Simulating Realistic Light Levels in Next Generation Image Generators

Brett Chladny
Renaissance Sciences Corporation
Richardson, TX
bchladny@rscusa.com

Kenny Hebert
Renaissance Sciences Corporation
Birmingham, AL
khebert@rscusa.com

Brad Colbert
Renaissance Sciences Corporation
San Diego, CA
bcolbert@rscusa.com

ABSTRACT

The dynamic range of modern day display systems have greatly improved over the past few years. However, they are still not capable of displaying the full intensity range that the human eye is capable of perceiving. The computational power of Graphical Processing Units (GPUs) that are incorporated into modern day Image Generators (IGs) have greatly improved over the past few years, particularly when operating on floating point values. Rendering the entire scene as accurate in-band radiance values enables sophisticated processing to be applied that can help compensate for the limited dynamic range of modern day display systems. The resulting realism can significantly improve training when high contrast scene content is present. Examples include landing on an aircraft carrier that is steaming into the sunset and spotting an entity that is in the direction of the sun. Furthermore, rendering light points using in-band radiance values improves training by providing pilots with realistic visual representations that can take the display's limited dynamic range into account. This can be critical to training when displaying navigational, runway, and anti-collision lights is required. This paper shows results from a new experimental IG that incorporates both accurately rendering the scene using floating point radiance values as-well-as post processing the resulting values to compensate for the dynamic range of the display system. This processing is accomplished by mimicking various aspects of the Human Vision System while still maintaining the commonly required 60 Hz update rate.

ABOUT THE AUTHORS

Brett Chladny has a Masters degree in Computer Science from the University of Missouri – Columbia. He began his career in 1998 at Silicon Graphics creating custom applications for customers. In 2002, Mr. Chladny took a position with MultiGen-Paradigm where he became the technical lead for the F-16 NVG and IR program at Luke AFB. In 2006, Mr. Chladny began working for Renaissance Sciences on the R&D team to develop methods to improve the physical accuracy of sensor and visual simulations. Mr. Chladny is now the Principal Engineer for an ongoing Navy Phase II SBIR project and Project Manager/Lead Engineer for the Deployable Sensor Scene Simulation System (DS4) application, Interservice Common Sensor Model (ICSM) software module and the IG component of NAVAIR's next generation Reduced Oxygen Breathing Device (ROBD) and Carrier Approach Landing Fidelity (CALF) projects.

Kenny Hebert has over 20 years of technical and product management experience in the Defense Intelligence and M&S domains. Prior to his current position as Vice President of Simulation Products with RSC, Mr. Hebert was the Product Manager for the Content Creation products at Presagis USA and a Regional Manager and Application Engineer with TERREX Inc. In addition, Mr. Hebert holds a Bachelors and Masters of Science degree from the University of Southern Mississippi and has over 15 years' military experience as an Intelligence Analyst with the United States Marine Corps.

Brad Colbert has over twenty years of experience in the research and development of physics-based sensor simulation and computer human interfacing. His current research is focused on bringing real-time physically accurate multi-spectral environments to simulations. Of specific interest to Mr. Colbert is both day and night lighting conditions, light interaction with the atmosphere, and modeling the photopic response as a sensor. Mr. Colbert holds a B.S. in Computer Science and Physics and is published in various papers.

Simulating Realistic Light Levels in Next Generation Image Generators

Brett Chladny
Renaissance Sciences Corporation
Richardson, TX
bchladny@rscusa.com

Kenny Hebert
Renaissance Sciences Corporation
Birmingham, AL
khebert@rscusa.com

Brad Colbert
Renaissance Sciences Corporation
San Diego, CA
bcolbert@rscusa.com

INTRODUCTION

Modern day computer hardware and graphics computing capabilities continue to increase at an exponential rate. As the hardware advances, Image Generator (IG) vendors are continuously reevaluating the best way to make use of these new performance improvements. Some vendors may choose to use these new hardware improvements to increase the number of polygons that are capable of being presented in the scene while maintaining 60Hz. Other vendors may choose to page more high resolution imagery from disk, and yet others, may use the hardware improvements to generate database content at run-time. While all of these approaches to improving simulation fidelity and realism have their unique benefits (and drawbacks) the authors of this paper have chosen to focus on a different approach to improve realism and fidelity. The authors have chosen to focus on exploiting the modern advances in hardware and graphics processing to mimic the human system and improve the *quality* of each pixel with enhanced atmospheric and illumination calculations in order to improve the realism of these pixels and their final red, green, and blue luminance levels on the display's surface.

