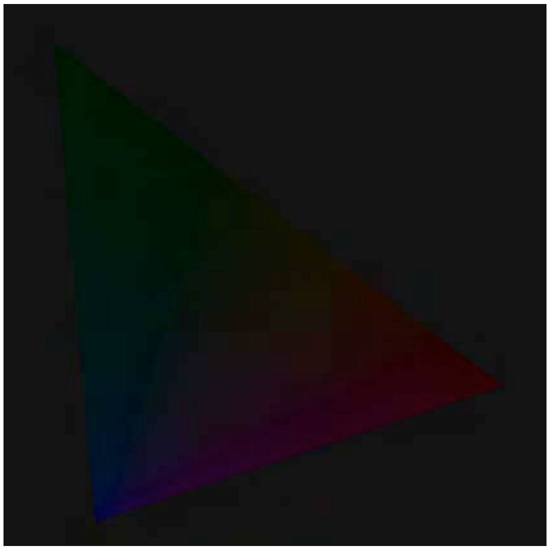

-8 stops

CG image

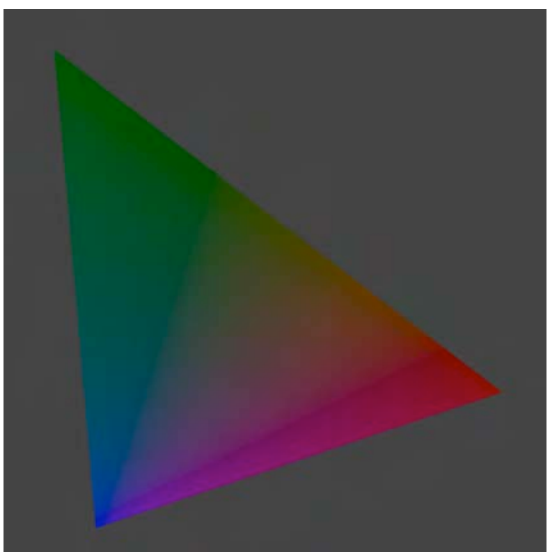

-4 stops

CG image

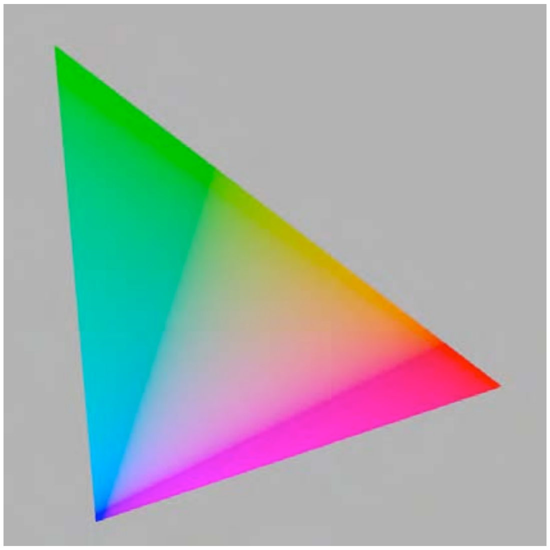

neutral

CG image

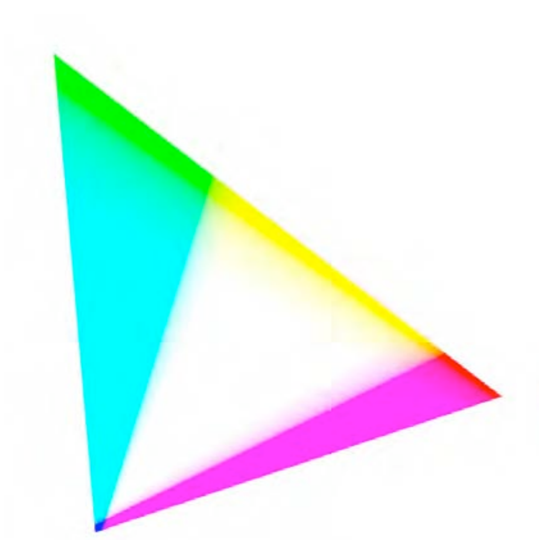

+4 stops

OpenEXR

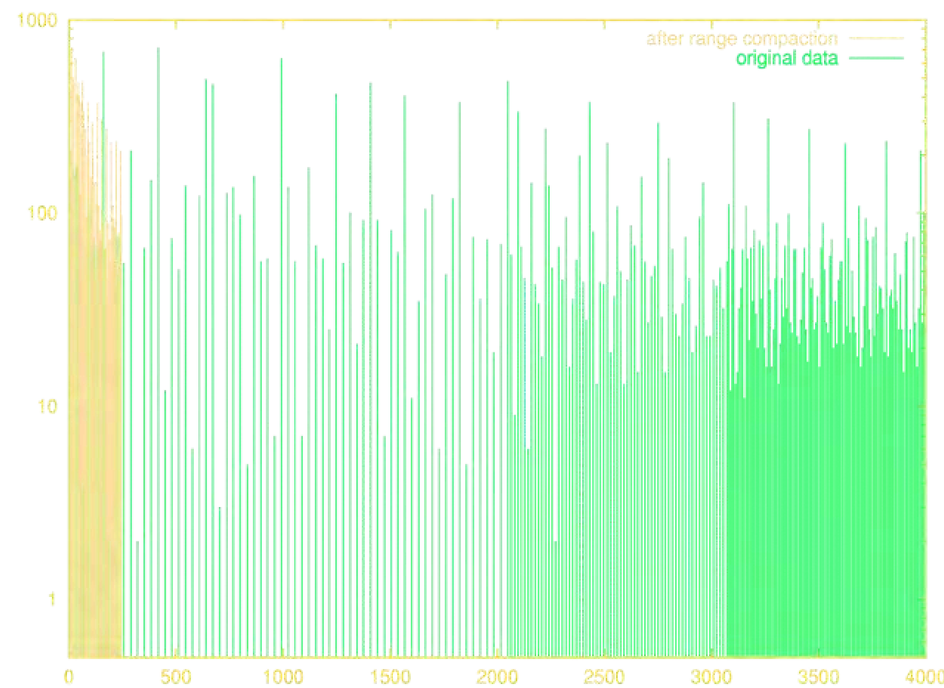

Wavelet Transform - Haar Wavelets

forward: **L = (A+B)/2**

H = A-B

reverse: **A = L+H/2+ { 1, if H is odd; 0, if H is even**

B = A-H

Original Data

16	15	14	13	12	11	10	9
15	14	13	12	11	10	9	8
14	13	12	11	10	9	8	7
13	12	11	10	9	8	7	6
12	11	10	9	8	7	6	5
11	10	9	8	7	6	5	4
10	9	8	7	6	5	4	3
9	8	7	6	5	4	3	2

15	1	13	1	11	1	9	1
14	1	12	1	10	1	8	1
13	1	11	1	9	1	7	1
12	1	10	1	8	1	6	1
11	1	9	1	7	1	5	1
10	1	8	1	6	1	4	1
9	1	7	1	5	1	3	1
8	1	6	1	4	1	2	1

14	1	12	1	10	1	8	1
1	0	1	0	1	0	1	0
12	1	10	1	8	1	6	1
1	0	1	0	1	0	1	0
10	1	8	1	6	1	4	1
1	0	1	0	1	0	1	0
8	1	6	1	4	1	2	1
1	0	1	0	1	0	1	0

13	1	2	1	9	1	2	1
1	0	1	0	1	0	1	0
11	1	2	1	7	1	2	1
1	0	1	0	1	0	1	0
9	1	2	1	5	1	2	1
1	0	1	0	1	0	1	0
7	1	2	1	3	1	2	1
1	0	1	0	1	0	1	0

