Achieving a Multi-spectral Virtual Reality

Franz Herbert and Jo Kirkenaer
Integrated Color Solutions, Inc., Encinitas, California, USA

Oliver Unter Ecker
Kula, Hawaii, USA

Abstract

With the advent of extremely powerful computers and graphics processors the worlds of 2D contextual images and 3D dynamic renderings are merging. Simulating a virtual environment, under multiple illuminants, and providing the capability to render realistic objects with selectable colors, special effects, and other color appearance phenomena, presents new challenges. The color and appearance attributes applied to an object are based on physics and the material properties of the phenomena. We combined multi-spectral imaging, BRDF technology and GPU-based rendering to meet these challenges and advance the state of the art.

A physically-based visualization system is presented which realizes photo-realistic renditions of 3-dimensional objects with complex real-world material characteristics in captured real-world panoramic environments. The system runs in real-time and permits interactive manipulation of material parameters, change of environments or environment characteristics such as current illuminant, with immediate visual feedback.

For a real-time system unprecedented visual accuracy is achieved through the use of multi-spectral, high-definition, and high-dynamic range descriptions of material reflectance and environments, and through a hybrid shading subsystem which draws on a number of recent rendering techniques including frequency-space reflection mapping based on Spherical Harmonics for accurate BRDF rendering.

The system's real-time abilities are enabled by recent programmable graphics card hardware and extensive use of programmed shaders. This paper describes the system along with some of the methods used to acquire material characteristics and environments.

1 Introduction

Real-world materials exhibit surface characteristics that are not captured satisfactorily by prevalent 3-D rendering (shading) technologies. BRDF-based rendering has long been identified as the technology to bring true realism to 3-D graphics but was until recently a forbiddingly expensive option and out of reach for real-time applications. Only a select few CPU-based academic rendering systems permitted BRDF-based rendering.

Recent developments in graphics hardware technology and theoretical work on rendering have almost over night enabled BRDF-based rendering on consumer class hardware. Graphics cards have morphed from fixed-function devices into programmable graphics processing units (GPUs), whose processing power can exceed a CPU's by magnitudes for certain tasks.

Figure 1. Rendering at daylight, fluorescent, and tungsten

Early real-time 3-D graphics was limited to simple light sources such as point and spot lights for lighting of objects. A much more realistic option is given by using captured panoramic images (environments) as light source, where every pixel is regarded a mini light. Environment maps, then, represent incident illumination and were used in real-time rendering to provide initially hard (mirror) reflections and later, in prefiltered forms, precalculated specular and diffuse reflections but were still not suitable for general BRDFs.

As BRDF-based rendering has been the “holy grail” rendering technique for some time and researchers sought to make it practical despite a fundamental lack of processing power, a variety of compromise techniques have been created, which achieve speed-ups through the use of more compact, if lossy, representations at the cost of reduced accuracy, introduced artifacts, and/or restrictions to only certain types of BRDFs.

The straightforward table-based approach, which implies a 3-D table for isotropic BRDFs or a 4-D table for more general anisotropic BRDFs, is free of such problems but proved difficult to optimize for use with environment lighting. Even with dramatically faster processing a theoretical break-through was still required to make high-quality BRDF rendering practical. Several recently suggested frequency-space techniques at last begin to fit the bill. One such technique, Ramamoorthi et al.’s[11] Spherical Harmonics Reflection Maps, forms the basis of our system’s 3-D BRDF rendering.

Full 3-D BRDFs are usually acquired using a goniospectrophotometer. We have developed our own vision-based acquisition system (not described in this paper) but accept a variety of input formats, incl. ASTM E1392-96.

While full 3-D BRDFs are desirable for static materials such as the gonioapparent “special effects” ingredients embedded in plastics, such as metallic flakes or iridescent or pearlescent pigments, a dynamic representation is desirable to describe a material which depends on parameters that a user will manipulate. For example, polymers whose appearance depend on thickness, finish and other parameters in a design system.

Our system implements an advanced hybrid shading model which incorporates both synthesized and measured BRDF rendering. Acquired optical material measurements, tabulated simplified 1-D and general isotropic 3-D BRDFs feed this shading model.

All shading occurs through GPU-based shader programs and runs two to three magnitudes faster than a CPU-based implementation would.

All imagery and BRDF data going into the system is spectrally-acquired and kept in a database which allows access of colors and BRDFs by name, delta-E specification given a target L*a*b* color.

