HDR SHOP

Contact
CHRIS TCHOU
USC Institute for
Creative Technologies
tchou@ict.usc.edu

PAUL DEBEVEC
USC Institute for
Creative Technologies

HDR Shop is a computer application (currently under development) designed to view and edit high-dynamic-range (HDR)[1] images: pictures that can capture a much greater range of light intensities than standard photographs or computer images. This approach is very useful for image-based lighting and post-render processing.

Photographs from traditional cameras do not record the amount of light over a certain level. All the bright points in a photo are white, which makes it impossible to detect any difference in intensity. The standard technique to acquire HDR images that capture this missing information is to take several photographs at different exposures (making each photo progressively darker, without moving the camera), until the bright lights no longer saturate. The sequence of photographs can then be analyzed to derive the light intensity of each point in the scene.

Whereas traditional image editors work with 8- or 16-bit images, HDR Shop is built from the ground up to work correctly with HDR images. All operations are done with linear floating-point numbers. In many cases, this simplifies the code, as well as providing more correct output.

For the purpose of real-time display, however, it is important to quickly convert linear floating-point images to 8-bit RGB with the appropriate gamma curve. The standard gamma formula involves an exponentiation, which is slow. In the interest of speed, we have found it useful to approximate this calculation by constructing a lookup table indexed by the most significant bits of the floating-point values. For common gamma values of 1.4 ~ 2.2, it suffices to use 16 bits (eight exponent bits and eight mantissa bits) to reduce the error below rounding error.

In addition to resampling, cropping, and mathematical operations, HDR Shop also supports transformations among most common panoramic formats, facilitating the use of HDR panoramas in image-based lighting[2]. HDR Shop can also automatically export a low-dynamic-range (LDR) copy of any image to an external image editor. Changes to the LDR image are then incorporated into the HDR image, so existing tools can be used to modify HDR images.

See also: www.debevec.org/HDRShop

References
1. Debevec, P. & Malik, J. (1997). Recovering high dynamic range radiance maps from photographs. *Proceedings of SIGGRAPH 97.*
2. Debevec, P. (1998). Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Proceedings of SIGGRAPH 98.*

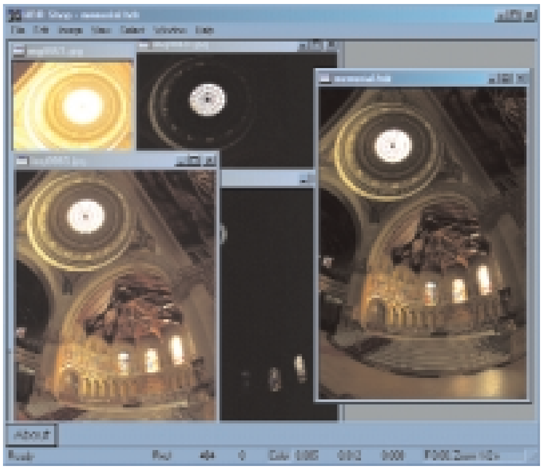
Figure 1. In HDR Shop, a sequence of low-dynamic-range images (left) can be compiled into a single high-dynamic-range image (right).

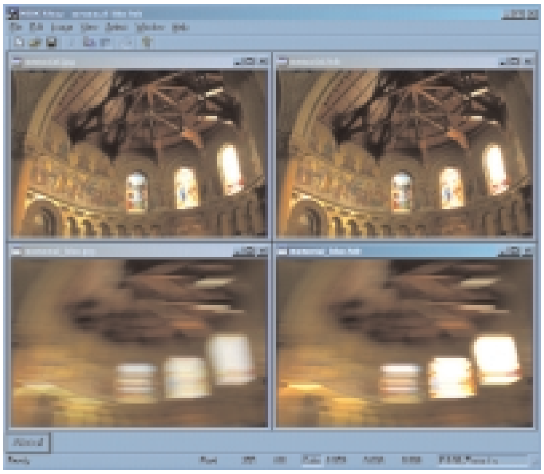
Figure 2. Comparison of HDR Shop's horizontal motion blur on a low-dynamic-range image (left) vs. a high-dynamic-range image (right).


Figure 3. St. Paul's Cathedral panorama, originally in cube-map format (left), converted in HDR Shop to latitude-longitude (upper right), mirrored ball, and light probe formats (lower right).