12	1	2	1	8	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0
8	1	2	1	4	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0

10	1	2	1	4	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0
6	1	2	1	4	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0

Transformed Data

8	1	2	1	4	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0
4	1	2	1	0	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0

Original Data

16	15	14	13	12	11	10	9
15	14	13	12	11	10	9	8
14	13	12	11	10	9	8	7
13	12	11	10	9	8	7	6
12	11	10	9	8	7	6	5
11	10	9	8	7	6	5	4
10	9	8	7	6	5	4	3
9	8	7	6	5	4	3	2

Transformed Data

8	1	2	1	4	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0
4	1	2	1	0	1	2	1
1	0	1	0	1	0	1	0
2	1	0	1	2	1	0	1
1	0	1	0	1	0	1	0

value	# occurrences	bits
2	1	6.0
3	2	5.0
4	3	4.4
5	4	4.0
6	5	3.7
7	6	3.4
8	7	3.2
9	8	3.0
10	7	3.2
11	6	3.4
12	5	3.7
13	4	4.0
14	3	4.4
15	2	5.0
16	1	6.0

value	# occurrences	bits
0	21	1.6
1	32	1.0
2	8	3.0
4	2	5.0
8	1	6.0

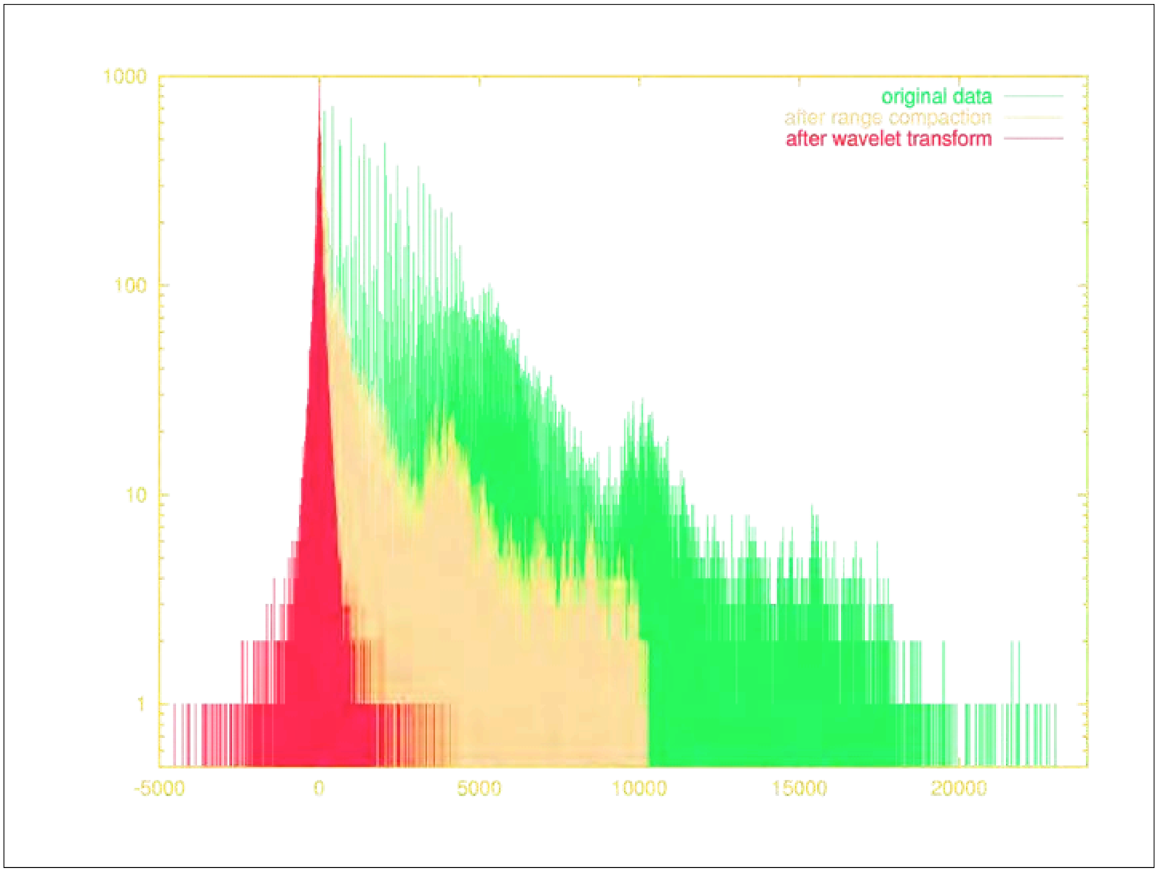

original data
after range compaction
after wavelet transform
1000
100
10
1
-5000
0
5000
10000
15000
20000