In coupling best-of-breed rendering technologies with a comprehensive application feature set and multi-spectral representations of all color going into the system, we achieve an unprecedented level of visual result accuracy and great suitability for a wide range of applications.

2 Related Work

Our shading work relates to three areas with a long history in computer graphics: environment mapping, BRDF rendering, and use of Spherical Harmonics (SH).

Environment mapping was first proposed by Blinn and Newell,[1] and prefiltering of environment maps was pioneered by Greene.[4]

Various angular (spatial) domain environment map prefiltering techniques, such as by Heidrich et al.[5] and Kautz et al.,[7] have since been proposed which generate static maps and were limited to Phong[9]-like BRDF models.

General BRDF-based rendering appeared in off-line render systems such as Greg Ward’s Radiance[12] but real-time techniques, such as one by Kautz and McCool,[6] were limited to non-complex illumination.

A principal key insight that paved the way for efficient environment-illuminated BRDF-based rendering came from Cabral et al.[2] with the observation that the integration over the visible hemisphere can be replaced by a dot product when transporting the problem into frequency space.

Ramamoorthi and Hanrahan[11] took a rigorous and innovative step when they applied their earlier developed signal processing framework and analysis[10] to the problem and proposed Spherical Harmonics Reflection Maps (SHRMs). Sharing the view-dependent map trait of earlier work by Cabral et al.[3] SHRMs represent a relatively compact representation and permit a tunable degree of accuracy.

Kautz et al.[8] proposed another SH-based approach which supports general 4-D BRDFs and runs entirely on the GPU but is restricted to (very) low-frequency environments.

3 Materials

3.1 Polymer Characterization

Measuring translucent samples poses a serious challenge. We decided to develop an empirical approach. Since nearly all translucent substrates today are made from various polymers, we had special chip samples made for characterization. Each substrate was made at three different thicknesses, 0.03”, 0.06” and 0.120”. Each of these samples was made with several levels of opacifier, TiO_2. Adding TiO_2 increases the opacity of the samples as well as in some cases changes the apparent color from a clear or yellowish polymer base to white.

Characterization of polymer specimens was done by measuring the spectral reflectance and transmission using a Datacolor Spectraflash 600 with spherical illumination, as well as the 5 angle gloss measurements. These data were used in the shading model. It is important that the translucency of the sample is preserved in the representation on display, as well as the surface gloss and texture. The latter can be controlled through the software as long as the substrate characterization is correct.

We measured each of these samples, thickness and TiO_2 concentrations spectrally over both black and white backing. Example responses from spectral reflectance measurements over black are shown for HDPE in Figure 5 as a function of TiO_2 concentration.

The samples were also measured for full spectral transmission. Calculated transmissive values are also shown in Figure 2. The y axis represents the green channel at D65.

From the spectral measurements we built a matrix of color values as a function of thickness and TiO_2 concentration. The RGB values to display the sample are then derived from the spectral values contained in the matrix, L*a*b* values calculated based on the selected illuminant, and the L*a*b* → RGB conversion done through the corresponding profile.

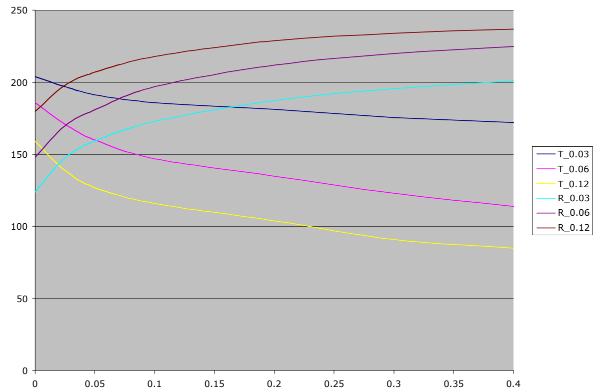

Figure 2. Lightness Response of HDPE

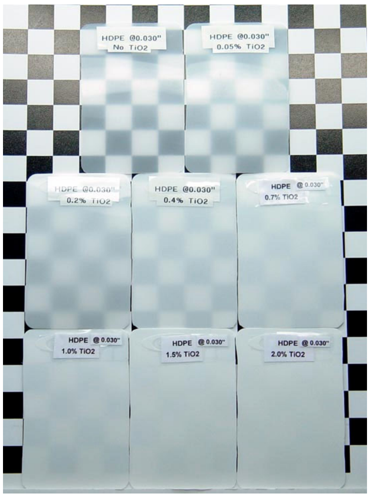

Figure 3. HDPE chips with different TiO_2 content

3.2 BRDFs

The BRDF (Bidirectional Reflectance Distribution Function) describes the way a material reacts to light. Every combination of incident and viewing angle is recorded across a hemisphere. Any shading done in the past was really a simulation of what the real BRDF would have achieved.

