

Real Illumination from Virtual Environments

Abhijeet Ghosh, Matthew Trentacoste, Helge Seetzen, Wolfgang Heidrich

The University of British Columbia[†]

Abstract

We introduce a method for actively controlling the illumination in a room so that it is consistent with a virtual world. In combination with a high dynamic range display, the system produces both uniform and directional illumination at intensity levels covering a wide range of real-world environments. It thereby allows natural adaptation processes of the human visual system to take place, for example when moving between bright and dark environments. In addition, the directional illumination provides additional information about the environment in the user's peripheral field of view.

We describe both the hardware and the software aspects of our system. We also conducted an informal survey to determine whether users prefer the dynamic illumination over constant room illumination in an entertainment setting.

Categories and Subject Descriptors (according to ACM CCS): B.4.2 [INPUT/OUTPUT AND DATA COMMUNICATIONS]: Input/Output Devices, Image display. I.3.7 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism, Virtual reality.

1. Introduction

Attempts to create a sense of presence and immersion in a virtual environment have been a major theme throughout the history of computer graphics. Research targeting this problem has so far had to deal with two major restrictions. First, conventional display technology has been incapable of representing the full range of dark and light intensities found in the real world. Tone mapping operators alleviate this problem to some degree but are unable to compensate fully for these shortcomings [LCS04]. For the most part, the dynamic range problem has been addressed recently by new high dynamic range (HDR) display technology [SHS*04]. Second, the viewing conditions are largely unknown, meaning that parameters such as the viewer's light and color adaptation cannot be considered in the image generation process.

The latter problem becomes particularly obvious with HDR displays, which can produce a range of intensities from moonlight to daylight, although they cannot reproduce the brightness of direct sunlight. The display behaves like a window into a virtual world, but a sense of immersion can only

be achieved if the illumination levels in the real and the virtual worlds are compatible. A night driving simulation, for example, should happen in a darkened room, while the same application in a daylight setting should take place in a bright room. Ideally, it would be possible to adjust the room illumination over time to simulate, for example, the car entering or leaving a tunnel.

In this paper, we propose to actively control lighting in the room according to the illumination in the virtual environment (see Figure 11). In our prototype implementation, we focus on methods that could conceivably be used in entertainment applications, such as gaming environments or home theaters. We use computer-controlled LED lights that are distributed throughout the room. All lights are individually programmable to a 24 bit RGB color. This setup allows us not only to raise or lower the ambient light in the room, but also to create some degree of directional illumination, which results in a low-resolution dynamic room illumination approximating an environment-map for the assumed viewing position. Although the light sources are located outside of the user's direct field of view, the directional illumination interacts with objects inside the field of view, such as

[†] E-mail: {ghosh, mmt, seetzen, heidrich}@cs.ubc.ca

the monitor or the wall behind it. This is the main concept behind our approach.

To evaluate the potential for the proposed method in an entertainment setting, we conducted survey of user preference. All participants of the survey preferred the system with dynamic, directional illumination over a room of constant brightness. The participants also believed that the additional cues provided by directional illumination helped them keep track of orientation in the virtual world.

In the following sections, we first review related work (Section 2) before describing the system, including hardware, calibration, and rendering algorithms (Section 3). We then describe several ways to acquire the relevant lighting information for both virtual worlds and film sequences (Section 4). Finally, in Section 5, we discuss the user survey that we performed to test the concept.

2. Related Work

A significant body of research in realistic rendering has focused on tone mapping operators for displaying a wide range of intensities on display technology with limited capabilities. Much of this work is intended primarily for still images (e.g., [Sch94, LRP97, TT99, DD02, RSSF02] and others).

Ferwerda et al. [FPSG96] pioneered work on tone mapping operators that *explicitly* take visual adaptation into account. Their work uses threshold vs. intensity functions to map threshold contrasts from the original intensity range to that of the display. Pattanaik et al. [PFFG98] developed a model that also incorporates supra-threshold brightness, color, and visual acuity. Both models could, in principle, be used to track changes in visual adaptation over time, although they are computationally rather expensive. Other researchers [SSS00, DD00, PTYG00, AWW03] have developed methods to model the time-dependent state of adaptation based on the recent history of viewing conditions in the virtual world.

One limitation of these methods is that they have no information of the user's *actual* state of adaptation, since it depends largely on the real-world illumination in the room, rather than the virtual world. An image generated for a dark-adapted viewer will not be perceived as realistic in a bright room. The use of HDR displays [SHS*04] largely removes the need for tone mapping operators, since they can reproduce intensities from mesopic to medium photopic vision levels. However, this does not solve the problem of unknown viewing conditions.