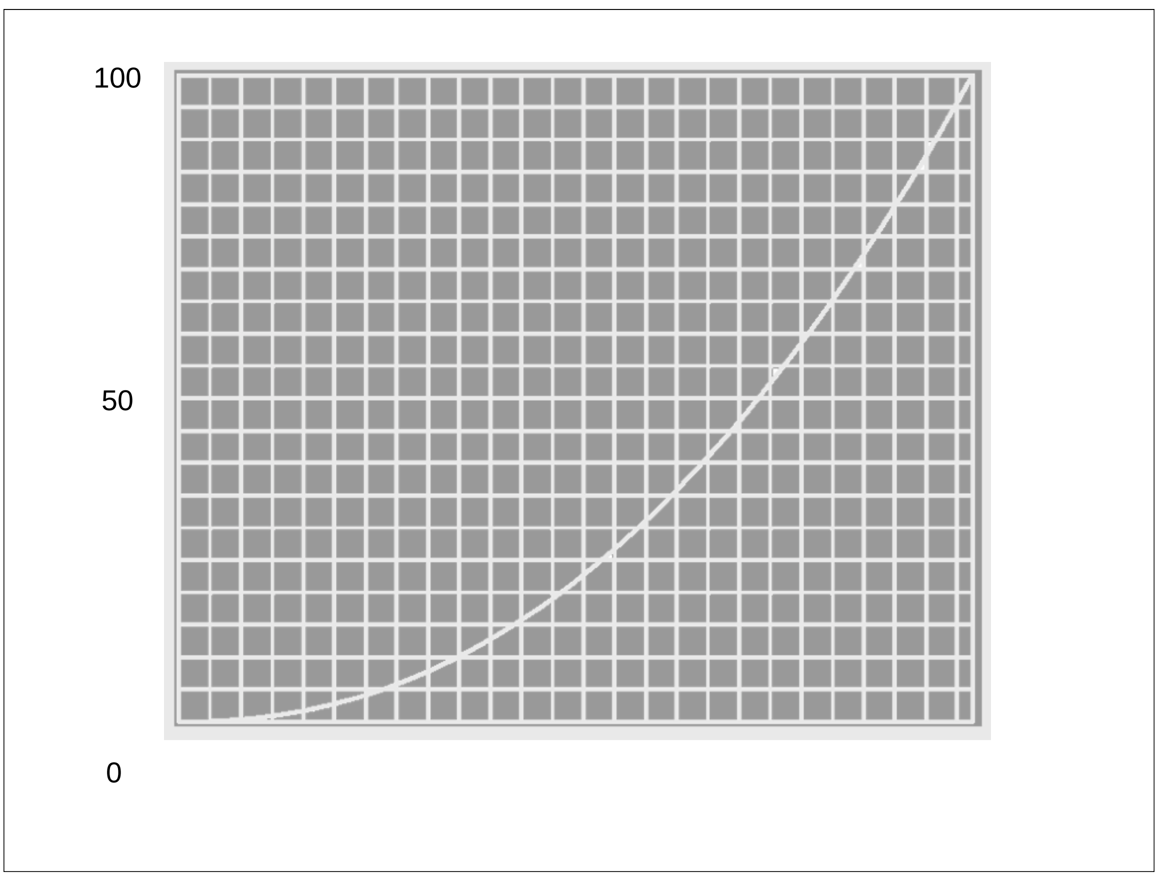

100
50
0

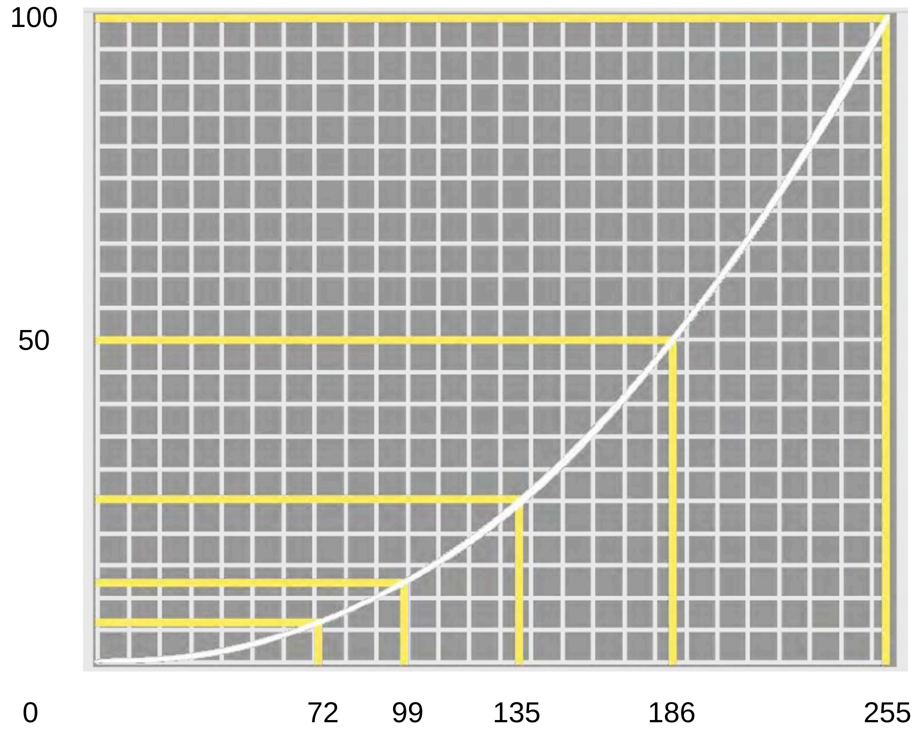

Implications of non-linear pixels

$$99 \text{ x } \cancel{2 =} 198$$

$$99 \text{ x } 2 = 135$$

$$(99/255)^{2.2} = 0.1247$$

$$(135/255)^{2.2} = 0.2468$$

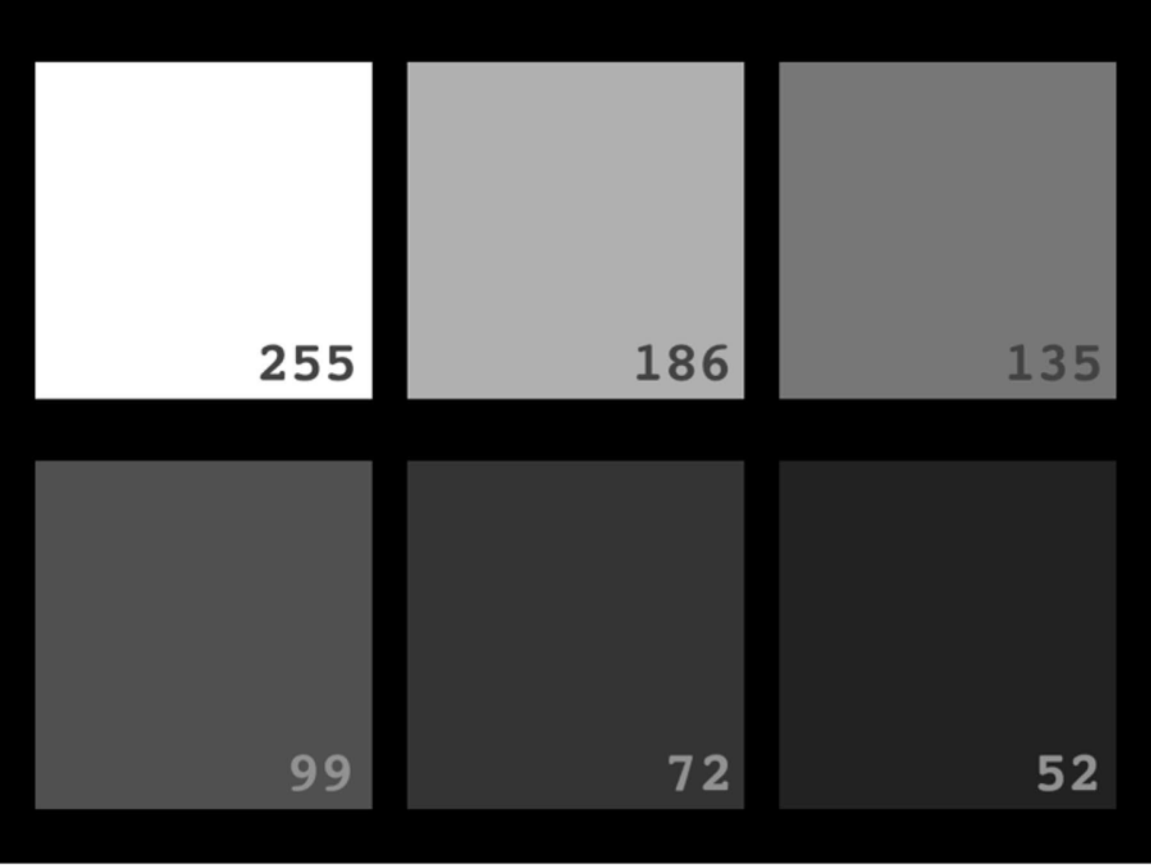
255
186
135
99
72
52

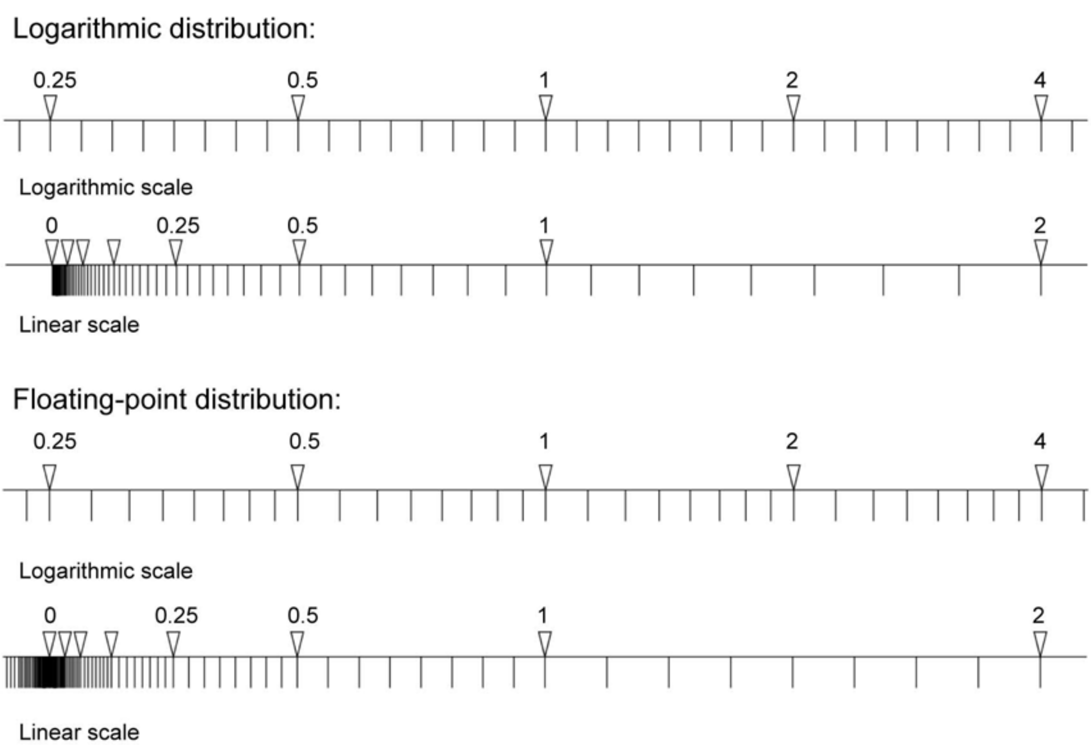
Logarithmic distribution:
0.25
0.5
1
2
4
Logarithmic scale
0
0.25
0.5
1
2
Linear scale
Floating-point distribution:
0.25
0.5
1
2
4
Logarithmic scale
0
0.25
0.5
1
2
Linear scale