The ultimate goal of the author's research is to improve the realism of visual and sensor simulations by trying to imitate the effects that our human visual system perceives in the real-world environment and display them in a simulated environment. In the real world, the human visual system is capable of adapting to vastly different luminance levels (i.e., light levels) that can extend to a range of up to 14 orders of magnitude (Matković, 1997). For instance, the sunlight on a clear summer day can be as much as 10 million times more intense than the illumination from the moon later that night (Ferwerda, et. al., 1996). Figure 1 shows the large range of luminances in the environment and some associated visual parameters as an example. Therefore, in order to mimic how our visual system perceives the vast ranges of illumination levels in a simulation, it seems logical that we would replicate the exact same luminance levels in the simulation display system. However, while physics-based rendering methods enable the accurate simulation and distribution of light energy rendered in day and night scenes, most simulation display systems are only capable of producing an illumination range of approximately 3 orders of magnitude (compared to the 14 orders of magnitude perceived by the human vision system). Therefore, accurate physics-based rendering does not automatically guarantee the realistic visual appearance of a displayed scene. Additional rendering techniques are required to mimic certain effects experienced in our human visual system (i.e., glare and adaption) while accounting for the limitations of the display.

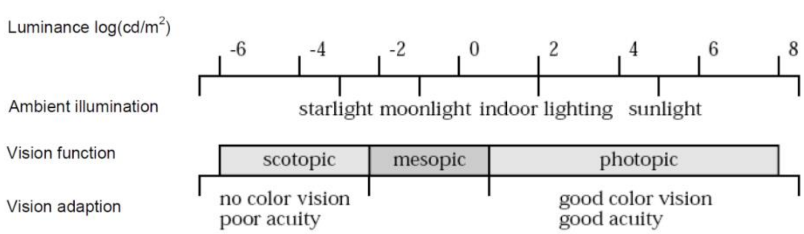

Figure 1: Range of Luminances and Visual Parameters. After Ferwerda, et. al., (1996)

This paper discusses how the authors have improved the realism of visual simulation by emulating the adaptation of the human visual system in an experimental IG being used in the Navy's Carrier Approach and Landing Fidelity (CALF) project. This involves the physics-based modeling of light propagation in the ambient environment in High Dynamic Range (HDR) while also modeling the variations in human perception in order to view the results on Low Dynamic Range (LDR) displays. The prediction of visibility and appearance of scene features, using our approach, provides significant improvements in realism and enhances the overall quality of training as can be seen in Figure 2.

Figure 2: Rendering Results from the Experimental IG

BACKGROUND AND RELATED WORKS

Before we can begin to discuss the implementation of our approach, it is first necessary to have a brief understanding of how the human visual system works and the mechanisms involved in the perception and adaptation to the different illumination levels that we experience in the real-world. The adaption of the human vision system is achieved through a combination of mechanical, photochemical, and neural processes in the visual system. The pupil, the rod and cone systems, and changes in neural processing play a major role in the adaptation mechanism. However, not all intensity levels are perceived equally well, and adaptation does not happen instantaneously.

It is convenient to think of the human vision system as a complex 'biological sensor', actually two sensors (rods and cones), that adapts to ambient illumination with the three vision functions:

- scotopic vision, that is active at a luminance range of approximately 10−6 to 10-2 cd/m2
- mesopic vision, that is active at a range of ~10-2 to 1 cd/m2, and
- photopic vision, that is active at a range of ~1 to 106 cd/m2 (Björn, 2002).

Figure 3 (below) illustrates how the spectral 'sensor' response curves of rods and cones change for the scotopic, mesopic, and photopic vision functions. Knowing and simulating the state of this biological sensor is a key component in determining what an observer should be able to distinguish in the virtual scene under different simulated lighting conditions. We contend that emulating the scotopic, mesopic, and photopic vision functions, the differences in perception at different luminance levels, including the dynamic characteristics, provides significant training advantages.