A vast body of literature in the perceptual psychology community deals with the impact of room lighting conditions. Many of these studies were performed to analyze the impact of room lighting on ergonomic factors such as screen visibility, eye strain, and so forth. There has also been work on using light sensors to adjust the display brightness and

contrast (e.g. [?, ?]). These studies try to either minimize the influence of room lighting on displays or compensate at the display end for the illumination conditions. We, on the other hand, deliberately focus on tying the room lighting into the viewing environment.

Other studies analyze the perceived brightness vs. luminance levels of images viewed under different room illumination (e.g. [BB67, DeM72]). Most existing studies only cover static illumination. An exception is [?], which discusses the need for adjusting for dynamic lighting changes in critical applications such as reading controls for cockpits. However, all studies we are aware of implicitly assume that room lighting is the primary factor for adaptation, and that it has a significant impact on the display surface itself. Both of these assumptions do not hold for HDR displays, and as a consequence the findings from these studies do not apply to our setting.

One way of integrating real and virtual objects is to change the illumination in the virtual world to make it consistent with the real world. Nayar et al. [NBB04] recently developed a lighting sensitive display that tracks changes of illumination in the room, including directional changes, and uses this information to shade virtual objects. Our work takes the opposite approach and adapts the room illumination to be consistent with a virtual world. We believe there is room for both approaches: while Nayar et al.'s method creates opportunities for new user interaction metaphors, ours is useful whenever a specific virtual world needs to be displayed.

There has been other work that uses separate light sources to augment computer displays. Philips produces a series of high-end flatpanel TVs which have light sources on the back to illuminate the wall behind the TV [Phi]. The lights are driven uniformly based on the average intensity of the screen content, thereby essentially reducing contrast between the wall and the TV screen. In contrast, our system creates directional illumination based on actual information from the virtual environment. Light sources have also been used to augment displays in theme park rides. There, lights are often used together with other physical props to show a fixed scene. We, on the other hand can deal with dynamically generated content, but aim only at creating low resolution information for the peripheral field of view.

Other related work includes Debevec et al.'s Light Stage 3 [DWT*02]. That work uses a number of computer-controlled lights to illuminate actors or objects such that they appear on camera as if they were in a certain real-world environment. An actor in Light Stage 3 sees only a number of point lights, while we aim at producing a smooth environmental illumination that can convincingly represent the real environment in the user's peripheral view. This goal requires a different physical setup, as well as different calibration and rendering algorithms.

Also related to our work are fully immersive, CAVE-like environments [CNSD93]. Unfortunately, CAVEs, like other

VR displays, have very limited contrast, which makes them unsuitable for representing the kind of adaptation processes we are interested in. Due to engineering constraints such as power consumption and heat production [SHS*04], it seems unlikely that immersive HDR environments will be feasible in the near or even medium term future. Space and cost present further obstacles. Such a system would also require high resolution, omnidirectional illumination information, which is hard to generate, for example, in live action film. We therefore focus on conventional, limited field-of-view displays with high contrast. We augment these with low-resolution directional illumination generated by a small number of light sources. These lights can easily be positioned in an ordinary room and do not preclude other uses of the space.

3. Method

The goal of our system is to illuminate the room so that it matches a low-frequency version of the virtual scene. Computer-controlled lights are programmed such that, for a specific real-world viewing position, the room illumination resembles a blurred environment map for the virtual world at the virtual viewing position.

Our method has three major components: the physical setup of the light sources in a room, calibration methods, and rendering algorithms. These aspects are discussed in the following sections.

3.1. Materials and Setup

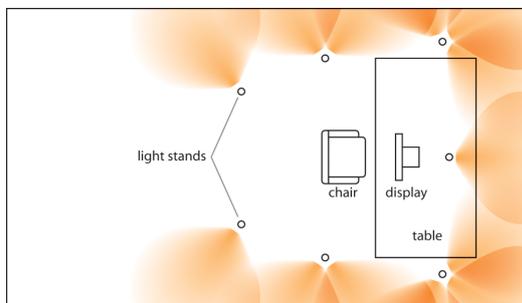


Figure 1: Room layout: additional lights are mounted below the ceiling, pointing upwards.

We assembled our prototype system in a separate room, approximately 15.5' long, 9' wide, and 9.5' high. The room remained as-is: the walls were kept in the original pastel color, and specular objects such as a whiteboard and the reflectors of the houselights were left in place. A window cover was used to block out daylight (Figure 11, left). The room contained several pieces of furniture, including two tables and several chairs. One of the tables was located at one end of the room and held the computer console. We used both

an 18" LED HDR display [SHS*04] and a standard 18" flat-panel (NEC Multisync LCD 1850e) for our experiments.

