Helsinki University of Technology
Telecommunications Software and Multimedia Laboratory
T-111.500 Seminar on computer graphics
Spring 2003

7.4.2003

High Dynamic Range Images

Timo Jyrinki
49728F

High Dynamic Range Images

Timo Jyrinki
HUT, Telecommunications Software and Multimedia Laboratory

Timo.Jyrinki@hut.fi

Abstract

This paper describes some of the principles in the computation and displaying of High Dynamic Range images. Various methods have been developed in computing High Dynamic Range images in higher precision and virtually unlimited range, and displaying the images on low-dynamic-range devices, such as CRT/LCD monitors. Tone mapping is a topic closely related to the latter sub-topic. Current state-of-the-art commodity hardware is able to do HDR-rendering.

1 INTRODUCTION

The need for HDR-rendering arises when comparing standard display devices, which have a dynamic range of luminance of about 100:1, with human vision which has a dynamic range of about 1000 000 000:1, from bright sunlight down to starlight (Devlin *et al.* , 2002). Also in computer generated images (CGI) created with global illumination algorithms a factor of 2000 between the highest and lowest intensity values is common and a factor of 30000 is sometimes reached (Schlick, 1994). Typically the images displayed on the screen are 24-bit - 8-bit per color component or a dynamic range of 256 intensity levels. The dynamic range in real-world environments thus far exceeds the range representable in 8-bit per-channel texture maps (Cohen *et al.* , 2001).

Because of the limitations in standard display devices, the computer generated HDR images still have to be displayed on the LDR display device somehow. Several techniques have been developed, including work in *Real-Time High Dynamic Range Texture Mapping* (Cohen *et al.* , 2001), *A Tone Mapping Algorithm for High Contrast Images* (Debevec & Gibson, 2002), *Photographic tone reproduction for digital images* (Reinhard *et al.* , 2002), *Gradient domain high dynamic range compression* (Fattal *et al.* , 2002), *A multiscale model of adaptation and spatial vision for realistic image display* (Pattanaik *et al.* , 1998), *Two methods for display of high contrast images* (Tumblin *et al.* , 1999) and *Adaptive gain control for high dynamic range image display* (Pattanaik & Yee, 2002). Some other contributors have been mentioned in a STAR report *Tone Reproduction and Physically Based Spectral Rendering* (Devlin *et al.* , 2002).

HDR rendering, i.e. increased accuracy throughout the graphics pipeline, is possible in real-time with the current state-of-the-art commodity hardware from ATI (Mitchell, 2002) and NVIDIA.

In the section 2 of this paper we will go through various ways of computing HDR images. It's followed by section 3, which describes a method which uses the current NVIDIA Geforce 2 -class LDR hardware to achieve HDR images, and further emphasizes the need for true HDR rendering. Section 4 discusses the concept of tone mapping and various methods of doing it, and finally section 5 explores the true real-time HDR rendering.

2 COMPUTING HDR IMAGES

Computing and acquiring images with high dynamic range can be done in various ways. One of the non real-time ways is to render images with a technique called Physically Based Spectral rendering (Devlin *et al.* , 2002). This technique was already noted early in the 1980ies but computer graphics mainstream has up to now avoided these techniques in favor of the seemingly more robust and less complicated colorspace approach.

Spectral rendering is a group of image synthesis methods which use a representation of the associated light spectrum for color values, as opposed to conventional systems that perform these calculations with tristimulus color values (red, green, blue). This can be beneficial as a part of HDR rendering to ultimately provoke the same responses that a viewer would have to a real scene. Problems that require Spectral Rendering for correct rendering behavior include problem of metamerism, dispersion in dielectric materials ('rainbow effect'), polarization of light and fluorescence. Metamerism means a pair of colors which differ spectrally but yield similar tristimulus values when viewed in some viewing conditions. In other viewing conditions they may seem different to the human visual system. Polarized light makes a difference especially in outdoor scenes, but the subject hasn't gotten a lot of attention yet. More obvious problems reside with current renderers, like scene complexity and global illumination issues.

Another way of creating a HDR image is to collect the data from a real-world environment. Utilizing a series of differently exposed photographs, enough data to represent a high dynamic range photo can be collected (Debevec & Malik, 1997). A lot of data can be acquired, as seen in Figure 1, where sixteen photographs were taken at 1-stop increments from 30 seconds to 1/1000 seconds. In the photos, sun is directly behind the rightmost stained glass window, making it especially bright.