half

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

S E M

S	**1 bit**	**sign bit**
E	**5 bits**	**exponent**
M	**10 bits**	**mantissa**

Value (normalized case):

$$v = (-1)^{S} * 2^{E-15} * (1 + M/1024)$$

half

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

S E M

Range:

30 f-stops

maximum **65504.0**

minimum $\left\{ \begin{array}{l} \textbf{0.00061 (normalized)} \\ \textbf{0.00000006 (denormalized)} \end{array} \right.$

Precision:

1024 increments per stop
0.1% relative error

Tone mapping

“neutral”

Tone mapping

Deluxe Vision

Tone mapping

Bleach bypass

Tone mapping

DLP

HDR at Valve

Gary McTaggart

High-Dynamic Range Rendering

Valve Design Team

Agenda

- What is HDR?
- Technology – HDR in Source
- Production – Creating HDR Content
- Challenges
- Benchmarking and visualization
- Lost Coast demo
- Playtest Day of Defeat
- Q&A and giving you a build to take away

“The ‘dynamic range’ of a scene is the contrast ratio between its brightest and darkest parts. A plate of evenly-lit mashed potatoes outside on a cloudy day is low-dynamic range. The interior of an ornate cathedral with light streaming in through its stained-glass windows is **high dynamic range**. In fact, any scene in which the light sources can be seen directly is high dynamic range.”

– Paul Debevec

HDR

- A High-Dynamic Range image is an image that has a greater contrast range than can be shown on a standard display device, or that can be captured with a standard camera with just a single exposure.

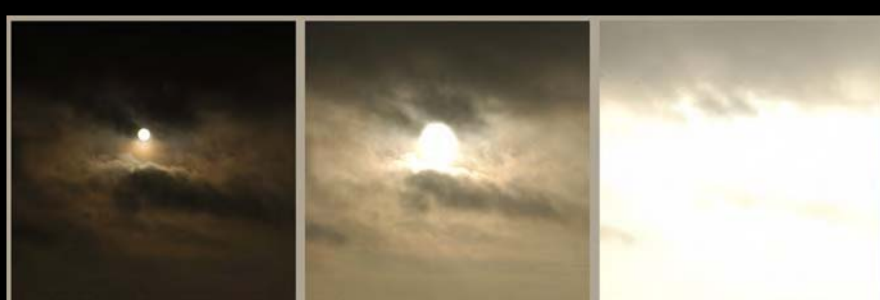

f22 @ 1/1600th second | f22 @ 1/250th second | f22 @ 1/40th second

HDR

- High-Dynamic Range rendering attempts to take an HDR image and produce a more realistic representation on a limited-range computer monitor.

Tonemap scale = 0.5 | Tonemap scale = 1 | Tonemap scale = 8

Lost Coast and HDR

- HDR as incremental technology
  - More efficient way to develop technology
  - Delivers value to customers and licensees
  - In all future games (DoD, Aftermath, ...)
- Part of the graphics roadmap
- Shorten the delay between hardware availability and game support and then target the delivery
- Give the GPU developers production code with which to evaluate their decisions

HDR in games

- So far, most video game “HDR” implementations are really just blooming.
- Content not authored for HDR since HDR is more of an afterthought.
- Since the definition of HDR varies within the game industry, we have a list of what we think is important for HDR.

HDR and Source

- Blooming – adding a blurred version of the bright parts of the scene to emulate a camera’s overexposure

VALVE

HDR and Source

- HDR skybox – authored by painting multiple exposures of the sky to allow for real-time exposure adjustment.

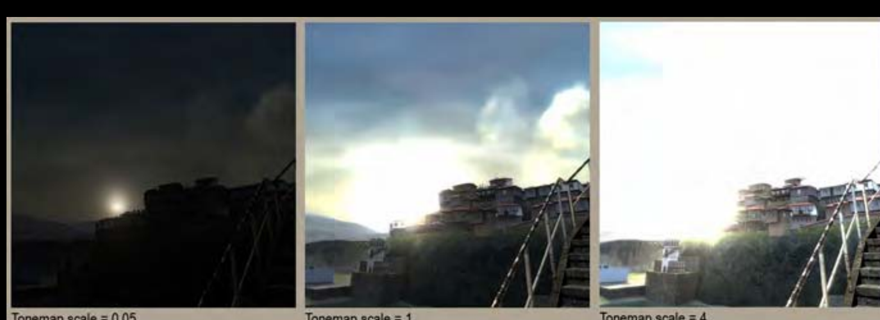

HDR and Source

- HDR cube maps – generated by the engine using the HDR skybox as well as the HDR light sources and the HDR light maps. This means that if an objects is reflecting the sun or some other bright part of the scene, you will see this in the full affect of the brightness in the reflection.

HDR and Source

- HDR water reflection/refraction – Where the reflection of the sky is really really bright, you get white hot spots along with blooming. If you are under water looking out, you'll get a similar effect.

VALVE

HDR and Source

- HDR light sources – lit like the real world instead of having to compress light values. Higher range of light values to light the world.
- HDR light maps – Generated from a radiosity process that takes light bounces/global illumination into account. You see this on the wall of the cathedral where the sun is blowing out the wall.

VALVE

HDR and Source

- Auto Exposure – Your eye adjusts to allow you to see details in dark scenes along with being able to see well outdoors. This is analogous to a camera's auto-exposure.
- This does impact game design (e.g. visibility in multiplayer games in moving from light to dark)

HDR Implementations

- HDR in real-time with hardware acceleration is cutting edge technology
- Lots of side effects and consequences were unknown when we started
- As a result, we implemented 4 different methods
- We will give an overview of the advantages and disadvantages of each method

HDR Implementation #1

- Demonstrated in Fall of 2003 at Shader Day
- Stored HDR textures as RGBA textures where Alpha contained a value between 0 and 16 to multiply RGB by.
- We learned from this demo that doing a full HDR implementation for HL2 would require a huge content and technology investment. We postponed HDR for HL2 in early 2004.