5 angle spectrometers measure 15°, 25°, 45°, 75°, and 110° in the incident plane. For many materials these 5 angles suffice to characterize the material. Figure 4 shows a gonioapparent material, Chromaflair (from Flex Products, Inc.), which is only properly characterized by a complete 3-D BRDF. 5 angle instruments are very popular in the automotive paint industry.

Figure 4. Red-Gold Chromaflair pigment (Photographs)

4 Multi-Spectral HDR Environments

We used a Spheron panoramic camera to obtain highly defined (up to 12k x 6k), high-dynamic range (32 bit floating-point) panoramic environment images.

The HDR image gets expanded to a complete spectral image with 31 or more wavelengths. A digital camera target with 576 patches is used to develop a set of polynomial equations of the 5th order that map RGB plus the higher order terms to the wavelengths of the measured target patches. Using this methodology an average deltaE94 of 0.77, peak 3.16 is achieved for all patches of the target. For higher accuracy a second capture can be performed using a light blue absorption filter.[13,14] Obviously this introduces problems with registration, white balancing, etc. In an ideal scenario the camera system would capture 6 or more channels with one exposure. The success of this method depends directly on the camera target and its “spectral diversity.”

The camera target was measured using a Barbieri Spectro 100xy. This spectrometer measures spectra containing 41 spectral bands, from 380 nm to 780 nm in 10 nm increments.

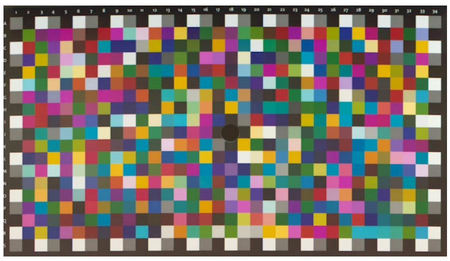

Figure 5. Digital Camera Target

The complete spectral panorama gets stored as a 41 component floating-point TIFF file. It gets converted to RGB using the user selectable illuminant and monitor calibration and ICC profile for that specific illuminant. The calibration sets the white point of the entire monitor for complete adaptation of the observer. The ICC profile converts illuminant-relative L*a*b* values to final monitor RGB values. The calibration ensures a perceptually linear behavior of the monitor regardless of the illuminant. Diffuse and specular environment maps get recomputed based on the current illuminant.

5 Shading Subsystem

Our shading subsystem implements a hybrid shading model which uses SHRMs for rendering of static measured 1-D and 3-D BRDFs and combines these with a physically-motivated dynamic model which approximates the BRDF of a polymer-with-glossy-finish material, taking as input the measurements obtained in our Polymer Characterization (3.1) step. The measured BRDFs are used for static "special effects" ingredients such as metallic flakes or iridescent, pearlescent pigments. Environments serve as the sole source of lighting in this system.

5.1 Shading Model

The shading model uses the following sum

$$I_{result} = I_d + I_{effect} + I_r F_r + I_{trans}(1 - F_r) \quad (1)$$

to incorporate these individual contributions

$$I_d = E_d D_{polymer} Tint_{polymer} (1 - Mod_{effect}) \quad (2)$$

$$I_{effect} = E_{effect} Mod_{effect} \quad (3)$$

$$I_{trans} = E_{trans} Col_{trans} \quad (4)$$

$$I_s = (E_s M_s + E_{mirror}(1 - M_s))\ gloss \quad (5)$$

where
I_{result} is the resulting intensity (pixel shade value);
F_r the Fresnel reflection coefficient, a function of the angle between eye and surface normal vectors;
I_{trans} the transmitted intensity;
I_d the reflected diffuse intensity;
E_d the Lambertian, or ideal diffuse, reflected radiance;
$D_{polymer}$ the measured, diffuse color of the polymer, a function of material thickness and TiO_2 load;
$Tint_{polymer}$ a tint color from a spot color database;
Mod_{effect} a material (special effect)-specific modulation map;
I_{effect} the reflected effect intensity;
E_{effect} the effect reflected radiance;
I_{trans} the transmitted intensity;
E_{trans} the transmitted radiance;
Col_{trans} the transmissive color, a function of material thickness and TiO_2 load;
I_s the reflected intensity;
E_s the specular reflected radiance;
M_s a heuristically-determined mirror reflection interpolation factor, designed to "sharpen" the overall reflection at near-grazing viewing angles, a function of the angle between eye and surface normal;
E_{mirror} the mirror reflection;
gloss the overall glossiness factor.