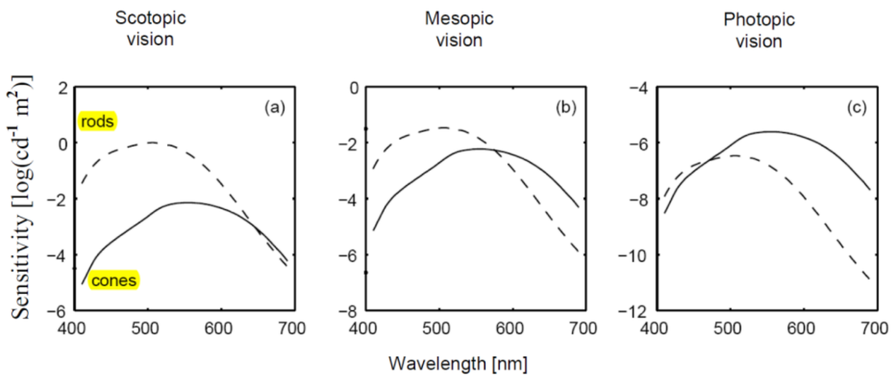

Figure 3: Spectral Response Curves of Rods and Cones for a) Scotopic, b) Mesopic, and c) Photopic Vision. After Ferwerda, et. al., (1996)

One of the most important aspects in modern simulations and serious games applications is the ability to accurately display realistic synthetic environments and scene content while maintaining real-time rendering performance. Accurate modeling of natural illumination and how the atmosphere modulates and contributes to this has become increasingly possible. Some of the reasons why this is advantages are as follows:

- Realistic radiance values tend to produce expected, reproducible results.
- When "tuning" a scene that starts with real-world values, it is easy to determine how much each "tweak" deviates from reality, which in turn helps determine the resulting numerical accuracy of the system.
- Less visual tricks are required to produce desired final results that appear and behave like the natural environment.

Moreover, the accurate and realistic display of synthetic environments are highly dependent on the correct depiction of illumination levels and atmospherics of the virtual scene. One of the growing methods for replicating the dynamic lighting conditions of the real-world in visual simulations and gaming platforms is to use High Dynamic Range Rendering (HDRR).

Most of today's digital images are stored with 24-bits per pixel, which allows for values from 0-255 for each of the red, green, and blue channels. However, in our natural (real-world) environment, the light intensities and brightness levels greatly exceed what can be accurately represented with 8-bits per-channel digital images. While it is convenient to use the integer values 0-255 to represent RGB pixel values, there are not enough gradations or steps to accurately represent the wide range of real-world light levels. Therefore, to more accurately replicate the natural environment and the way our Human Vision System (HVS) interprets the scene, color and luminance values need to be represented as floating point values that cover a much larger range (ie., the ratio between the highest and lowest values of luminance in a scene or image and is usually expressed in orders of magnitude) (Luksch, 2007). High Dynamic Range Rendering provides more realistic and convincing visual effects since the pixel values more accurately reflect the light levels of the real-world.

Although HDR imagery and rendering techniques are relatively new, the use of HDRR has been rapidly expanding in the gaming and simulation industries; largely as a result of the equally rapid advancements in GPU capabilities. This has prompted significant research in recent years on different real-time HDRR techniques. For instance, Petit & Bremond (2009 & 2010) proposed a noteworthy method to compute in real-time a HDR illumination in virtual environments that was based on physical lighting models. Their method allowed for the re-use of existing virtual environments as input to compute the HDR images in photometric units and rendered with a tonemapping operator on a Low Dynamic Range (LDR) display. Their approach, which conducted the HDR computation and tonemapping operations in *OpenSceneGraph* using pixel shaders, provided enhanced perceptual realism of the scene at a low cost to performance and operation in the rendering pipeline. They demonstrated the results of their HDR pipeline with an interactive driving simulator that provided enhanced (realistic) lighting.

Similarly, Delacour (2009) developed a simulation framework with a display and environment that was designed to mimic the visual perception that a human would have of the scene in real life. His goal was to build a simulator that was based entirely on physics and the physiological perception of light using HDRR techniques combined with the interactions of the physical simulation interfaces. In addition to providing accurate displays of the synthetic environment in the simulation framework proposed by Delacour (2009), his approach also provided simulation designers the ability to accurately simulate the lighting within the simulation platform to improve the decisions on Human Machine Interface design.

HDRR provides the ability to represent each pixel with a larger contrast ratio and dynamic range to more accurately represent real-world conditions. The intensity of the pixels can even have a dynamic range greater than the range available on the display. In the cases where the values of the HDR exceed the values of the display, a transformation is usually applied in the rendering pipeline to convert the HDR values to LDR values for the display (Krawczyk, Myszkowski, & Seidel, 2005). This transformation is commonly referred to as tonemapping and is done using tonemapping operators (TMO) (Yoshida, et. al., 2005).