Schlick's implementation, *High Dynamic Range Pixels* (Schlick, 1994), uses nonlinear quantization of pixel values in only 24 bits per pixel. The generation

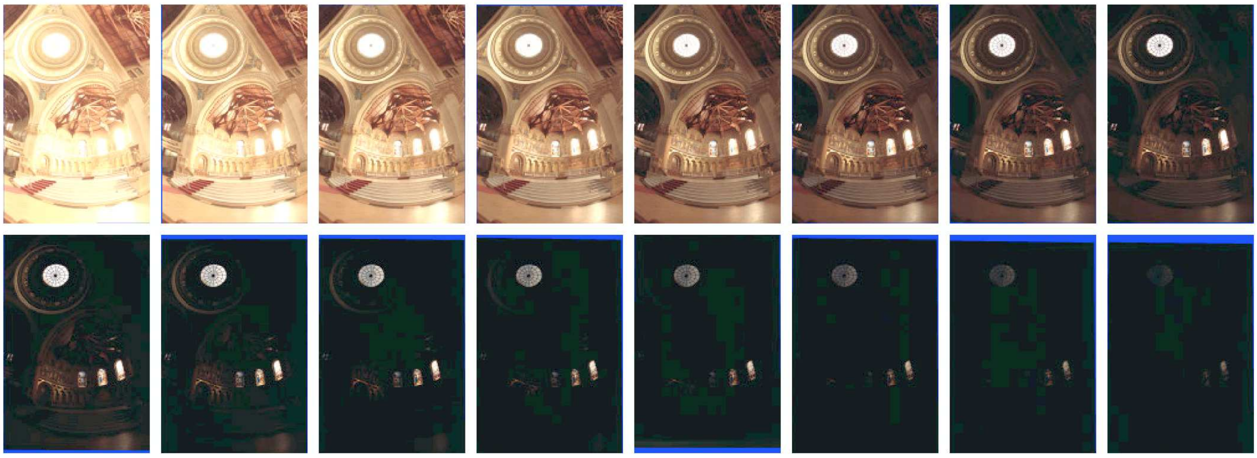

Figure 1: 16 photographs at different exposures (Debevec & Malik, 1997)

of images are done in software so that no specialized hardware is used. In other works a logarithmic function has sometimes been used, which maps the high dynamic range (e.g. 4000) to a lower range by enlarging the gaps for adjacent high values and keep smaller gaps in the lower range. This is done because the human visual system is much more sensitive to small differences between low values than between high values. Instead of logarithmic function, Schlick uses a slightly different bias function which has a parameter by which the strength of deformation applied to the incoming range can be varied. Furthermore, the function has desirable features for making decoding/encoding easier. This technique can be used for keeping HDR image data in a 24-bit image format.

3 DISPLAYING HDR IMAGES ON LDR DEVICES

Displaying HDR images on current display hardware is possible either by a pure software solution, by exposing the limited resources of lower class 3D-hardware using e.g. *Real-Time High-Dynamic Range Texture Mapping* (Cohen *et al.* , 2001), or finally, by using state-of-the-art commodity hardware with support for 64-bit, 96-bit or 128-bit floating point color spaces which have the required dynamic range throughout the rendering pipeline (see section 5). With recent hardware, the whole rendering process is then done in high dynamic range in hardware, and only reduced to an LDR output device at the very end of the rendering pipeline.

Real-Time High-Dynamic Range Texture Mapping (Cohen *et al.* , 2001) utilizes the possibility of allowing texels from multiple textures to be combined during texture fetching or in the framebuffer to produce the final rendered color. The method is illustrated in Figure 2. An exposure level may be defined, which then affects the resulting image. Thus 8-bit textures can be used to create hardware