The lighting system consists of 24 RGB LED lights (ColorKinetics iColor Cove), each of which can be individually programmed to a 24 bit RGB color value. Instead of pointing the light sources directly at the viewer, which would create high intensity illumination from very specific directions, we aimed the lights at the walls in order to diffuse the light output over a large range of directions. This corresponds to our goal of creating a lighting system that can produce a low-frequency version of the illumination in the virtual world.

We used seven poles with stands to mount the 24 light sources. The lights were positioned and oriented such that they predominantly illuminated the ceiling, as well as the walls to the left, right and in front of the viewer (see Figure 1). Experiments quickly confirmed our intuition that illumination from behind has only a comparatively small impact, and hence we used only a few light sources for those directions. No illuminators were aimed towards the floor for similar reasons. Our arrangement roughly mimics the change in photoreceptor density in the retina from foveal to peripheral view. However, it might be interesting future work to design the physical setup by formally taking into account the resolution of the human eye [?].

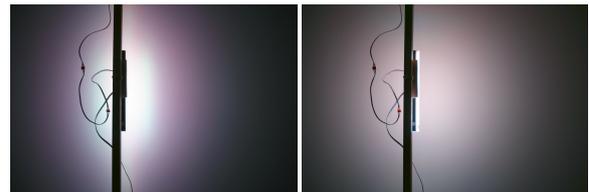


Figure 2: Left: light pattern generated by a single iColor Cove light. Note the narrow light spot and color banding. Right: pattern generated using a diffuser.

To create a smoothly varying illumination pattern we used strong diffusers at the light sources, which also reduce color separation of the RGB elements (Figure 2). The diffuser for each light consist of 2" diameter transparent acrylic tubing that was cut in half along its axis, and spray-painted lightly on the outside with white plastic paint (Krylon Fusion for plastic). To avoid internal reflection losses we used reflective film to coat the inside of the light source.

The typical light output of each light is specified as 52.4 lumens for full white. This is not bright enough to match the top intensity of the HDR display, but since every light is set to an average intensity over a moderately large cone of directions, this limited top intensity has not been an issue[†]. It should also be noted that brighter LED than the ones

[†] Typical HDR scenes have only a few small bright regions (corresponding to windows or skylights) that are to full white. Experi-

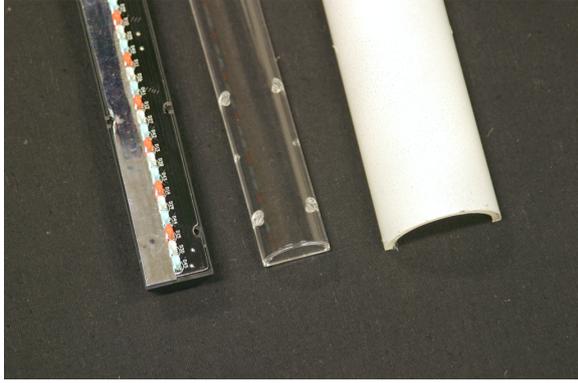


Figure 3: Left: opened iColor Cove light source. The left side of the circuit board is already covered with reflective foil. Center: clear cover for the light source. Right: diffuser built from acrylic tubing and white spray-paint.

we use are readily available. We chose the iColor Cove system primarily because it includes off-the-shelf electronics for computer control.

3.2. Calibration

Some calibration steps are necessary in order to control the system in a way consistent with the virtual world and the image shown on the display. Geometric calibration is necessary to determine the positions of the light sources relative to the viewer, and their spread, which is modeled as a Gaussian. Photometric calibration is subsequently performed to match white points and illumination levels between the light sources and the display.

Light Position. The rendering algorithm, as described below, requires information about the impact of every individual light on the illumination as seen from the location of the viewer. To obtain this information, we place a reflective ball at the intended viewer position to act as a light probe. We take photographs with a web camera (Creative NX Ultra) while switching on one light at a time. The resulting environment maps appear in Figure 4.

We then model the impact of every light source by fitting a Gaussian to the environment map. It is centered around the direction corresponding to the brightest point in the environment map, and its standard deviation is chosen such as to minimize the RMS error. Other directional bases such as cosine lobes could be used instead of Gaussians. However, we found Gaussians to be convenient, since they can capture more distant contributions caused by indirect illumination,

ments we conducted with the HDR display indicate that overall light output is typically less than 10% of peak intensity.

while smoothing over high frequency details such as object boundaries.