The importance of tonemapping operations in HDRR is widely understood and many research studies have already proposed algorithms to perform efficient TMOs that achieve real-time performance (Devlin, et. al., 2002). In one example, Krawczyk, et. al., (2005) presented an efficient way to maintain the most significant perceptual effects in HDRR (i.e., glare, bloom, visual acuity, night vision, etc.,) with real-time local contrast compression (tonemapping) for Low Dynamic Range displays. The result was a unified model that included all of those effects into a common framework that ran on common hardware and could be added into any real-time rendering system to enhance the realism and believability of the displayed environment.

OVERVIEW

The authors' solution divides simulating the human vision system into two separable components. First, the scene is rendered using true radiance values as opposed to the unitless 0 to 1 values typically used for out-the-window scenes. To support the human vision system, the authors render the scene using radiance values in the photopic red, green, and blue wavebands. The rendered scene is captured as floating point values and processed in a way that simulates the human visual system's response to these radiance values. The radiance value to final output pixel values incorporates graphics processing similar to that performed by a traditional sensor post-processor. It is worth noting that much of the processing involved in simulating the human vision system is similar to that used to simulate Night Vision Goggles (NVGs) and Forward Looking InfraRed (FLIR) sensors.

The authors have developed two software libraries, called SERE and ICSM, to facilitate these two tasks. The intent of these libraries is that they can be dropped into virtually any scene graph to improve visual and sensor representations of the virtual scene. SERE provides two basic functions: direct rendering of the sky, sun, stars, etc. and providing shader code to be used by the scene graph to render true radiance values. In the Figure 4 below, SERE's domain is represented by the arrows in the scene and those that point to the observer. SERE supports rending in the wave lengths from photopic blue through long-wave infrared. ICSM facilitates the capturing of the scene into a floating-point Frame Buffer Object (FBO). ICSM also provides the math model and graphics processing needed to emulate the human vision system, NVG sensors, and FLIR sensors. In the figure below, this domain is referred to as the observer (detector) transformations on the right.

While SERE and ICSM are crucial to the work presented in the paper, the focus of this paper is on leveraging these libraries in a training environment as opposed to the implementation details of these libraries. SERE and ICSM have been integrated into a run-time experimental IG being used as the IG software for the Navy's CALF project. This is an Office of Naval Research (ONR) sponsored project that seeks to determine, among other things, the level of visual fidelity that is required to perform carrier landing training effectively. Beyond the experimental IG's use of SERE and ICSM, it behaves similar to other fielded Navy IGs. It provides a CIGI interface to allow host control and can synchronously render across multiple computers in order to generate a contiguous image that provides a wide field of view.

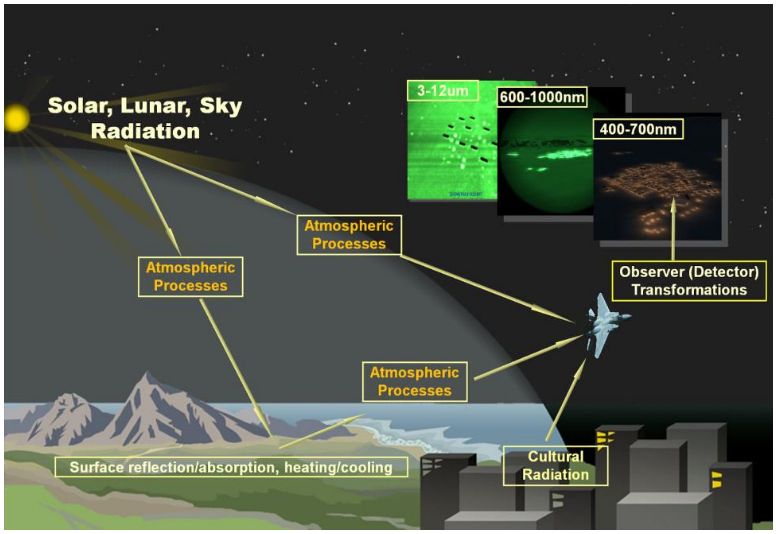

Figure 4: Components of Radiometric Simulations

REAL-WORLD ILLUMINATION LEVELS

In order to model the lighting environment using values that are representative of the real world, let us first discuss the levels that need to consider. Full daylight is roughly 10 orders of magnitude greater than clear star light conditions. Dividing the environments into two operational conditions, night and day respectively, large spans of intensities still remain. Excluding the extremes, overcast conditions in star light and manmade lighting, the illumination spans roughly four orders of magnitude. With manmade lighting (typical) it can be around 7 orders or greater. During daytime on a clear day, sunrise and sunset to high noon, we encounter roughly 3 orders of magnitude in illumination. Table 1 illustrates typical illumination levels.