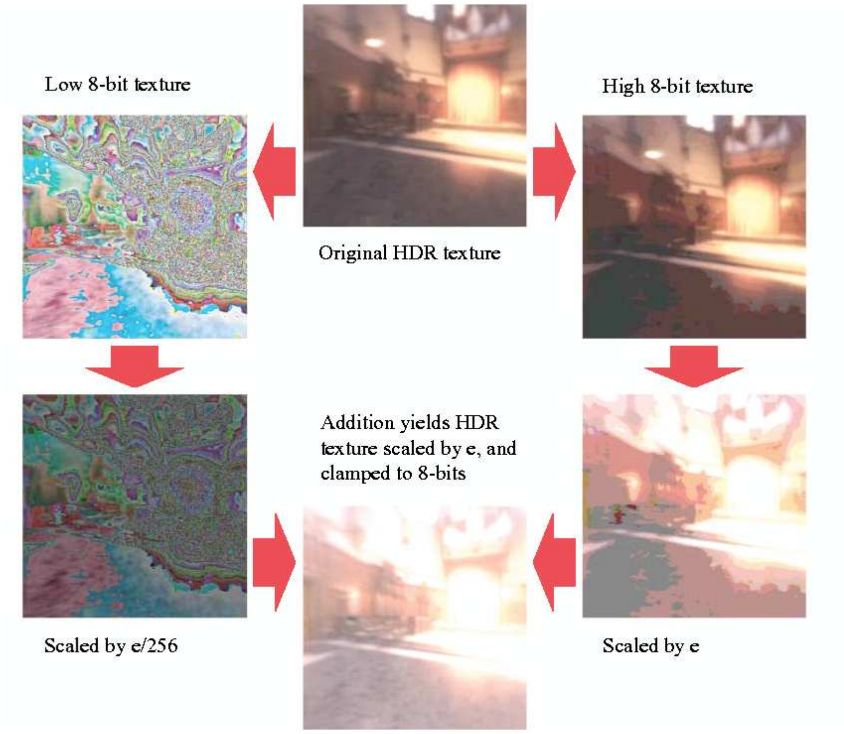

Figure 2: Real-Time HDR Texture Mapping (Cohen et al., 2001)

accelerated real-time HDR graphics.

True HDR-rendering throughout the pipeline is often needed, when multiple passes on framebuffer information is done. The precision in ordinary 24-bit framebuffers is not enough when the data is re-used and rounding errors occur, causing artifacts and loss of visual impressiveness. Examples include various volumetric effects, like ground fog, spherical fog and fading of distant objects - these effects can't be done smoothly when buffering and reusing per-pixel color information. Also, rounding errors can be noticed e.g. in bump mapping even with 64-bit floating point precision versus 128-bit, as seen in Figure 3 (64 bits means 16 bits per color channel). In floating point formats, some of the bits are reserved for sign and exponent, reducing the increase in dynamic range over traditional 24/32-bit integer formats. In the figure, a normal texture is used to simulate a displaced surface and there are shades of gray that are very close to each other. If two close values are subtracted with low precision formats, the result could be 0 rather than a small floating point value, altering the end result.

A 24-bit framebuffer is enough for the final display device, but the visual quality has to be manually fixed in order to produce a good-looking image to the framebuffer. By using HDR data and hardware in order to have all the data in high precision right until it's going to be shown on the LDR device, no visual quality will be lost if only the final stage of HDR->LDR conversion is done