White Point and Intensity Calibration. An important part of the calibration step is to match the white points of the display and the lighting system. At the same time, we need to establish the relative intensities of light sources and display. For both tasks, we use an 18% gray card commonly used in photography. Under the assumption of uniform hemispherical illumination ($L_i(\omega_i) = \text{const}, \omega_i \in \Omega$), the reflected radiance of a (diffuse!) 18% gray card is

$$L_o(\omega_o) = \int_{\Omega} L_i \cdot \frac{0.18}{\pi} \cos \theta_i d\omega_i = 0.18 \cdot L_i.$$

Since uniform hemispherical illumination can be approximated by setting all our lights to the same intensity, the calibration task is implemented as a uniform adaptation of the intensity of the lights until the gray card reflection matches an 18% monitor gray. Note that the response function of the display needs to be taken into account, i.e., the monitor color is set to 18% of the top *intensity*, not 18% of the top pixel value. The color matching can be automated by using the same web cam as above. We cover one half of the screen with the gray card, and show 18% red on the other half of the display. We then adjust the red intensity of the light sources using binary search until the camera observes the same intensities for both the monitor image and the gray card. The same steps are repeated for the green and blue channels.

From this procedure we recover the relative scaling factors for the light sources that correspond to the full intensity of the individual color channels on the display. During rendering, these can either be used directly to adjust the light intensities, or their reciprocals can be applied to the image shown on the display. If an absolute white point calibration is desired, the former method should be used, and the monitor should be calibrated with standard tools.

Since the lighting setup is indirect, the color of the walls and other large objects does influence the color temperature of the illumination. If the walls or other large objects in the room show a great variation in color temperature, then the effective contributions of the individual lights have a different color. In that case, the color difference for the individual lights needs to be calibrated first by sequentially switching them on, and comparing the color of the resulting illumination on the gray card. Only then can the intensity be calibrated as described above. Note that the white point and intensity calibration should ideally be repeated every time a light source moves, or even when large, colored objects get moved around in the room. Fortunately, all calibration steps are automated and can be completed within a few minutes.

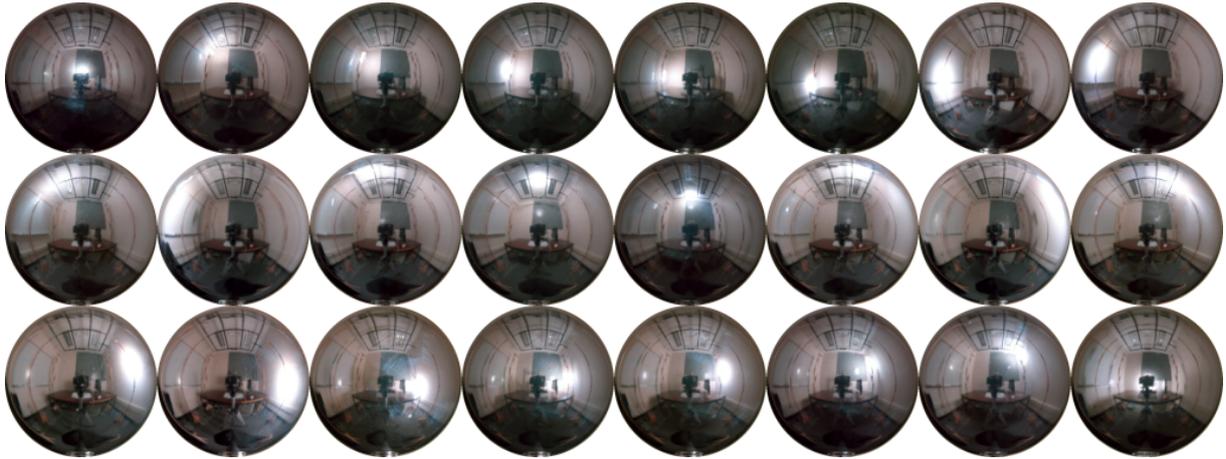


Figure 4: Light probe images acquired for each of the 24 light sources at the intended viewing position.

3.3. Rendering

There are two algorithms for driving the calibrated lighting system during rendering. The first option is to uniformly adjust the light source intensity to the average intensity of the scene. The second option is to drive every light source individually by sampling the scene’s illumination in the region affected by the light. In both cases, we use an environment map for the location of the viewer as a representation of the scene illumination. This is a convenient choice, since many realtime graphics applications already create those maps for shading objects near the viewer. It also does not require any scene geometry, and could therefore easily be painted by an artist to augment old footage (see Section 5). Ideally, the environment map should be in high dynamic range format, but low dynamic range information can also be used, especially if the environment map is split into different layers. This is discussed further in Section 4.

In the case of uniform illumination, we simply integrate the intensities in the environment map, and use the resulting value to drive all light sources.

In the case of directional illumination, we precompute an importance sampling pattern for the Gaussians that we fit to every light in the calibration process (Section 3.2). For every frame, we then sample the illumination from these patterns, and use the resulting integral to drive the light sources. We can also blend uniform and directional illumination to control the degree of directional dependence.