All the E_x values are provided thru prefiltered maps. All "function-of" parameters are realized through look-up tables, where the look-up frequency is per-pixel if the variable is an angle, and per-frame otherwise. All noted colors start out as spectral colors but are transformed into device-RGB for use in shading.

5.2 Render Algorithm

The pseudo code for the overall render algorithm is

```
for each frame
  for each material group
    if illuminant (il) or environment (env) changed
        color manage env for il (see 4.1)
        SH-project env
        create diffuse and specular maps for env
        If SHRM for il does not exist
            create SHRM for env and BRDF
    if viewpoint (vp) changed
        extract reflection map from SHRM for vp
    if glossiness (gloss) changed
        prefilter specular map for gloss
    set dynamic parameters into shaders
    setup render state (textures, shaders, etc.)
    render
```

(A persistent caching scheme, which is tied into this algorithm but not included here for clarity, realizes map reuse.)

The following subsections elaborate on the map-related action items in this algorithm.

5.3 SH-Projection

Spherical Harmonics form a function basis analogous to the Fourier basis, except on the sphere. Any square-integrable function defined over the unit sphere can be approximated to an arbitrary degree of accuracy through a SH series of appropriate order.

All our prefiltering operates in frequency space and SH-projection transforms an (HDR) environment from its native latitude/longitude angular format into SH coefficients in preparation for prefiltering. The order of the transform specifies how many SH terms will be used and how accurately the environment will be represented.

We use order 100 projections (implying 101^2 = ~10k coefficients) for environments, which allows for use of relatively highly-specular BRDFs in the prefiltering phase. BRDFs are projected at an order that depends on their highest frequency.

5.4 Prefiltering

Prefiltering in angular space entails the convolution of an environment with a BRDF and is of quadratic complexity per pixel. This is replaced in frequency space by a dot product of SH-projected environment and BRDF and is of linear complexity per pixel.

We use prefiltering to compute Lambertian and Phong BRDF reflection maps which are used in the approximating component of our shading model, to compute reflection maps for our 1-D measure BRDFs, and as a step in the creation of SHRMs for our 3-D BRDFs.

5.5 SHRM Creation

SHRMs are (angular-space) cubemaps which contain for each pixel SH-coefficients, which encode view-dependent reflection maps (the prefiltering result). The angular resolution of our SHRMs is typically 64x64x6, and the SH order typically between 2 and 10.

We create global SHRMs, assuming a distant viewer.

5.6 Extracting Reflection Maps from SHRMs

Static (view-independent) reflection maps are obtained from SHRMs by simply expanding the series for each coefficients-bearing "pixel" in the SHRM, where the Spherical Harmonics used in the expansion are indexed by the given viewpoint in spherical coordinates.

6 Display Subsystem

Precise calibration and profiling of the display is critical for accurate display of color. The objective is to present the color representation of a sample on the computer display, and obtain a visual color match to the original specimen displayed in a viewing booth next to the computer display. Expanding this concept to virtual reality implies that the display needs to be calibrated to any specified white point or matching illuminant. Since calibration is done spectrally, it is possible to achieve this multi-spectral calibration.

Standard calibration and profiling software calibrates the display to either a gamma of 1.8 (a traditional value in the graphic arts), or to a gamma of 2.2, more native for CRT display, according to

$$Y=((G/255^{Gamma})*Y_{max} \qquad (6)$$

Where G is the input gray value (0 – 255), and Gamma represents either 1.8 or 2.2, and Y_{max} represents the maximum brightness of the display at the given white point. However neither of these gamma values gave the desired color match required.

Based on visual appearance and display calibration, we modified the calibration gamma to match a linear L* as in equation 2. This is shown in Figure 9 and compared to gamma values of 1.8 and 2.2.

$$Y=((((G/255)*100+16)/116)^{3})*Y_{max} \qquad (7)$$

Figure 6 also shows normalized Y values from Apple Cinema and Studio Displays with the predicted Y_{pred} curve. Measured values show excellent correlation to the predicted value.