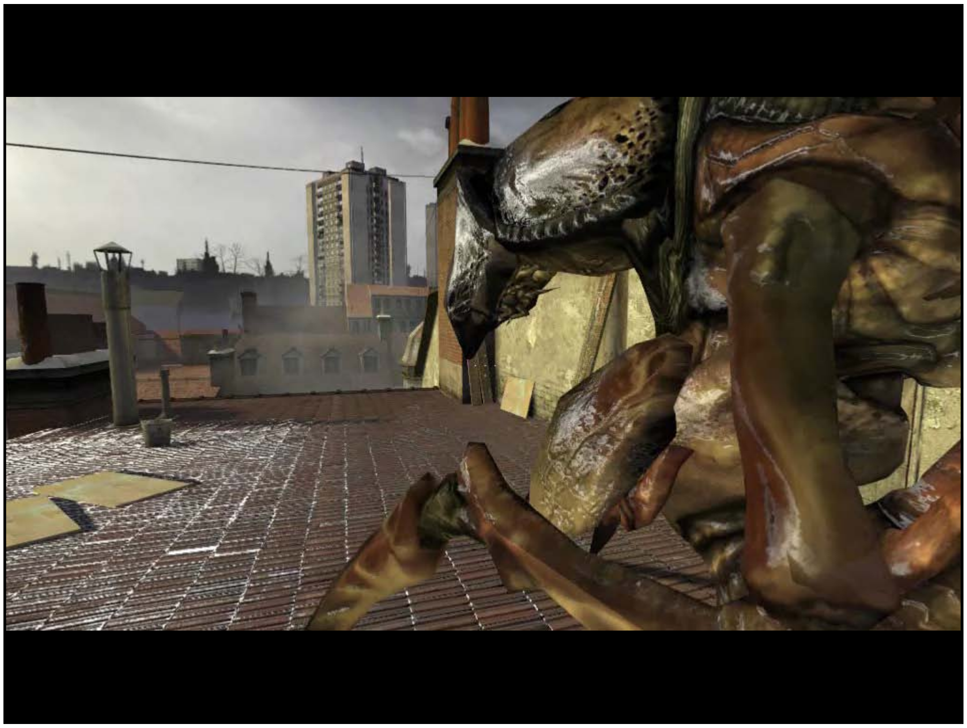

HDR Implementation #1

- Advantages:
  - MSAA worked
  - works on all DirectX 9 hardware
  - exposure is a simple scale in pixel shaders
- Disadvantages:
  - alpha blending and fog are tricky
  - shaders are complex
  - bilinear interpolation of HDR texture maps doesn't work

VALVE

HDR Implementation #2

- This is the first post-Half-Life 2 HDR experiment and was used for the first Lost Coast demos.
- This method uses two MRT buffers/textures for all HDR resources.
  - The first buffer stores visible data.
  - The second buffer stores overbright data.

VALVE

HDR Implementation #2

- Pluses:
  - Main motivation: alpha blending works.
  - works on all DirectX 9 hardware
  - bilinear interpolation works
- Minuses:
  - MSAA doesn't work: MSAA and MRT don't work together on current hardware
  - can't use the hardware fog
  - HDR textures, render targets, etc take twice as much space.

HDR implementation #3

- Used for E3 2005 Lost Coast demo
- HDR textures and render targets are all full floating-point RGB.
- this is the future
- Improved tone mapping
- HDR refraction

HDR Implementation #3

- Pluses:
  - generates the best data, we use this mode to generate our procedural HDR data
  - you get actual overbright data to bloom with
- Minuses:
  - Requires floating-point alpha blending
  - Performance not that good for supported hardware.
  - Tons of memory used for all resources.
  - no MSAA
  - Current hardware lacks precision to do fog

HDR Implementation #4

- Lost Coast and Day of Defeat will ship with this!
- Similar to 2 buffer version, but without over-bright data
- Store HDR textures in fp16 (linear color space) when filtered fp16 textures are supported.
- Store HDR textures in 4.12 linear color space otherwise.
- Had to implement this version along with previous version to support more hardware.

HDR Implementation #4

- Pluses:
  - MSAA works!!!
  - works on all dx9 hardware (runs well even on 9600)
  - best performance out of all of our methods
    - is a small performance hit over HL2's shipped LDR solution
  - doesn't take much memory (only extra memory is for HDR textures and bloom buffers)
  - fog interpolators work
- Minuses:
  - Refraction doesn't carry HDR data through

HDR and Authoring Content

- Only hand-authored HDR textures are skybox.
- Techniques for creating HDR skyboxes
  - Use a 3D rendering program
    - Many 3D packages read and write HDR formats
  - Use multiple exposures photo-reference
    - Use HDR Shop to build HDR images.
  - Hand-authored
    - HDR Shop + Photoshop
    - Currently HDR paint programs are inadequate. (including Photoshop CS2)

HDR and Authoring Content

- Lighting for HDR is different
- LDR: light for a single exposure
  - use supplemental fill lights to brighten dark areas
- HDR: light more like the real world, with dynamic exposure
  - Lost Coast uses one primary light, the sun.
  - Would not work without radiosity to bounce the light around.
- We actually light levels separately for LDR and HDR

HDR and Authoring Content

- Bloom amount and exposure range
  - Are settable on a per map or per map area basis by level designers to override the automatic settings.
- Within this set range we use auto-exposure to adjust the exposure level.

HDR and MODs

- When Lost Coast ships, HDR support in SDK will also be released
- MODs will need to author LDR and HDR lighting for maps, and create HDR sky boxes (or use ours)
- LDR and HDR lighting contained in the same BSP
- There's gameplay to explore for stealth and in exposure control (e.g. flash bangs)

HDR Challenges

- Lack of HDR Photoshop
- Non-orthogonal implementations
  - AA and FP
  - Texture compression doesn't work with HDR
- Performance in NVIDIA floating point path
- Necessity of an integer path because of ATI

Benchmarking and Visualization

- Driver versions
- Histograms
- Auto-exposure
- Exposure level
- Bloom scale
- Split screen
- Enabling floating point

HDR at Electronic Arts

Habib Zargarpour

High-Dynamic Range Imaging

SIGGRAPH2006

Need for Speed Most Wanted

Habib Zargarpour
http://www.zargarpour.net
Senior Art Director, Electronic Arts
Need for Speed Team

High-Dynamic Range Imaging and real time rendering

Purpose of this course: To familiarize with the process and issues of using High Dynamic Range (HDR) rendering in real time on a next generation game engine. We will be using the case study of Need for Speed Most Wanted for examples.

Main contributors: The entire team at EA Black Box contributed to this game. Most notably, specifically relating to the HDR rendering aspects on the Xbox 360, Gareth Jones was the main contributor to this work.

Although using specific examples, the principles will apply to any real time rendering application. We hope to bring to light pitfalls and unpredictable problems that we ran into along the way to help avoid such issues in the future.


Need for Speed Most Wanted
High-Dynamic Range Imaging and real time rendering
Xbox 360 with full HD 720p rendering and 4x Anti-Aliasing
Screen shot from Xbox 360 game engine shows normal and spec mapping and HDR bloom
3
SIGGRAPH2006