4. Content Creation

The environment maps used to control the light sources can be generated in a number of ways, depending on the application.

Synthetic Environments. The system is easy to integrate into fully synthetic scenes, such as in computer games or animated films. In both cases, the environment map for the viewer location is readily computable. Many games already generate these environment maps for shading objects in the scene [HS99]. Future games will likely generate these environment maps in an HDR format. Current games typically use low-dynamic-range representations due to the lack of support for floating point textures and framebuffers in older graphics hardware. However, the environments are often split into multiple layers, corresponding to different parts of the scene. This layered information can be used to reconstruct HDR lighting information.



Figure 5: Left: room lighting corresponding to a passing street lamp on the left side in the driving game NFS underground 2. Right: lighting corresponding to overhead lights and lights on the walls inside a tunnel in the game.

For our experiments we used footage from Electronic Arts’ racing game “Need for Speed Underground 2” (Figure 5), which features a particularly wide range of differently illuminated environments. We used captured environment maps generated by the existing shading system. The layers of the environment map, which correspond to lights, sky, and objects in the scene, were scaled by different factors and added up. The resulting HDR information was used to program the lighting system.

Legacy Film and Video Footage. For some applications it is interesting to retro-fit conventionally shot film and video material with environment lighting information. Sometimes it might be possible to automatically extract the required information from the image sequences themselves. For example, Nishino and Nayar’s approach for extracting environment maps from reflections in eyes [NN04] might be used for sequences in which human or animal faces are visible at a high enough resolution.

In other cases, the required information can be generated manually, since only approximate low-resolution lighting information is necessary. A uniform brightness can be estimated for a set of key frames and interpolated across the sequence. The generation of directional information is more tedious, but as shown by Sloan et al. [SMGG01], artists can control lighting in a scene by painting spherical environment maps. This could be done for a set of key frames, and the resulting environments could then be interpolated using a method similar to the one proposed by Cabral et al. [CON99]. We leave this for future work.



Figure 6: *Left: uniform lighting corresponding to direct daylight outside a tunnel in the HDR driving video. Right: uniform room lighting corresponding to lighting inside the tunnel.*

For our experiments we manually created environment map information for HDR video footage of a car ride, and played it back on an 18” LED HDR display [SHS*04]. The video was shot with an HDR camera (HDRC CamCube from IMS Vision), and compressed using the method described by Mantiuk et al. [MKMS04]. It features both direct daylight and a dark tunnel (Figure 6). We augmented this video with environment lighting information by uniformly adjusting the brightness level for every frame. We mostly used two brightness values, one for the inside of the tunnel and one for the daylight part. At the entrance and exit of the tunnel, we interpolated the uniform intensity linearly over a few seconds.

Filming New Footage. When shooting new films, a light probe can be used to capture the surrounding environment. Unlike in relighting applications [Deb98, DWT*02], the environment maps for our system (Figure 7) should ideally be centered at the viewer position, that is, the main camera used for filming the scene. This should make it feasible to record a light-probe video with all additional components outside the field-of-view of that of the main camera. In some



Figure 7: *Left: directional lighting corresponding to the bright windows in the Grace Cathedral HDR environment. Right: directional lighting corresponding to the alter in the Grace Cathedral.*

cases it might be necessary or desirable to post-process the sequences to account for specific lighting effects.

5. Experiments

To test the concept, we conducted a user survey with a set of three experiments. As the evaluation criterion, we chose user preference rather than other possible criteria such as perceived realism. This choice was made due to our primary interest in entertainment applications for the current system; ultimately the proposed method can be successful for these applications if and only if potential users like the results independent of how realistic they believe those are. This means that other, more formal studies with a different aim are required before the system could be used, for example, for design purposes.

The first experiment focused on the impact of uniform changes in room illumination, while the second one emphasized the directional aspects of our approach. The third experiment was designed to test whether the lighting system could also be useful in combination with conventional low-dynamic-range displays. The participants were 12 graduate and undergraduate students, none of whom work in computer graphics or related areas. All participants had normal or corrected-to-normal vision. The participants entered the room at least 5 minutes before the start of the actual experiments in order to allow them to adapt to a slightly dimmed environment. Questions were asked after individual experiments. After all experiments had been completed, the participants had the opportunity to provide additional comments.