Table 1. Typical Illuminance Levels at the Ground (RASC, 2008)

Condition	**Illumination Level (*lux*)**	**Illumination Level (*fc*)**
Clear night sky (no Moon)	0.00005 - 0.002	0.000005 - 0.0002
Clear urban sky with light pollution	0.015	0.0014
Beginning of astronomical twilight	0.1	0.009
Overcast urban sky with light pollution	0.15	0.015
Full Moon	0.2 (typical) to 1 maximum	0.019 (typical) to 0.093 max
Urban road with artificial illumination	2	0.19
Open parking lot	11 - 22	1 - 2
Car dealership lot	200	19
Sun(rise/set) on a clear day (ambient illumination).	400	40
In shadow at noon, clear sky	~5000 - 10000	500 - 1000
Full sunlight	100,000	~10,000

DISPLAY LIMITATIONS

Display systems, whether direct view as in flat panel displays or indirect view as in projected displays, are greatly limited in the dynamic range of the intensities that they can produce. The next few paragraphs are mainly intended as background information and to paint the picture of the problem space and are not intended to be a full analysis of the challenges faced by display systems.

Manufacturers of projection systems commonly tout large dynamic range capabilities for their products. These ranges are measured in strictly controlled scenarios and typically do not stand up to the challenges of actual use. It is also common for display manufacturers to quote numbers representing their display systems dynamic contrast. Dynamic contrast is the ratio of the brightest instance (frame) as compared to the darkest instance (frame) usually using the full field of view. A contrast measurement that better reflects the real world use cases is referred to as static contrast. Static contrast is the measurement of the brightest part of a single frame as compared to the darkest part of the same frame. Some numbers for dynamic contrast from leading manufacturers span 6 orders of magnitude. However, these same projectors are barely capable of having static contrast ranges of 4 orders of magnitude on a flat surface and 2 orders of magnitude in cases where screen cross illumination is heavy.

In practice, 4 fL (foot-Lamberts) is used as a typical calibration luminance in many simulation installations. Many factors push users to select this value. One is that this value has been found to be a good setting that allows for both day and night time scenarios with little or no projector setting changes. Another factor is to prolong the life of the bulb in the projectors (for those that have bulbs). To put this 4 fL luminance value into perspective, below is an equation for asphalt viewed at night under artificial lighting conditions in a car dealership parking lot:

$$L_v = E_v \times R$$

Where:

L_v is the luminance, in foot-lamberts [fL]

E_v is the illuminance, in foot-candels [fc]

R is the reflectivity, expressed as a fractional number

3.8 fL = 19 fc * 0.2 (20%)

* 20% is a typical reflectance value of older asphalt.

The limited dynamic range of display system may seem problematic, but once the user has viewed the simulation for some time, the human visual system tends to adapt and accept this as a daytime image. Of course, there are systems currently deployed or coming online which are calibrated to higher intensities for daytime use and can be adjusted for night time training. However, these increases do not represent orders of magnitude improvements, but simply modest incremental improvements. One reason for this is that increasing the top end brightness commonly increases the bottom end (black level) as well. In bright scenarios, this may not be as noticeable, but for night time scenes the black levels can become visible and distracting.

SIMULATING THE HUMAN VISION SYSTEM

The experimental IG utilized for this effort uses ICSM to simulate the human vision system. In many ways, the ICSM software plays a role similar to that of traditional hardware sensor post-processor systems, except that it is largely implemented on the GPU. The input interface from the experimental IG to ICSM is largely accomplished by "handing off" at-aperture radiance frames. In this interfacing approach, the implementation details associated with scene generation are largely hidden from ICSM (See Figure 5).

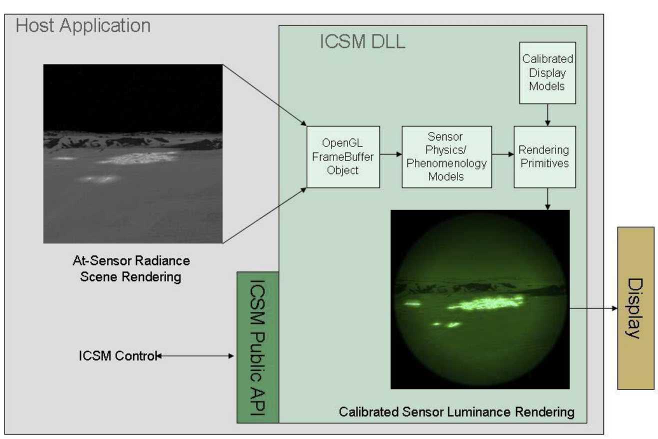

Figure 5: Sensor as Post-Processor

ICSM provides many features as part of its simulation of the Human Vision System. Since ICSM is not the focus of this paper, only a few of the more relevant feature are discussed below