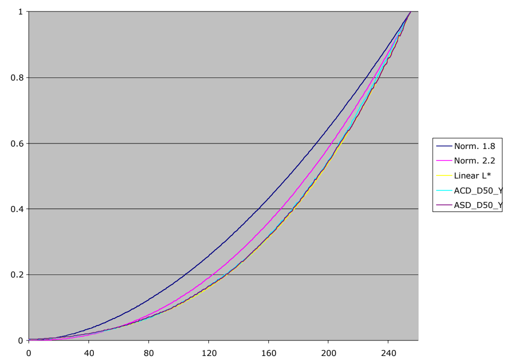

Figure 6. Display Gamma

We also found that the white point calibration contributes significantly to the perceived color, not only at maximum brightness, but also throughout the whole tonal range. To achieve constant white point, our display calibration software needed to iterate at multiple points across the complete tonal range, and verifying and validating on several points after completion. We added up to 15 points, with more points used in the highlights than in the shadows, and the ability to add more points where required.

Figure 10 shows chromaticity values (x, y) as a function of 8-bit gray value (0-255). This figure is for D65 calibration, and the straight, horizontal lines are the ideal D65 x and y coordinates. Note that in the shadows, below Y values of approximately 15 cd/m^2 the values fall off due to lack of instrument sensitivity in the shadows. We found that only spectrometers exhibited enough sensitivity in the shadow region to accurately calibrate the display to the given white point throughout most of the tonal range. The colorimeters tested were not able to hold both x and y below approximately the quartertone.

Because of the lack of instrument sensitivity in the shadows, we do not try to adjust the black point. Between absolute black and the first control point, the color is the natural color of the display, though generally too dark for the eye to see.

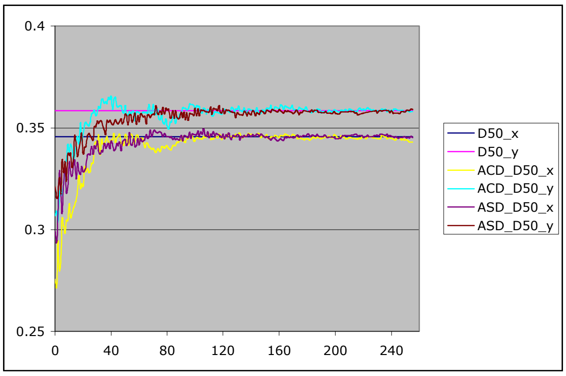

Figure 7. White Point Chromaticities

Since calibration is done to a given white point based on its chromaticity coordinates, the display can be calibrated to (almost) any white point by entering the actual chromaticity coordinates of the white point being matched. Our standard calibration is done at Daylight D65, indoor fluorescent F2, and tungsten Std A illuminants. For the graphic arts we also calibrate to Daylight D50.

Profiling is done as the software measures and characterizes the current state of the display after calibration. A standard LUT type ICC profile results, containing the calibration information for the Video Display Card.

7 Implementation

Speed and visual fidelity were the goals for our implementation. Panoramic transformations and all SHRM generation and prefiltering are carried out on the CPU while all rendering happens on the GPU.

The system is multi-platform and runs under Mac OS X, Windows 2000/XP, and Linux.

7.1 Fast Environment Transforms

Panaromic transformations, such as from latitude/longitude to cubemap format, are carried out through CPU code that draws on a vector processor-optimized (AltiVec on PowerPC processors, SSE2 on IA32 processors) pixel access library which is capable of bilinearly filtering samples at speeds that are limited by quad word-accessed memory read/write bandwidth.

Spectral-to-L*a*b* conversion of environments with concurrent bit-depth and sample count reduction (from spectral (31 bands) 32-bps IEEE floating-point to L*a*b* (3 samples) 8-bit) uses built-in multiply-accumulate vector machine instructions for even higher speeds, close to the memory read bandwidth.

The transform code is multi-threaded and will take advantage of multiple processors.

7.2 GPU Programs

Shaders, comprised of vertex and fragment programs, have been implemented which support the system on 2nd generation (ps.1.4-compatible, such as ATI Radeon 9000) and 3rd generation (ps2.0-compatible, such as ATI Radeon 9700/9800 and NVIDIA GeForce FX series) GPUs. Hardware limits on the number of available texture coordinate sets mandated the use of multi-pass rendering to implement our shading model, even on the latest GPUs.