Uniform Illumination. We designed the first experiment to test whether the participants would prefer a dynamic, but directionally uniform illumination level over constant room illumination. To this end, the participants were shown the HDR driving video (see Section 4) on the HDR display once in a dark room, and once with a uniform brightness change generated by the lighting system. The participants were then asked to indicate their preference for either the constant or the dynamic illumination on a 5-level scale (strong preference for the dark room, weak preference for the dark room,

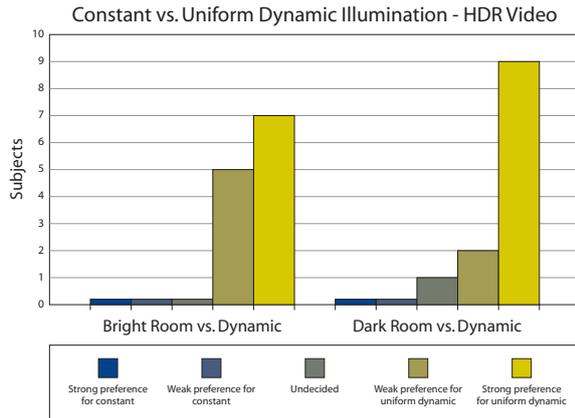


Figure 8: User preferences regarding constant or uniform dynamic illumination for HDR video.

undecided, weak preference for dynamic lighting, and strong preference for dynamic lighting). The same experiment was repeated for a room illuminated at normal brightness levels. Again, both variants were shown back-to-back, and the participants were asked to state their preference.

The answers given by the participants are summarized in Figure 8. All participants preferred or strongly preferred dynamic illumination over a constant brightness level. In general, the preference was even stronger in the comparison with a dark room, although one participant was undecided. This stronger preference can be explained by the ability of the HDR display to produce light levels that are starting to be uncomfortable in very dark environments.

From this results we can conclude a significant preference of the participants for dynamic illumination compared to a constant light level.

Directional Illumination. The second set of experiments tested whether the users would prefer directionally localized illumination changes over uniform brightness changes. This experiment was based on using a simple viewer similar to Quicktime VR [Che95]: the program loads an HDR panorama, shows it on the HDR display, and lets the user look around with a simple mouse interface for rotation (see Figure 7). Note that a ‘dynamic’ illumination approach with uniform intensities for all lights would result in a constant illumination pattern for this application, since total brightness would not change under rotations of the viewing direction.

The participant was then asked to use the rotation interface to locate the brightest point in the panorama. The application was run 20 times, and each time a different panorama was selected at random. The application also randomly decided whether or not to use the lighting system to create directional information (if not, the lights were switched to a medium intensity). After 20 runs, the participants were

asked two questions: first, whether they felt that the directional illumination helped them with orientation (on a scale from strongly disagree to strongly agree), and second whether they preferred the directional or the uniform illumination.

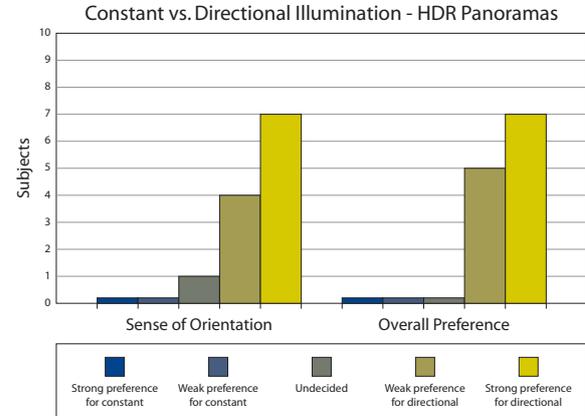


Figure 9: User preferences for directional vs. uniform illumination in an HDR panorama viewer.

Figure 9 shows the answers for both questions. All participants preferred or strongly preferred the directional illumination over the uniform one. With one exception, all participants answered the question regarding an improved sense of orientation identically to the way they answered the question for overall preference. One participant was uncertain whether the directional lighting had improved his sense of orientation, but nonetheless preferred directional over uniform lighting.

From the results of this experiment, it is clear that dynamic directional illumination is preferable over both constant and uniform adaptive lighting (note again, that the latter would have produced constant illumination as well in this application).

Low-Dynamic-Range Footage. Finally, we also wanted to determine whether the lighting system is useful in combination with conventional low-dynamic-range displays. To this end, the participants were shown the footage from “Need for Speed Underground 2” (see Section 4) on a conventional display. The segment contained a tunnel sequence, in which widely spaced street lights caused the scene to get brighter and darker at regular intervals. First, we showed the segment under constant illumination, and then with the dynamic, directional illumination, and the participants were asked for their preference.