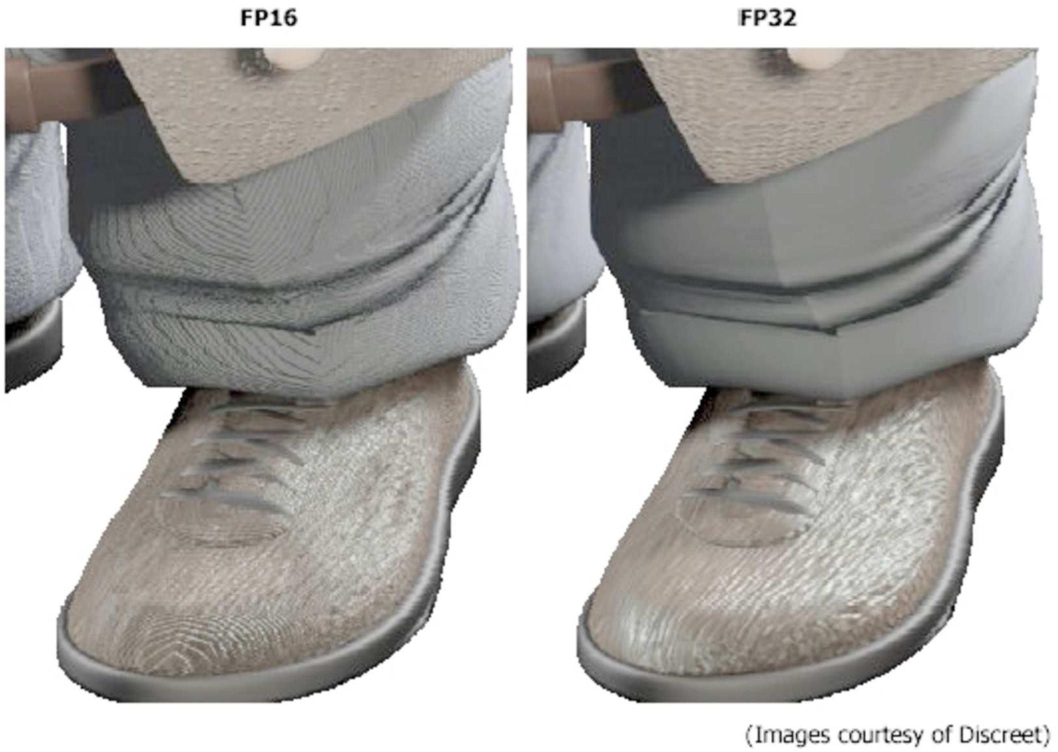

Figure 3: Bump Mapping Artifacts

properly. Whenever "true" HDR rendering is done, tone mapping is used to bring the HDR image to the display device.

4 TONE MAPPING

Tone mapping is what brings the computed HDR image to the display device. Since human visual system has such a wide dynamic range, some form of reproduction procedure to create perceptually accurate images on display devices is needed (Devlin *et al.* , 2002). The knowledge of Human Visual System (HVS) is still not thorough, but based on the facts we know new psychophysically-based visual models are being developed to address tone mapping problems. Most of the tone mapping techniques available today are concentrated on singular aspects for given purposes. The fact which makes tone mapping possible is that HVS has a greater sensitivity to relative rather than absolute luminance levels. Thus, dark/light adaption of HVS may be simulated using tone mapping to bring the desired response for the viewer.

According to *Tone Reproduction and Physically Based Spectral Rendering* (Devlin *et al.* , 2002), two types of tone reproduction operators can be used: spatially uniform (aka single-scale or global) and spatially varying (aka multi-scale or local). Spatially uniform operators apply the same transformation to every pixel regardless of their position in the image. Spatially varying operators apply a different scale to different parts of an image. Implementations on these two types of tone reproduction operators usually do not account for temporal differences such as adaptation of HVS over time, and are hence time independent methods.

An example of usage of spatially varying operator is presented in *Gradient Domain High Dynamic Range Compression* (Fattal *et al.* , 2002). They use luminance gradients, attenuating large gradients in order to bring areas closer together in the dynamic range. Different kind of tone mapping is used in different zones for compression. The results demonstrate that this method is capable of drastic dynamic range compression, preserving fine details and avoiding common artifacts (halos, gradient reversals, loss of local contrast). An example of the result of Fattal's method compared with a method of Ward Larson et al. and Tumblin and Turk is presented in Figure 4. The top-most picture is the result of Fattal's method, which preserves details in both high and low luminance areas. The two other pictures show either loss of detail in some areas or various artifacts.

Another recent example is the so-called "Zone System" (Reinhard *et al.* , 2002), which can be used for digital images or images captured using high dynamic range photography (Debevec & Malik, 1997). The method first applies a scaling that is analogous to setting exposure in a camera. After that, the method applies automatic dodging-and-burning to accomplish dynamic range compression. Dodging-and-burning is a concept from photography, where this printing technique withholds some light from a portion of the print during the development or adds more light. In computer graphics terms, they use local adaptation of perceptually-driven tone mapping operators. They tested various methods and finally selected the best one, which used Gaussian profiles that scale at different image positions accordingly. A comparison picture with a couple of other methods is shown in Figure 5.

One time dependent method is presented in *A model of visual adaptation for realistic image synthesis* (Ferwerda *et al.* , 1996). They developed a computational model simulating the adaptation in Human Visual System. In HVS, adaptation is achieved through the coordinated action of mechanical, photochemical, and neural processes. The model captures the changes in threshold visibility, color appearance, visual acuity and sensitivity over time, all of which are caused by the visual system's adaptation mechanisms.

As of this writing, no time dependent spatially varying operator had been seemingly proposed in scientific papers, though it might yield better results when properly implemented.