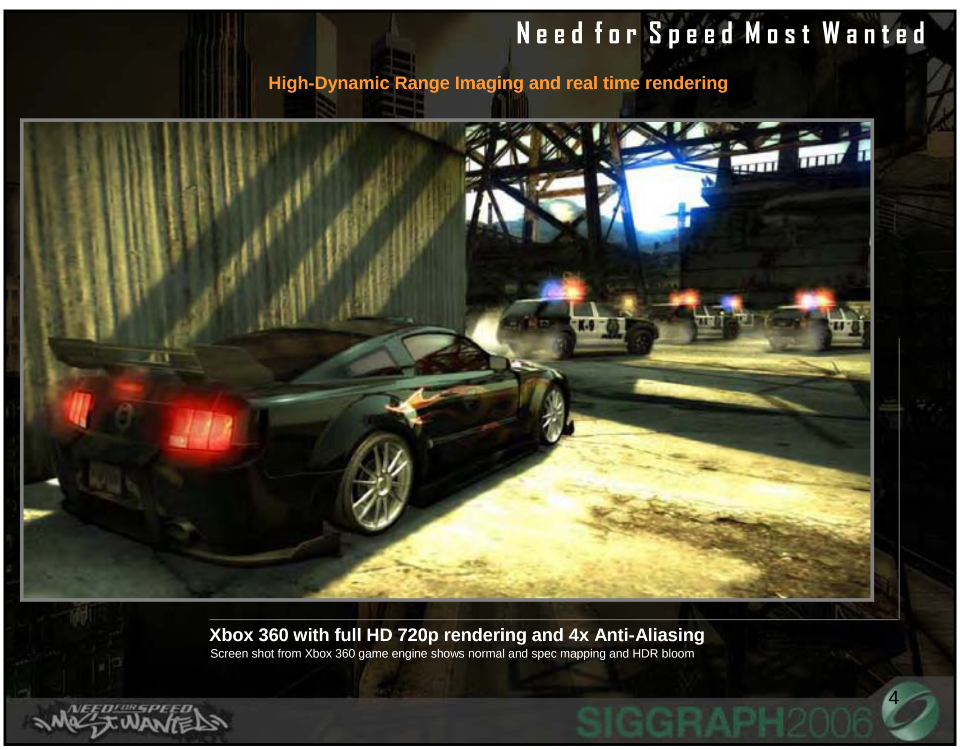
Need for Speed Most Wanted
High-Dynamic Range Imaging and real time rendering
Xbox 360 with full HD 720p rendering and 4x Anti-Aliasing
Screen shot from Xbox 360 game engine shows normal and spec mapping and HDR bloom
4
SIGGRAPH2006

Lighting of sky is done in conjunction with the atmosphere and fog with dynamic scattering

Screen shot showing Good sky exposure and bloom

- Sky texture in HDR buffer
- Histogram used to determine exposure
- Resulting bright areas are isolated for blooming.
  - Blooming depends on exposure amount
  - Exposure is adjusted in real time according to the brightness on screen

Bloom can also affect bright reflections on the surface of the car

Cast shadows also affect cars

Dynamic real time lighting Sun rises and sets in 30 minute cycles

All objects are lit in real time and cast shadows move with the sun

- Weather Effects and rain

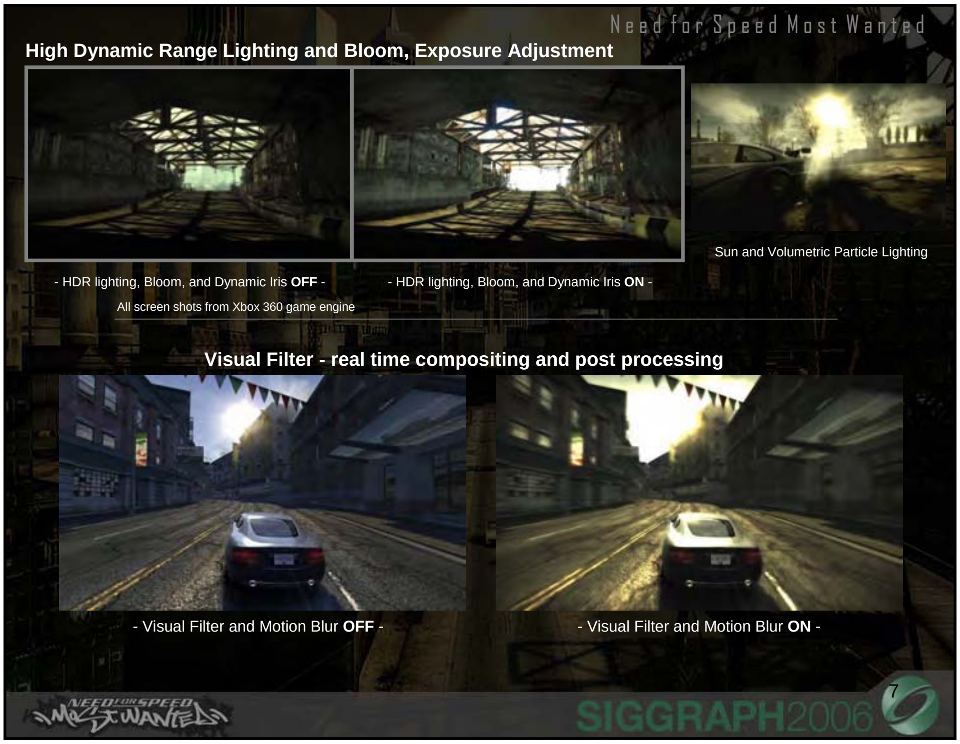
Need for Speed Most Wanted
High Dynamic Range Lighting and Bloom, Exposure Adjustment
Sun and Volumetric Particle Lighting
- HDR lighting, Bloom, and Dynamic Iris **OFF** -
- HDR lighting, Bloom, and Dynamic Iris **ON** -
All screen shots from Xbox 360 game engine
Visual Filter - real time compositing and post processing
- Visual Filter and Motion Blur **OFF** -
- Visual Filter and Motion Blur **ON** -
7
SIGGRAPH2006


Need for Speed Most Wanted
Tree Lighting
Exposure adjustment in the shade makes sky naturally blown out
Real time lighting on trees softly distributes the sunlight according to the sun position
Visual Filter - real time compositing and post processing
- Visual Filter OFF -
- Visual Filter ON -
8
SIGGRAPH2006