In this experiment, the preference for directional dynamic illumination was very strong, as indicated in Figure 9. One participant was undecided, as also revealed in his written comment: “[I] did not like the flickering lights [when pass-

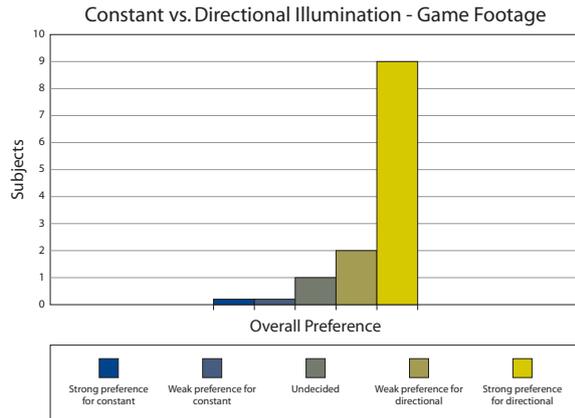


Figure 10: User preferences for directional vs. uniform illumination when watching low-dynamic-range video game footage.

ing by the street lights in Need for Speed] – very realistic, but very annoying (it’s distracting enough when you’re driving in real life). [It is] more annoying in real life. [I] really enjoyed the feeling of motion”. This subject did not express similar concerns in the directional experiment using the HDR display, where he strongly preferred the directional illumination. We believe that this ambivalence is in some sense caused by the lighting system overpowering the conventional display, which cannot produce the same intensities as the HDR display.

Based on the overwhelmingly positive response of the other participants, we do believe that the system has potential even in a low-dynamic-range setting. However, more studies are required to determine how the lights should be controlled in this scenario to avoid irritations for some users.

Other comments from participants not attributed to a specific experiment included the following: “The dynamic lighting immerses you in the experience”, “[Dynamic lighting] makes you really feel like you’re there, especially in light-dark transitions”, and “One day, this will be in the movies!”

6. Discussion and Conclusions

We have introduced an approach for actively controlling the lighting in a room to match illumination in a virtual world. In doing so, we are able to reproduce illumination levels similar to the ones we experience every day in the real world. This triggers natural adaptation processes in the human visual system, for example, when moving between bright and dark environments. In addition, we can generate directional illumination patterns, such as light that appears to come from the sides or from behind.

Our user survey shows overwhelming support for this

concept in combination with an HDR display: *all* of our participants preferred the lighting system over constant room illumination. We believe that this combination of HDR display and lighting system comprises the best setup, since it makes it possible to create similar brightnesses both on the display and in the surrounding room.

Even in combination with a conventional low-dynamic-range display, the participants were predominantly positive about the lighting system. With the display dimmer than the light sources, one subject was irritated by the dynamic illumination, although he strongly preferred the system in combination with an HDR display. This indicates that, while the lighting system is promising even in combination with conventional displays, the algorithms for driving the system in such a setting require more research. One solution could be the use of a tone mapping operator for the light sources.

We believe that the work presented here also creates a variety of promising directions for future research. One important area is artistic tools for content creation, in particular for augmenting existing film material with information about directional illumination. We have described some ideas on this topic in Section 4, but more sophisticated approaches should be possible.

At the moment, we focus on entertainment-style applications, where user preference is arguably all that matters. An interesting topic for future work is to analyze whether the system can also be helpful in task-oriented applications, for example ones that require navigation in space. Our user survey indicates that this might be possible since the dynamic illumination can help with orientation. However, more studies are required to fully assess the potential of the proposed method in such applications.

On the hardware side, several variants of the current system are possible. At the moment, we use 24 individually packaged LED light sources. One could easily imagine repackaging those lights into a single housing with multiple independently controllable spotlights, to be hung from the ceiling. Such a package could also contain light sensors for the calibration, eliminating the need for an external camera. Since LED lights are becoming more and more popular as standard room illumination, one could also imagine directly plugging into an existing home automation system to control those light sources. We envision that such systems could be used in home theaters and gaming environments, while screening rooms and higher end systems would use dedicated light sources such as in our system.

7. Acknowledgments

We would like to thank Electronic Arts for providing the image material from “Need For Speed Underground 2”. We would also like to thank Paul Debevec and the Computer Graphics group at K.U.Leuven for the HDR environment maps and Rafal Mantiuk for the HDR driving video. The

first author was supported by an ATI Technologies Fellowship.