Figure 4: Tone mapping example (Fattal et al., 2002)

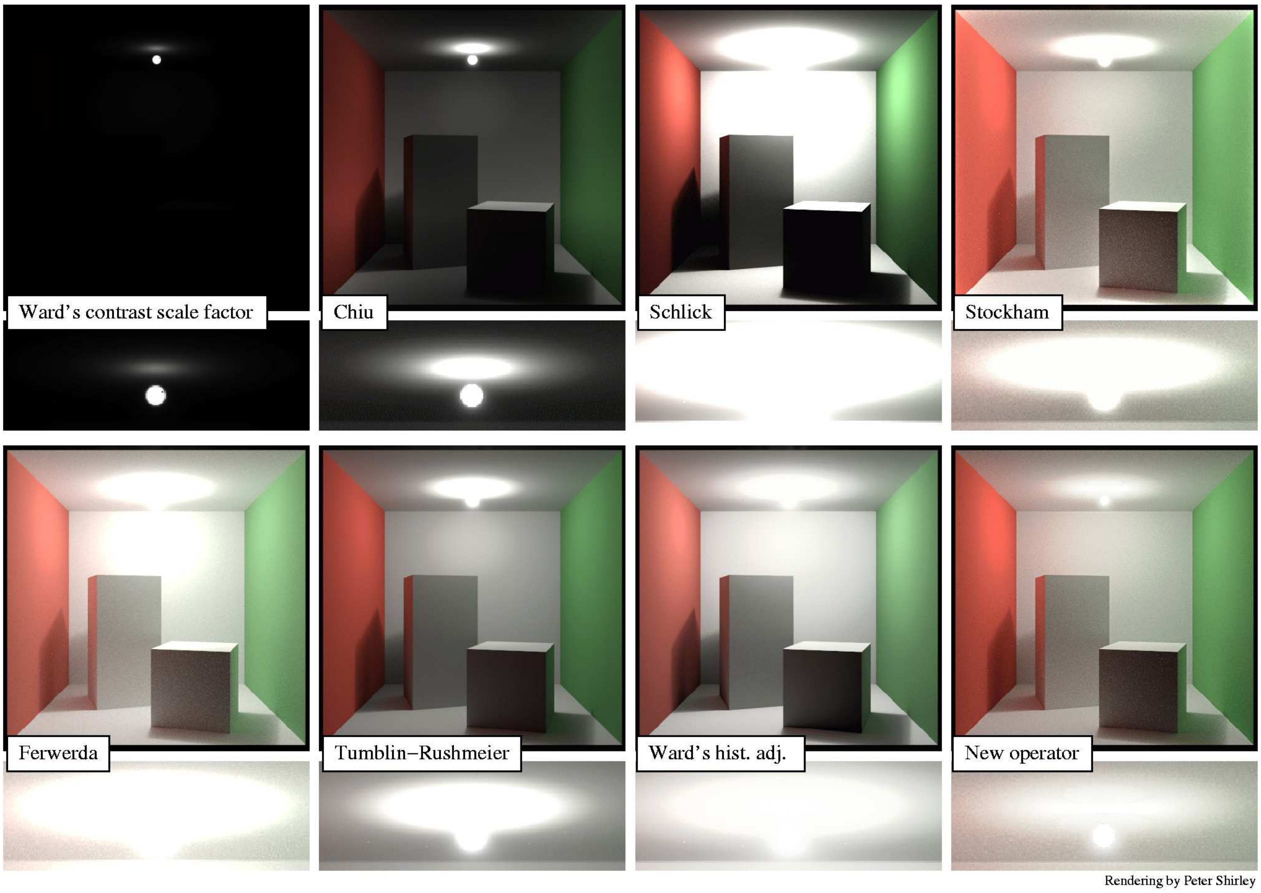

Figure 5: Tone mapping example (Reinhard et al., 2002) - Reinhard's new method is demonstrated in the bottom right corner

5 CURRENT STATE OF THE ART HARDWARE IMPLEMENTATIONS

The current commodity hardware can do HDR-rendering in *real-time*, as ATI has demonstrated by creating a real-time version (Mitchell, 2002) of a non real-time animation that was presented at SIGGRAPH98 (Debevec, 1998). In the ATI's version, a 64-bit image (16-16-16-16 for RGBA) is rendered by using a 16-16-16-16 planar reflection map and HDR diffuse, specular and rough specular illumination maps as appropriate for each material in the demo. Additional filters are used (see below) in the high dynamic range to enhance the image, and finally tone mapping is used to present the image on the monitor. The exposure level can be changed in real time, as the LDR image (8-8-8-8 bits that can be read by the Digital-to-Analog Converter (DAC)) is created from the HDR data each frame. This is illustrated in Figure 6.