Need for Speed Most Wanted

High Dynamic Range Lighting and Bloom, Exposure Adjustment

Visual Filter - real time compositing and post processing

- HDR lighting, Bloom, and Dynamic Iris **OFF** -

All screen shots from Xbox 360 game engine

Need for Speed Most Wanted SIGGRAPH2006 9

Need for Speed Most Wanted

High Dynamic Range Lighting and Bloom, Exposure Adjustment

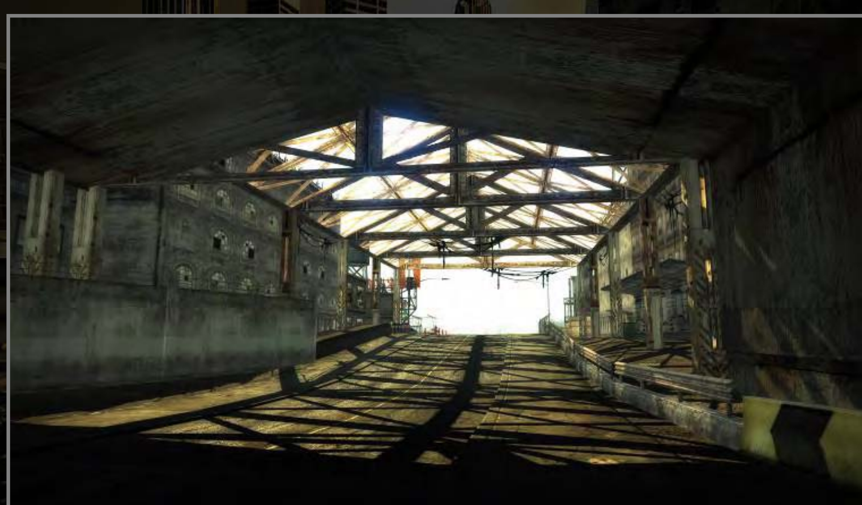

Visual Filter - real time compositing and post processing

- HDR lighting, Bloom, and Dynamic Iris **ON** -

All screen shots from Xbox 360 game engine

Need for Speed Most Wanted SIGGRAPH2006 10

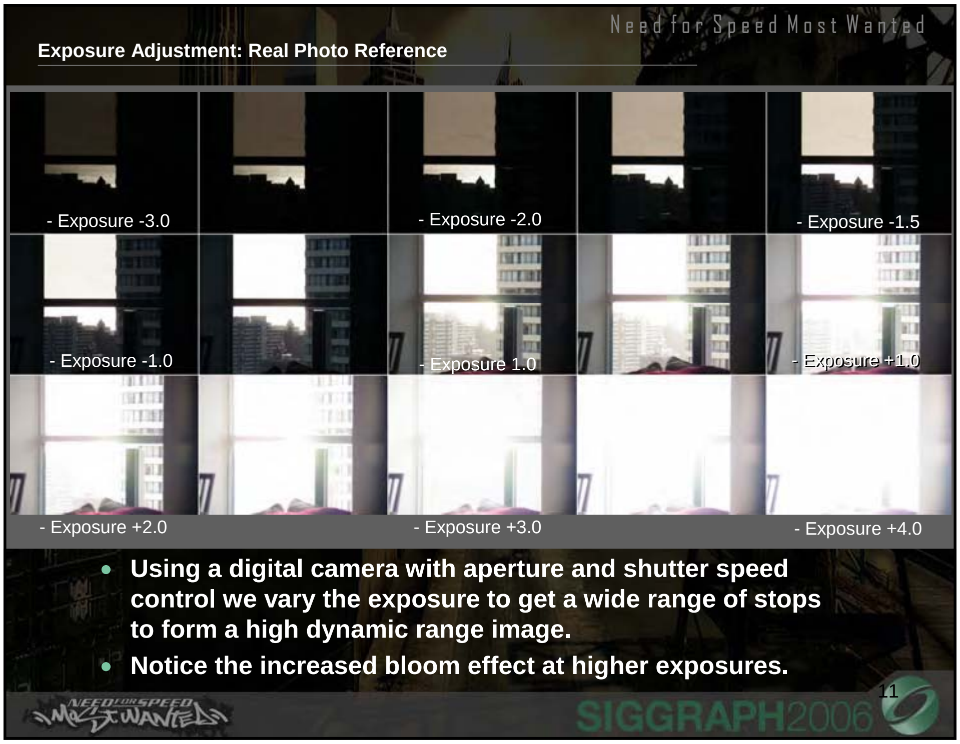
Need for Speed Most Wanted
Exposure Adjustment: Real Photo Reference
- Exposure -3.0
- Exposure -2.0
- Exposure -1.5
- Exposure -1.0
- Exposure 1.0
- Exposure +1.0
- Exposure +2.0
- Exposure +3.0
- Exposure +4.0
• Using a digital camera with aperture and shutter speed control we vary the exposure to get a wide range of stops to form a high dynamic range image.
• Notice the increased bloom effect at higher exposures.
11
NEED FOR SPEED MOST WANTED
SIGGRAPH2006


Need for Speed Most Wanted
Exposure Adjustment: Real time within the game engine
- Exposure 0.4
- Exposure 1.0
- Exposure 1.5
- Exposure 2.0
- Exposure 3.0
- Exposure 4.0
- Exposure 5.0
- Exposure 9.0
Increasing the exposure will automatically cause areas that are over-exposed to contaminate the area around them.
You can see the similar behavior compared to the real photo reference
12
NEED FOR SPEED MOST WANTED
SIGGRAPH2006

Radiosity Lighting

- Lighting levels need to be balanced between exterior and interior environments
- The relative ratio is critical for achieving exposure changes when going from one to the other

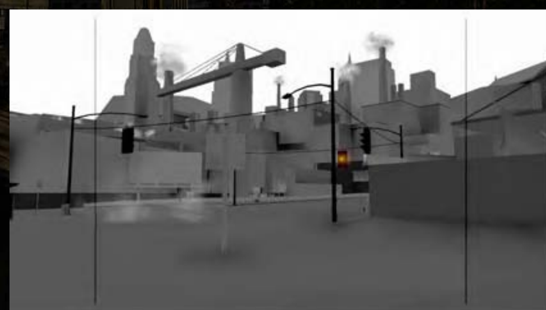

- Example Ambient occlusion lighting pass for an exterior environment

High Dynamic Range Lighting and Bloom, Exposure Adjustment

Lessons: Tunnel Lighting Exposure

- HDR tunnel lights reflect onto the car

- Exposure begins to auto-adjust

- **Interior lighting exposure has to be correct relative to the exterior lighting, otherwise no bloom would occur when exiting a tunnel.**

Using HDRI to select Exposure and Bloom

- HDRI and Bloom **OFF**

- HDRI and Bloom **ON**

- Because Blooming affects it's surrounding pixels it needs to be calculated after the exposure range has been determined and the overexposed area isolated.
- The dark part of the image above causes the virtual iris to open up causing the sky to be overexposed and bloom

Using HDRI to select Exposure and Bloom

- Exposure set for exterior

- Exposure set for interior

- Exposure set for exterior

- Exposure set for interior

Using HDRI to select Exposure

- Sunset with exposure adjustment **OFF**

- Sunset with exposure adjustment **ON**

- Dynamic Time of Day changes the level of light in an exterior environment
- With Auto-iris adjustment there is less concern about certain times of day or weather conditions becoming too dark or too bright
- The relative lighting between shade and sunlit areas, however, does need to be tuned accurately

Auto-Iris Exposure adjustment

- Auto-Iris adjustment **OFF**

- Auto-Iris adjustment **ON**

- Bloom **OFF**

- When the exposure is set close to mid-range levels, toggling Auto-iris adjustment will only slightly change the exposure
- Lens Bloom effect can be turned off separately

Need for Speed Most Wanted

HDRI and Interior lighting

- Exposure set for exterior

- Exposure Auto-adjusted for interior

- Example on the right shows dangers of letting the exposure get too bright
- Compression artifacts and texture imperfections that are normally too dark to notice become glaringly visible

- Exposure set too high

SIGGRAPH2006 19

Need for Speed Most Wanted

High Dynamic Range Lighting and Bloom, Exposure Adjustment

Pitfalls:

What is wrong with the car? Is the car too dark?

- HDR lighting, Bloom, and Dynamic Iris **ON** -

SIGGRAPH2006 20

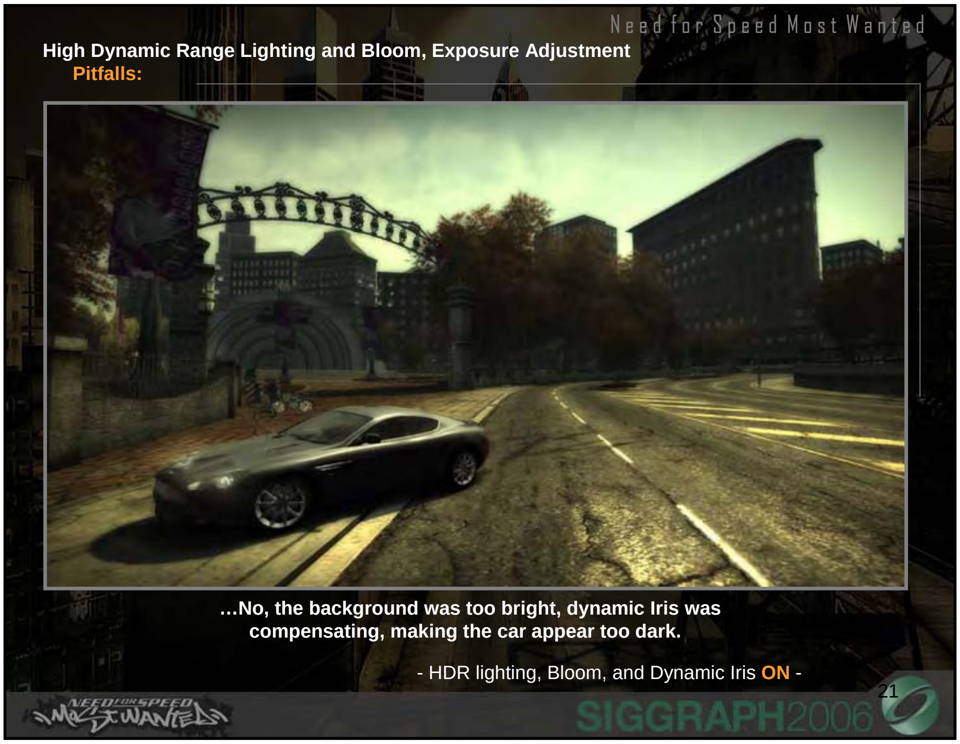
Need for Speed Most Wanted
High Dynamic Range Lighting and Bloom, Exposure Adjustment
Pitfalls:
…No, the background was too bright, dynamic Iris was compensating, making the car appear too dark.
- HDR lighting, Bloom, and Dynamic Iris ON -
21
SIGGRAPH2006