References

- [AWW03] ARTUSI A., WIMMER J. B. M., WILKIE A.: Delivering interactivity to complex tone mapping operators. In *Eurographics Symposium on Rendering* (2003), pp. 38–44. 2
- [BB67] BARTELSON C. J., BRENEMAN E. J.: Brightness perception in complex fields. *Journal of the Optical Society of America* 57, 7 (Mar. 1967), 953–957. 2
- [Che95] CHEN S. E.: Quicktime VR – an image-based approach to virtual environment navigation. In *Proc. of ACM SIGGRAPH '95* (1995), pp. 29–38. 7
- [CNSD93] CRUZ-NEIRA C., SANDIN D. J., DEFANTI T. A.: Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proc. of ACM SIGGRAPH '93* (1993), pp. 135–142. 2
- [CON99] CABRAL B., OLANO M., NEMEC P.: Reflection space image based rendering. In *Proc. of ACM SIGGRAPH '99* (1999), pp. 165–170. 6
- [DD00] DURAND F., DORSEY J.: Interactive tone mapping. In *Eurographics Workshop on Rendering* (2000), pp. 219–230. 2
- [DD02] DURAND F., DORSEY J.: Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2002)* 21, 3 (2002), 257–266. 2
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. of ACM SIGGRAPH '98* (1998), pp. 189–198. 6
- [DeM72] DEMARSH L. E.: Optimum telecine transfer characteristics. *J. of the SMPTE* 81, 10 (Oct. 1972), 784–787. 2
- [DWT*02] DEBEVEC P., WENGER A., TCHOU C., GARDNER A., WAESE J., HAWKINS T.: A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (Proc. of SIGGRAPH '02)* 21, 3 (2002), 547–556. 2, 6
- [FPSG96] FERWERDA J. A., PATANAIAK S. N., SHIRLEY P., GREENBERG D. P.: A model of visual adaptation for realistic image synthesis. In *Proc. of ACM SIGGRAPH '96* (1996), pp. 249–258. 2
- [HS99] HEIDRICH W., SEIDEL H.-P.: Realistic, hardware-accelerated shading and lighting. In *Proc. of SIGGRAPH 1999* (1999), pp. 171–178. 5
- [LCS04] LEDDA P., CHALMERS A., SEETZEN H.: A psychophysical validation of tone mapping operators using a high dynamic range display. In *Symposium on Applied Perception in Graphics and Visualization* (2004). 1
- [LRP97] LARSON G. W., RUSHMEIER H., PIATKO C.: A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. on Visualization and Computer Graphics* 3, 4 (1997), 291–306. 2
- [MKMS04] MANTIUK R., KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Perception-motivated high dynamic range video encoding. *ACM Transactions on Graphics (Proc. SIGGRAPH '04)* (2004), 733–741. 6
- [NBB04] NAYAR S. K., BELHUMEUR P. N., BOULT T. E.: Lighting sensitive display. *ACM Transactions on Graphics* 23, 4 (2004), 963–979. 2
- [NN04] NISHINO K., NAYAR S. K.: Eyes fo relighting. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2004)* 23, 3 (2004), 704–711. 6
- [PFFG98] PATTANAIAK S. N., FERWERDA J. A., FAIRCHILD M. D., GREENBERG D. P.: A multiscale model of adaptation and spatial vision for realistic image display. In *Proc. of ACM SIGGRAPH '98* (1998), pp. 287–298. 2
- [Phi] PHILIPS: Ambient light technology. <http://www.flattv.philips.com/>. 2
- [PTYG00] PATTANAIAK S. N., TUMBLIN J., YEE H., GREENBERG D. P.: Time-dependent visual adaptation for fast realistic display. In *Proc. of ACM SIGGRAPH 2000* (2000), pp. 47–54. 2
- [RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2002)* 21, 3 (2002), 267–276. 2
- [Sch94] SCHLICK C.: Quantization techniques for visualization of high dynamic range pictures. In *Proc. of Eurographics Workshop on Rendering '94* (1994), pp. 7–20. 2
- [SHS*04] SEETZEN H., HEIDRICH W., STUERZLINGER W., WARD G., WHITEHEAD L., TRENTACOSTE M., GHOSH A., VOROZCOVS A.: High dynamic range display systems. In *ACM Transactions on Graphics (Proc. of SIGGRAPH '04)* (Aug. 2004), pp. 760–768. 1, 2, 3, 6
- [SMGG01] SLOAN P.-P., MARTIN W., GOOCH A., GOOCH B.: The lit sphere: a model for capturing NPR shading from art. In *Proc. Graphics Interface 2001* (2001), pp. 143–150. 6
- [SSS00] SCHEEL A., STAMMINGER M., SEIDEL H.-P.: Tone reproduction for interactive walkthroughs. In *Computer Graphics Forum (Proc. Eurographics)* (2000), pp. 301–312. 2
- [TT99] TUMBLIN J., TURK G.: LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Proc. of ACM SIGGRAPH '99* (1999), pp. 83–90. 2



Figure 11: Left: photograph of the room housing our system, with all lights switched on. Center: illumination programmed to resemble the Grace Cathedral environment. Right: A user viewing the Grace Cathedral environment on an HDR display in a room illuminated by our system. Note how the room illumination is consistent with the virtual environment shown on the screen.