One of the effects possible on the HDR-capable hardware is post-process glow (also known as the bloom effect). A glow applied to the bright parts of the image gives a pleasant, soft effect and emphasizes the perceived brightness of the image. The glow effect simulates light scattering in the eye or the optics viewing the virtual scene, in the case of high contrast differences. An example of glow usage is also seen in the ATI's HDR-image (Figure 6), where it's used in the bright spots of the globe(s). In the image, a glow filter is applied to the image to achieve the "blooming". The blooms are added back to the HDR rendering, and this sum is vignetted (softened) and tone mapped down to the displayable range.

To achieve full HDR-rendering on the hardware, floating point precision color has to be used. Currently, if real-time HDR-rendering is done, it's mostly being done in 64-bit integer or floating point colors, as in the example above. In the future 128-bit FP colors will probably be used, as they are already supported in all the newest consumer graphics chipsets from the major vendors NVIDIA and ATI (Geforce FX 5200/5600/5800 and Radeon 9500/9600/9700/9800) and in the standards like OpenGL. Each color channel has 32-bit FP precision, which consists on 1 sign bit, 8 exponent bits and 23 mantissa bits. This allows a dynamic range from 2^{-128} to 2^{127}, or 23 bits (mantissa) versus the old 8 bits per channel. In real-time graphics, this will allow images of greater quality than today, but with an expense of the requirement of more processing speed and especially memory bandwidth.

6 CONCLUSIONS

HDR-rendering as a whole is a wide subject and there are many approaches to the problem of generating and visualizing HDR images. Some may just approach the problem by trying what looks good, others try the old methods known from the 150 years of history of photography, and yet others try to go deeper in the