Gain

HDRR commonly uses tone mapping for converting from floating point pixel values to 24-bit integer values that are sent to the display system. ICSM performs this task by computing the average scene radiance and approximates what the observer should see in their simulated adaption state on the target display system. It should be pointed out that if the display system is capable of generating real-world luminance levels that are present in a given scene, ICSM is able to generate pixel values that result in accurate luminance levels on the display's surface. Most display systems are not capable of displaying clear starlight or daytime illumination levels, so some amount of tonemapping is generally required.

Night Adapted Vision

The human vision system functions significantly different at night then it does during the day. Acuity and color perception are significantly different at night than they are during the day. If the display system is incapable of reproducing all of the required values of a night-time scene or the observer is not able to adapt to night time vision, due to the training room's ambient lighting conditions, ICSM can facilitate simulating reduced acuity and a shift in color perception.

Glare

Full scene glare or bloom is a technique that video games have made popular in recent years. In some implementations, it produces surreal looking image that are less stark than when not employing this technique. Although full scene bloom can produce impressive looking results, it can also produce unrealistic images when overdone. The premise behind this technique is valid. When our eyes view something that is too bright for our current adaption state, light from this source will scatter within our eyeball and the light will appear to bloom around this object. Our eye is an imperfect imaging sensor. Our lens has imperfections in it and the liquid inside and outside our eye is not perfectly clear. The older we get, the more these imperfections become noticeable.

The authors contend that glare is an important part of some training exercises, such as formation flying and landings on an aircraft carrier. In the case of formation flying, it is never desirable to have the sun directly behind one of your wingmen, as this will make it nearly impossible to see them in the real world. A similar problem occurs when landing on an aircraft carrier that is steaming into the rising or setting sun. Without adding a glare effect to the simulation of these two tasks, the trainee will be able to see his wingman and the landing aids on the carrier deck without issue. This is because the display system cannot represent the true intensity of sun. Therefore, current simulators commonly provide negative training in these situations. The left half of Figure 6 contains an image of the glare effect that is part of ICSM. In this image, the bright sun is reflecting off the ocean, partially obscuring the landing aid upon approach. The amount of bloom that is applied must be scaled according to the capabilities of the display system.

Another common occurrence of glare can be light points, particularly bright light points at night. A good example of this is the bloom we see around approaching car headlights on a dark night. Ideally, a display system would be able to generate luminance values that matched that of light points in the real world. Unfortunately, such display systems do not exist for visual simulations. However, knowledge of the true radiant intensity of a light point and the current adaption state of the observer can enable software to place an appropriate amount of bloom around light points, much like when simulating an NVG sensor. The right half of Figure 6 shows the landing lights on a carrier deck turned up so bright that they bloom. This effect must also be tuned to the display system's capabilities.

Figure 6: Glare Effect

Simulating Analog Output

Light levels in the real-world are analog, whereas display systems are commonly digital devices that have a limited number of intensity values that they can reproduce. The intensity differences between these discrete values can sometimes be seen when simulating subtle gradations, such as the sky near the horizon.

When you closely examine a photographic image that contains a smooth gradation, you commonly see a small amount of noise in the image. Without this noise, the image would likely contain color banding. In the authors' solution, the entire scene is rendered using floating point radiance values into a buffer and then converted into quantized discrete integer values. By adding a very small amount of noise when converting from a floating point values to integer values, banding artifacts are avoided.

The authors have also developed a method for adding a dynamic noise pattern each frame, which not only prevents the observer from detecting the noise pattern but also enables the observer to perceive intensity values that the display system is unable to produce. As an example, let's say the desired pixel value is (127.5, 127.5, 127.5). Let's assume the display hardware can only reproduce the pixel values (127, 127, 127) and (128, 128, 128). By showing the observer the pixel value (127, 127, 127) on all even frames and (128, 128, 128) all odd frame, the observer would perceive the intended value of (127.5, 127.5, 127.5). By extrapolating this idea, the authors have achieved perceived results that are significantly closer to the original floating point values in the rendered scene and have eliminated banding artifacts.