Need for Speed Most Wanted
Exposure Adjustment and Bloom Threshold
- Bloom Threshold 0.5
- Bloom Threshold 0.2
- Bloom Threshold 0.05
- Bloom Threshold 0.01
Reduction of the Bloom Threshold will cause more contamination of the dark areas from the bright sky
22
SIGGRAPH2006

Special Thanks

EA Blackbox Team

XBox 360 Rendering:

Greg D'Esposito
Gareth Jones
Colin O'Connor
David Lam

Producers:

Michael Mann
Larry LaPierre

Associate Art Director:

Eduardo Agostini

http://www.zargarpour.net

Contrast Processing

Rafal Mantiuk

Tone Mapping in Contrast Domain 1/3

- Operate on image contrast instead of luminance
- Similar to the gradient methods, e.g. [Fattal'02]

Mantiuk et al., A Perceptual Framework for Contrast Processing of HDR Images. TAP 2006.

Tone Mapping in Contrast Domain 2/3

Input Image

↓

Luminance to Contrast

↓

Tone Mapping

↓

Contrast to Luminance

↓

Output Image

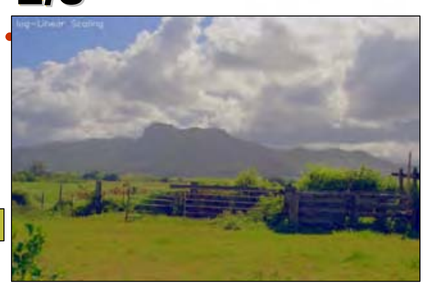

Tone Mapping in Contrast Domain 3/3

Contrast – differences between neighboring pixels on all levels of the Gaussian Pyramid

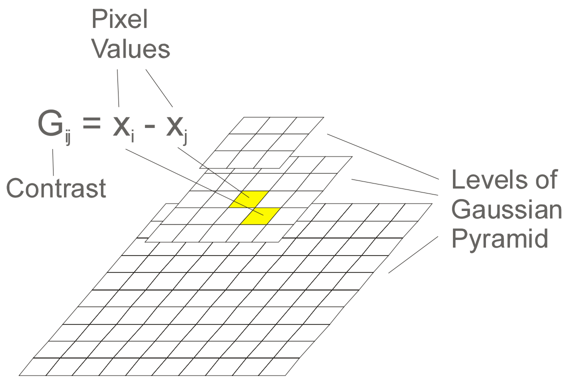

http://www.mpi-inf.mpg.de/resources/hdr/

Lightness Perception

Grzegorz Krawczyk

Lightness Perception in Tone Reproduction

Model the lightness perception corresponding to conditions in the real world.

The theory:
- An Anchoring Theory of Lightness Perception [Gilchrist et al. 1999]

Key concepts:
- Frameworks – areas of common illumination
- Anchoring – luminance → lightness mapping

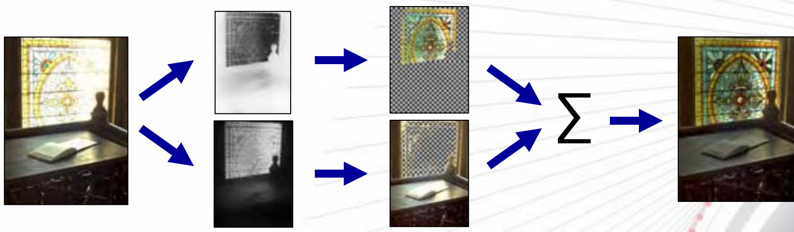

Lightness Perception in Tone Reproduction for HDR Images
Krawczyk et al., Computer Graphics Forum vol.24, 2005.

Frameworks

Figure: Frameworks of common illumination.

Frameworks allow for lightness estimation in complex images.

Perceptual organization:
- semantic grouping
- good continuation
- grouping of illumination
- proximity

Frameworks:
- defined by probability maps
- each pixel may belong to several frameworks

Anchoring

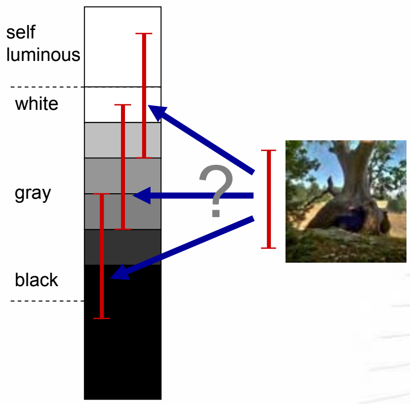

Figure: Perceived gray shades.

Mapping between luminance value and value on a scale of perceived gray shades.

Two possibilities:

- Anchoring to middle-gray
- Anchoring to white

Experimental evidence favors anchoring to white.

Estimation of Anchor

Anchoring to white rule:

- tendency of the highest luminance to appear white
- tendency of the largest area to appear white

Self-luminosity:

- small white disc surrounded by a large dark area appears luminous

Area related approach to anchor estimation:

- Remove the top 5% of pixels of highest luminance
- Measure the highest luminance of the rest of the pixels

Net Lightness Calculation

Shift original luminance (Y)

- according to local anchors (W)
- proportionally to belongingness (P)

$$L(x,y) = Y(x,y) - \sum_i W_i \cdot P_i(x,y)$$

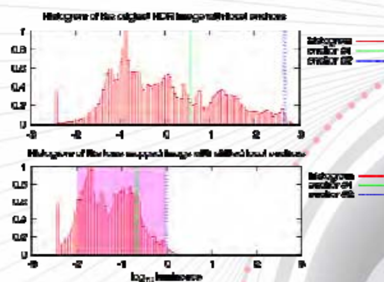

Tone Mapping: Tree

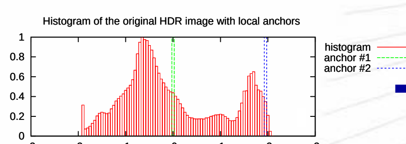

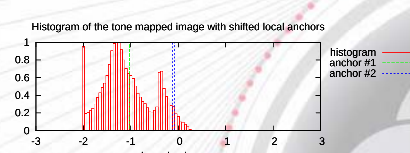

Tone Mapping: Desk

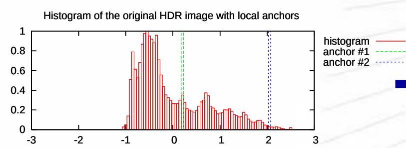

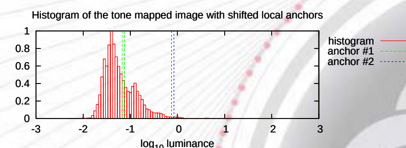

Conclusions

- Computational model of lightness perception theory
- Formalize method for extracting frameworks
- Application of the anchoring theory to tone reproduction
- Simulation of lightness constancy failure
- Lightness calculation in tone mapping