LIGHT STAGE 2.0

Contact
TIM HAWKINS
Institute for Creative
Technologies
University of Southern California
timh@ict.usc.edu

JONATHAN COHEN
CHRIS TCHOU
PAUL DEBEVEC
Institute for Creative
Technologies
University of Southern California

At SIGGRAPH 2000, we presented an apparatus for capturing the appearance of a person's face under all possible directions of illumination. The captured data can be directly used to render the person into any imaginable lighting environment, and can also be used to build photo-real computer graphics models that capture the unique texture and reflectance of the face. We have recently been developing the next generation of this lighting apparatus, which we call Light Stage 2.0.

Light Stage 2.0 is a much faster and more precise version of its predecessor[1]. The original device allowed a single light to be spun around on a spherical path so that a subject could be illuminated from all directions, and regular video cameras were used to record the subject's appearance as the light moved. This system had two major problems. First, since the light was moved around by pulling on various ropes, it was hard to be sure what the precise location of the light was at any given time. Second, because the device could not be spun very fast, and because of the limit of 30 frames per second imposed by the video cameras, it took over a minute to do a data capture. Since the subject must remain still during the data capture, this meant we could only capture people in very passive expressions, and even then multiple trials were often needed.

With Light Stage 2.0 (shown in Figure 1), we can capture all of the different lighting directions much more rapidly, with only a single rotation of a semicircular arm, and with greater accuracy. Thirty strobe lights arrayed along the length of the arm flash repeatedly in rapid sequence as the arm rotates. High-speed digital cameras capture the subject's appearance. This allows all directions of illumination to be provided in about four seconds, a period of time for which a person can easily remain still. It is also much easier to capture facial expressions that would be very difficult to maintain for an extended period of time (smiling, frowning, wincing, etc.).

We are currently working on integrating geometry capture to provide a complete model of the subject. For this, we use digital LCD projectors to project different structured patterns onto the subject, quickly recording the appearance of the subject under each of the patterns with our high-speed cameras. From these structured-light data, the geometry of the subject is easily recovered. These data together with the reflectance data may provide more complete and photo-real models of faces than ever before.

In the next few months, we will be researching new ways of analyzing the large amount of reflectance field information captured in a Light Stage 2.0 scan and adapting the datasets for use in facial animation. We would also like to make our capture process even faster, with the goal of being able to capture both geometry and reflectance information in about five seconds. Our future plans include new prototype lighting devices that will allow similar datasets to be captured many times a second. This will allow an actor's performance to be recorded and then rendered photo-realistically into virtual environments with arbitrary lighting, where the performance can be viewed from arbitrary angles.

Another approach is to directly illuminate an actor with light sources aimed from all directions whose intensity and color is controlled by the computer. In this case, if the incident illumination necessary to realistically composite the actor into a particular scene is known in advance, the actor can be filmed directly under this illumination.

Reference
1. Debevec, P., Hawkins, T., Tchou, C., Duiker,H.-P., Sarokin, W., & Sagar, M. (2000). Acquiring the reflectance field of a human face. In *Proceedings of SIGGRAPH 2000.*


Light Stage 2.0, with seated subject.

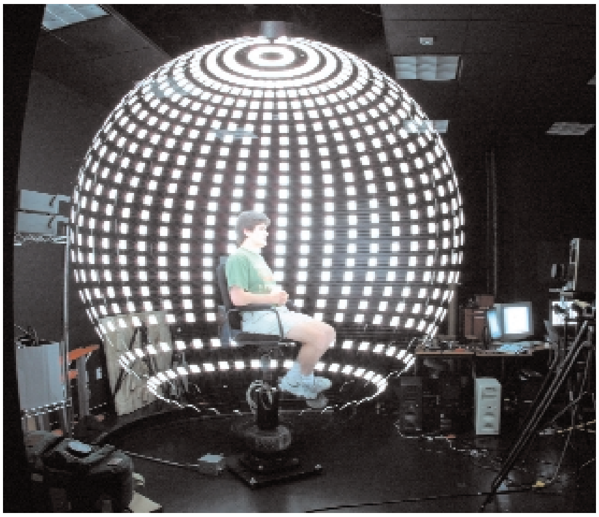
A 10-second-exposure photograph of Light Stage 2.0 acquiring a 4D reflectance field dataset of the subject's face.