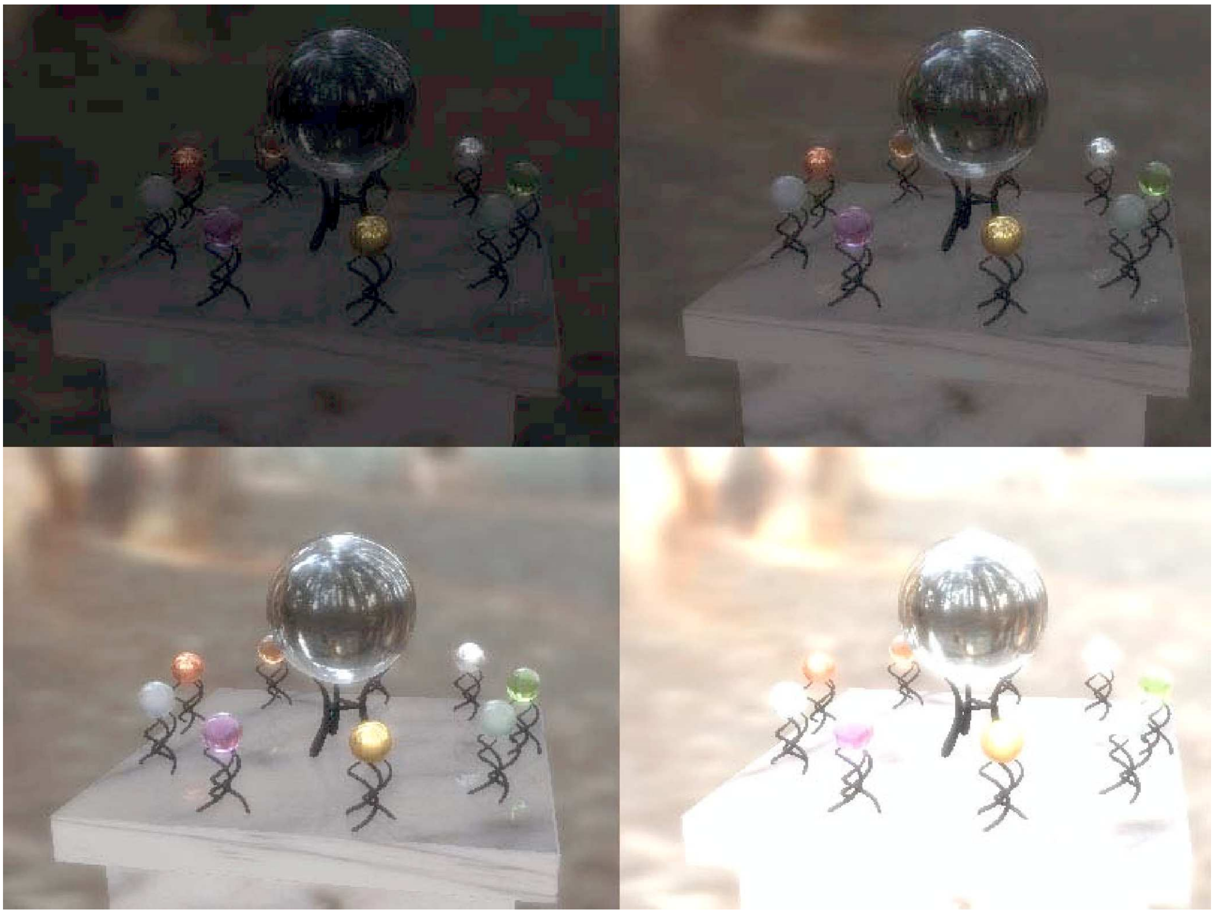

Figure 6: ATI's HDR example (Mitchell, 2002)

workings of the human visual system (HVS) in order to find out how the "biological tone reproduction" works. Regardless of the approach, HDR images are bringing more realistic computer images to the display devices in the future. For example, HDR rendering allows the graphic designer to truly set some part of the image/texture "very bright", without the need of manually fixing all of the environment so that the resulting bright portion would truly look bright. With the freshly released new commodity graphics hardware, and the upcoming models in the following couple of years, new levels of real-time graphics may be achieved. John Carmack, the creator of the Doom and Quake series of graphics engines, has said at Quakecon 2002 keynote speech that the next evolution in graphics development will be tone mapping, amongst with higher precision of color.

REFERENCES

Cohen Jonathan, Tchou Chris, Hawkins Tim, & Debevec Paul. 2001. Real-Time High Dynamic Range Texture Mapping. *Pages 313–320 of: To appear in Eurographics Rendering Workshop, London, 2001.*

Debevec Paul. 1998. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Pages 189–198 of: Proceedings of the 25th annual conference on Computer graphics and interactive techniques.* ACM Press.

Debevec Paul, & Gibson Simon. 2002. *A Tone Mapping Algorithm for High Contrast Images.*

Debevec Paul E., & Malik Jitendra. 1997. Recovering high dynamic range radiance maps from photographs. *Pages 369–378 of: Proceedings of the 24th annual conference on Computer graphics and interactive techniques.* ACM Press/Addison-Wesley Publishing Co.

Devlin Kate, Chalmers Alan, Wilkie Alexander, & Purgathofer Werner. 2002. *Tone Reproduction and Physically Based Spectral Rendering.*

Fattal Raanan, Lischinski Dani, & Werman Michael. 2002. Gradient domain high dynamic range compression. *Pages 249–256 of: Proceedings of the 29th annual conference on Computer graphics and interactive techniques.* ACM Press.

Ferwerda James A., Pattanaik Sumanta N., Shirley Peter, & Greenberg Donald P. 1996. A model of visual adaptation for realistic image synthesis. *Pages 249–258 of: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques.* ACM Press.

Mitchell Jason L. 2002. *Radeon 9700 Shading.* Tech. rept. ATI Research.

Pattanaik Sumanta, & Yee Hector. 2002. Adaptive gain control for high dynamic range image display. *Pages 83–87 of: Proceedings of the 18th spring conference on Computer graphics*. ACM Press.

Pattanaik Sumanta N., Ferwerda James A., Fairchild Mark D., & Greenberg Donald P. 1998. A multiscale model of adaptation and spatial vision for realistic image display. *Pages 287–298 of: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. ACM Press.

Reinhard Erik, Stark Michael, Shirley Peter, & Ferwerda James. 2002. Photographic tone reproduction for digital images. *Pages 267–276 of: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. ACM Press.

Schlick Christophe. 1994. High Dynamic Range Pixels. *Pages 422–430 of:* Heckbert Paul S. (ed), *Graphics Gems IV*. Boston, MA: Academic Press Professional.

Tumblin Jack, Hodgins Jessica K., & Guenter Brian K. 1999. Two methods for display of high contrast images. *ACM Transactions on Graphics (TOG)*, **18**(1), 56–94.