Forty-two shaders were written and are dynamically employed by one to four rendering passes (the exact count depending on the configuration of the render setup and shading parameters such as percentage of transparency of the polymer and the capabilities of the GPU).

Each pass takes advantage of multi-texturing, typically specifying six (maximum on 2nd generation GPUs) textures. Aside from environments and prefiltered maps, conventional modulation and decal textures, bump maps, and LUTs for physical parameters (such as for the Fresnel coefficients) are passed into the shaders as textures.

7.3 SH-Projection Details

SH-projection occurs on the CPU with an implementation that uses quad-extended double (320 bit) precision to allow the generation of order 100 SH projections.

We successfully countered ringing tendencies in projections by sampling the angular-space source (environment or BRDF) with a Gaussian kernel and super-sampling the coefficient matrix used in SH rotation.

8 Results

We report on speed and achieved visual accuracy of our system.

Interactive frame rates are achieved for models with polygon counts exceeding 100,000 and containing multiple material assignments. Changes to dynamic parameters (such as material thickness, color, TiO2 load) are part of the model and incur no speed penalty. Changes that trigger prefiltering (such as changing into a different virtual environment) incur a penalty which is minimized through the use of caching, optimized image transforms, and frequency-space prefiltering.

On a dual-processor 1.0 GHz (x2) G4 computer, we achieve 400+ MB/s for panoramic transformations, and 1+ GB/s for spectral-to-L*a*b* conversions.

Our multi-spectral HDR environments are large (sometimes > 1 GB) but are transformed into color-managed textures suitable for rendering on the GPU in seconds and from-scratch prefiltered maps are typically obtained in under a minute. For the special case of 1-D BRDFs (and an existing SH-projected environment), prefilter times go down to 2-3 seconds. Angular-space prefilter algorithms, in contrast, often go into hours (and days, for the kind of resolutions we often deal with).

The visual fidelity our system achieves is harder to quantify. We conducted many psycho-physical experiments comparing plastic chips (see Figure 10) and bottles, both rendered and viewed in a viewing booth with 3 illuminants next to each other together with color matchers from the plastics industry. We also compared automotive paint chips under the same illuminants with great success.

Figure 8. Brushed metal, wood, alien skin using 3-D BRDFs

Figure 9. Silver & Pearl using 1-D BRDFs

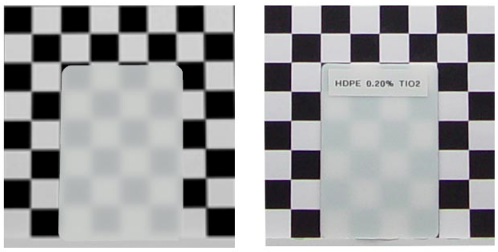

Figure 10. HDPE chip, rendering versus photo

9 Conclusions and Future Directions

We presented a visualization system which uses multi-spectral input sources and renders objects with complex surface characteristics under complex illumination, drawing on prefiltered view-dependent frequency-space environment maps.

We seek a fully GPU-based implementation, which also carries out the SH-based filtering portion (SH-projection, -rotation, and –expansion) of our render algorithm. Latest GPUs sporting 1000+ instructions for fragment programs support this goal. Filtering on the GPU will remove the need for prefiltering and hence eliminate altogether the one-time computational expense and (quite significant) storage requirements that come with SHRMs.

We also plan to couple SH-based BRDF rendering with a straightforward angular-space method for the highest-frequency portion of the BRDF to improve accuracy when rendering highly specular materials.

With BRDF rendering thusly enhanced we plan to synthesize 3-D BRDFs for our dynamic materials, superseding the current approximating shading model.

Finally, as soon as GPUs allow, we want to replace bump mapping with displacement mapping for greater visual fidelity when rendering bumpy surfaces.

Acknowledgements

We thank Ravi Ramamoorthi for helpful discussions regarding SHRMs and the Stanford University Computer Graphics Lab implementation of SHRMs.

We thank Thibault Lepoutre for help with our implementation.