INTEGRATION INTO NAVY EXPERIMENTAL IMAGE GENERATOR

The experimental IG used for this effort is a software application that was developed by RSC to provide simulation of the human vision system, light intensifiers, and thermal imagers. The goal was to create an application that performs complex radiometric computations and sophisticated viewer transforms while maintaining real-time frame rates (60 Hz) at high resolutions (2560x1600). There is no question that SERE and ICSM consume CPU and GPU resources. However, the authors' work has shown that modern day off-the-shelf gamer class computer hardware is fast enough to render required database content, perform this processing, and still meet real-time simulation requirements.

This IG has been implemented to enable the underlying scene graph to be changed relatively easily, while still providing many of the same features, including its CIGI interface and synchronization capabilities. The underlying rendering engine is currently OpenSceneGraph (OSG), although the authors have also successfully leveraged Diamond Visionic's GenesisRTX is this role.

SERE has been integrated predominantly in two ways. First, an OSG specific SERE adapter has been created. This adapter facilitates the construction, loading, and updating of SERE shaders that can be attached to any OSG node in the scene graph. Second, SERE is told when it should directly render the sky dome, sun, and stars at the beginning of each frame.

Figure 7 illustrates how ICSM's processing is integrated into the experimental IG's frame loop. The time axis starts on the left. The light gray background indicates that the IG's code is being executed; the dark green background indicates that ICSM's code is being executed; and a red line around a buffer indicates that buffer is the active OpenGL render target. Which buffer is active, and the contents of the buffer are based on the state of the IG at the end of each block in the diagram. It is important to note that when the IG is rendering the scene, the ICSM buffer is active. Similarly, when ICSM is rendering, the frame buffer is the active buffer.

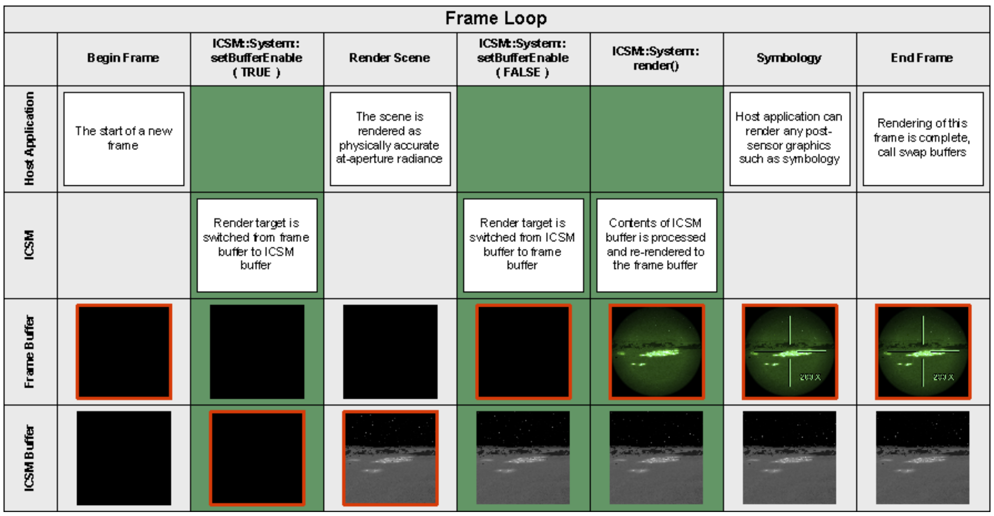

Figure 7: Frame Loop

PERFORMANCE

In order to minimize latency, the IG uses a single thread for all app, cull, and draw processing. The resolution used for the CALF project is 2560x1600 and the GPU employed is the NVIDIA GTX 680. As stated above, the IG performs accurate floating point radiance based rendering and simulates the human vision system. All of the processing that is required to construct, render, post-process, and syncronize a single frame consistently stays below 11 ms on each of the 6 IGs used for the CALF project.

CONCLUSION AND FUTURE WORK

The goal of the authors' portion of the ONR project was to create an experimental IG that was capable of incorporating aspects of the human vision system to increase the realism and accuracy of the simulated environment while still maintaining the commonly accepted requirements for frame rate and resolution. To achieve this, the authors emulated the adaption of the human visual system in the experimental IG to accurately render the scene using floating point radiance values and post-processing techniques to compensate for the limited dynamic range of typical display systems. This involved the physics-based modeling of light propagation in the simulated scene in High Dynamic Range (HDR) while also modeling the variations in human perception in order to view the results on Low Dynamic Range (LDR) displays. The results of this ongoing effort have shown significant improvements in the visual fidelity and perceptual accuracy (i.e., realism) of the rendered scene.