A Real-Time High Dynamic Range Light Probe

Contact
Jamie Waese
USC Institute for
Creative Technologies
13274 Fiji Way, 5th Floor
Marina del Rey, California 90292
USA
waese@ict.usc.edu

Paul Debevec
USC Institute for
Creative Technologies

In order to successfully composite computer graphics elements into live action scenes it is important that the lighting of the CG object match the lighting of the scene into which it is being composited. One technique that has been used to reproduce the incident illumination in a live-action scene is to acquire a high-dynamic-range photograph of a mirrored ball placed in the scene and then use this light-probe image as a source of illumination for image-based lighting.[1]

Previous Work

Currently, in order to create a high dynamic range image of a mirrored ball one must take an iterative series of photographs with the exposure value of each image being stopped down by a given increment from the exposure value of the one before. Later, each of the images are assembled into a single high dynamic range image using a program such as HDR Shop[3]. If an artist wished to accurately illuminate a CG object traveling through a complex lighting environment, it would be necessary to capture these iterative photographs at numerous locations (ideally at every frame) along the object's path. Clearly, this would be an ambitious task.

Technique

One solution for creating a real-time high-dynamic range light probe is to develop a system in which multiple exposures of the same image can be captured within a single video frame. We did this by modifying a five point multi-image filter (a faceted lens that is commonly used to create photographic kaleidoscope effects), and applying successively increasing values of neutral density gel to four of the five facets of the filter (3, 6, 10 and 13 stops). This modified filter effectively produces a single image that is divided into five identical regions, with the center region capturing a "direct" view and the four outer regions stopped down to their respective exposure values. This modified filter is placed on a video camera that is mounted along with a mirrored ball on a span of angle iron (see Figure 1).

Assuming the relation between the camera and the ball never changes, the light probe only needs to be calibrated once. To compensate for the angle shift introduced by parallax effects from the facets of the multi-image filter, one can compute the arctangent of the distance between facets divided by the distance between the lens and the silver ball. By determining the number of degrees each facet is offset from the center, we are able to warp each region of the filter according to the direction space of its view of the ball. In our case, each facet's view of the ball was computed to be 2.7 degrees off from center.

More accurate calibration can be done with the help of a light stage,[2] which provides a "master key" for factoring out lens distortion and imperfections in the mirrored ball. However, we found that simply computing the pixel shift and then overlapping each region of the filter was sufficient for assembling a usable image.

In order to capture high dynamic range light probe data at every frame along a path, one presses "record" on the video camera and carries the light probe along the desired path. A computer program then imports each recorded frame, isolates the five distinct images in the frame, aligns them according to predetermined calibration data, and then assembles the aligned images into a high dynamic range omnidirectional measurement of incident illumination.

Results

Figure 2 shows a raw, unprocessed image from the light probe.

Figure 3 shows several exposures of a high dynamic range image that were assembled from a single light probe frame.

Figure 4 shows a CG object, lit with captured light from the real time high dynamic range light probe.

Conclusion

This new technique will permit artists to composite CG objects into dynamic complex lighting environments, accurately reproducing high dynamic range lighting parameters for each frame. In the future, this technique would benefit from greater precision in applying the neutral density gels to the multi-image filter, a smaller camera rig, and higher resolution video cameras.

References

1. Debevec, P. (1998). Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings SIGGRAPH 98*.
2. Debevec, P., Hawkins, T., Tchou, C., Duiker, H.P., Sarokin, W., & Sagar, M. Acquiring the reflectance field of a human face. In *Proceedings SIGGRAPH 2000*.
3. Tchou, C. & Debevec, P. (2001). HDR shop. *SIGGRAPH 2001 Conference Abstracts and Applications.*

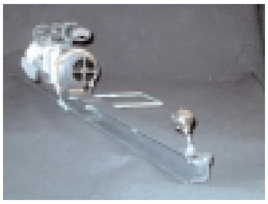

Figure 1. A real time high dynamic range light probe.

Figure 2. Five exposures of a mirrored ball in a single image.

Figure 3. Five exposures of a high dynamic range image captured in a single frame.

Figure 4. A CG model that is synthetically illuminated with light captured with the real time light probe.