References

1. J. Blinn, J., and M. Newell, Texture and Reflection in Computer Generated Images. Communications of the ACM 19, pg. 542–546 (1976).
2. B. Cabral, N. Max, and R. Springmeyer, Bidirectional Reflection Function for real-time rendering in dynamic, low-frequency lighting environments. ACM Transactions on Graphics 21, 3, pg. 527–536 (1987).
3. B. Cabral, M. Olano, and P. Nemec, Reflection Space Image Based Rendering. In Proceedings of ACM SIGGRAPH 99, 165–170From Complex Surfaces. In Proceedings of ACM SIGGRAPH 92, pg. 255–264 (1999).
4. N. Greene, Environment Mapping and Other Applications of World Projections, IEEE CG&A, 6(11) pg. 21-29 (1986).
5. W. Heidrich, and H. Seidel, Realistic, Hardware-Accelerated Shading and Lighting, SIGGRAPH '99, pg. 171-178.
6. J. Kautz and M. Mccool, Interactive Rendering with Arbitrary BRDFs using Separable Approximations, Eurographics Workshop on Rendering, pg. 281-292 (1999).
7. J. Kautz, P. Vazquez, W. Heidrich, and H. Seidel, A Unified Approach to Pre-filtered Environment Maps, Eurographics Workshop on Rendering, pg. 185-196 (2000).
8. J. Kautz, P.-P. Sloan, and J. Snyder, Arbitrary BRDF Shading for LowFrequency Lighting Using Spherical Harmonics. In 13th Eurographics Workshop on Rendering, pg. 301–308 (2002).
9. B.-T. Phong, Illumination for Computer Generated Pictures. Communications of the ACM 18, 6 (June), pg. 311–317 (1975).
10. R. Ramamoorthi, and P. Hanrahan, A Signal-Processing Framework for Inverse Rendering. In Proceedings of ACM SIGGRAPH 2001, 117–128 (2001).
11. R. Ramamoorthi, and P. Hanrahan, Frequency-space Environment Map Rendering. ACM Transactions on Graphics 21, 3, pg. 517–526 (2002).
12. G. Ward, The RADIANCE Lighting Simulation and Rendering System, Computer Graphics (Proceedings of '94 SIGGRAPH conference), pg. 459-472 (1994).
13. F. H. Imai and R. S. Berns, A comparative analysis of spectral reflectance reconstruction in various spaces using a trichromatic camera system, in IS&T/SID Seventh Color Imaging Conference, pg. 21-25 (1999).
14. F. H. Imai, M. R. Rosen and R. S. Berns, Comparison of spectrally narrow-band capture versus wide-band with a priori sample analysis for spectral reflectance estimation, in IS&T/SID Eight Color Imaging Conference, pg. 234-241 (2000).
15. H. Kotera, H. Motomura and T. Fumoto, Recovery of fundamental spectrum from color signals, in Recent Progress in Color Processing, The Society for Imaging Science and Technology, pg. 324 – 327 (1997).
16. Y. Miyak, K. Tonsh, T. Iked, N. Tsumura, T. Senzak and J. Hayashi, Development of Goniospectral Imaging Systems for Digital Archives, in IS&T/SID Tenth Color Imaging Conference, pg. 229-232 (2002).
17. F. Herbert, J. Kirkenaer and J. Ladson, Absolute and Relative Colorimetric Evaluation for Precise Color on Screen. Electronic Imaging, Color Imaging VIII: Processing, Hardcopy, and Applications (2003).
18. J. Leachtenauer, G. Garney, and A. Biache, Image Processing, Image Quality, Image Capture, Systems Conference, Portland, OR, pg. 124 – 129 (2000).
19. N. Katoh and T. Deguchi, Recent Progress in Color Management and Communications, Chapter II – Color Management: Characterization, pg. 49 – 55. (1998)
20. J. Leachtenauer, A. S. Biache and G. Garney: Image Processing, Image Quality, Image Capture, Systems Conference, Savannah, Georgia, 179 – 183 (1999).
21. R. Chung and S.Kuo, in IS&T/SID Fourth Color Imaging Conference, pg. 10 – 14 (1996)
22. N. Katoh, Recent Progress in Color Science, pg. 203 – 209 (1997).

Biography

Franz Herbert has been developing award-winning software like ColorBlind® and Remote Director™ for 20 years. He received his Diplomingenieur in Informatik from the Technical University of Vienna, and his Master's degree in Computer and Systems Engineering from Rensselaer Polytechnic Institute in Troy, NY. He is a member of IS&T and ACM.