The software and experiment discussed in this paper is still on-going. Currently, only one of potentially three experiments has been conducted, and the data from the first experiment is still being analyzed. The authors plan to use the results of the first experiment to identify what aspects of the visual scene still need improvement. Furthermore, the authors are hopeful that this paper, and the results of the mentioned experiments, will help push the state of the art for realism in next generation Image Generators.

ACKNOWLEDGEMENTS

The authors would like to thank the Office of Naval Research for providing the funding for this research. However, the views presented in this paper are those of the authors and do not necessarily represent the views of DoD or its Components.

REFERENCES

Björn, L. O. (2002). *Photobiology: The Science of Light and Life*. New York:Springer.

Cohen, J., Tchou, C., Hawkins, T., & Debevec, P. (2001). *Real-Time high dynamic range texture mapping* (pp. 313-320). Springer Vienna.

Darling, B. A., Ferwerda, J. A., Berns, R. S., & Chen, T. (2011). Real-time multispectral rendering with complex illumination. In *Color and Imaging Conference* (Vol. 2011, No. 1, pp. 345-351). Society for Imaging Science and Technology.

Delacour, J. (2009). *Light and Human Vision Based Simulation Technology*. OPTIS LA FARLEDE (FRANCE). Retrieved May 30, 2014, from http://ftp.rta.nato.int/public/PubFullText/RTO/MP/.../MP-MSG-069-03.doc.

Devlin, K., Chalmers, A., Wilkie, A., & Purgathofer, W. (2002). Tone reproduction and physically based spectral rendering. *Eurographics 2002: State of the Art Reports*, 101-123.

Ferwerda, J. A., Pattanaik, S. N., Shirley, P., & Greenberg, D. P. (1996). A Model of Visual Adaptation for Realistic Image Synthesis. In *Proceedings of the 2009 SIGGRAPH 23rd annual conference on Computer graphics and interactive techniques* (pp. 249-258). ACM.

Krawczyk, G., Myszkowski, K., & Seidel, H. P. (2005, May). Perceptual effects in real-time tone mapping. In *Proceedings of the 21st spring conference on Computer graphics* (pp. 195-202). ACM.

Luksch, C. (2007). Realtime HDR rendering. *Graduation Project. Institute of Computer Graphics and Algorithms, TU Vienna*. Retrieved May 30, 2014, from www.cg.tuwien.ac.at/research/publications/2007/Luksch_2007_RHR/Luksch_2007_RHR-RealtimeHDR%20.pdf

Mantiuk, R., Krawczyk, G., Mantiuk, R., & Seidel, H. P. (2007, February). High dynamic range imaging pipeline: Perception-motivated representation of visual content. In *Electronic Imaging 2007* (pp. 649212-649212). International Society for Optics and Photonics.

Matković, K. (1997). Tone Mapping Techniques and Color Image Difference in Global Illumination. Unpublished Ph.D. Thesis, Technische Universität Wien, Vienna, Austria. Retrieved May 30, 2014, from http://ftp.cg.tuwien.ac.at/research/theses/matkovic/matkovic.pdf

Nijasure, M., Pattanaik, S., & Goel, V. (2005). Real-time global illumination on gpus. *journal of graphics, gpu, and game tools*, *10*(2), 55-71.

Petit, J., Brémond, R., & Vienne, F. (2009, October). High dynamic range rendering for driving simulation. In *Proc. RSS 2009: International conference on Road safety and simulation, Paris, France* (pp. 5-7).

Petit, J., & Brémond, R. (2010). A high dynamic range rendering pipeline for interactive applications. *The Visual Computer*, *26*(6-8), 533-542.

Royal Astronomical Society of Canada (RASC) (2008). Guidelines for Outdoor Lighting in Dark Sky Preserves (RASC-DSP-GOL). Adopted by the RASC in March 2008. Retrieved June 1, 2014, from www.darkskiesawareness.org/files/RASC%20DSP%20GOL.pdf.

Wang, L., Wang, X., Sloan, P. P., Wei, L. Y., Tong, X., & Guo, B. (2007, April). Rendering from compressed high dynamic range textures on programmable graphics hardware. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games* (pp. 17-24). ACM.

Yoshida, A., Blanz, V., Myszkowski, K., & Seidel, H. P. (2005, March). Perceptual evaluation of tone mapping operators with real-world scenes. In *Electronic Imaging 2005* (pp. 192-203). International Society for Optics and Photonics.