

VALIDATION OF GPU LIGHTING SIMULATION IN NATURALLY AND ARTIFICIALLY LIT SPACES

Nathaniel L Jones and Christoph F Reinhart
Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT

Daylight in buildings is both aesthetically pleasing and a sustainable means of offsetting costs for space conditioning and electric lighting. However, poor use of daylight can cause glare that impedes worker productivity. Traditional means of predicting lighting levels for indoor spaces through simulation are time-consuming, which inhibits exploration of daylighting potential in new buildings.

This paper compares results from two lighting simulation engines, Radiance and Accelerad, with measurements taken in physical spaces. Radiance is an established backward ray tracer that runs on the CPU, whereas Accelerad is a recent porting of Radiance's algorithms for the GPU. Vertical eye illuminance, daylight glare probability, and monitor contrast ratio serve as metrics for comparison. Radiance and Accelerad produce similar errors in visual comfort metrics of around 10%, and Accelerad generates solutions between 3 and 24 times faster than Radiance in the tested scenes. These speedups are expected to scale up on new generations of graphics hardware.

INTRODUCTION

Daylight in interior spaces is viewed as aesthetically pleasing, beneficial to maintaining alertness and productivity, and a sustainable means of offsetting heating and electric lighting costs. However, misuse or overuse of daylight can lead to veiling glare on monitors and discomfort or disability glare that impedes worker productivity. Currently, short of building elaborate physical mock-ups, predictions of glare conditions can only be made by time-consuming ray tracing simulation.

This study compares two simulation methods with measurements taken in physical spaces. The simulation engines used are Radiance, a well-validated central processing unit (CPU)-based ray tracing engine originally developed by Greg Ward at Lawrence Berkeley National Laboratory (Larson & Shakespeare, 1998), and Accelerad, a graphics processing unit (GPU)-based tool developed by the authors (Jones & Reinhart, 2014a; Jones & Reinhart, 2014b). Accelerad produces physically-based images and simulated sensor data for daylit and artificially lit scenes with comparable accuracy to Radiance and in much less time.

Simulations performed in Radiance and Accelerad are compared to calibrated high-dynamic range (HDR) photographs taken in two spaces that were furnished to resemble typical workspaces. The first space is windowless and lit by sources of known distribution. The second space is entirely daylit and is captured in photographs taken at 32 times under varying sky conditions. In both cases, a monitor displaying a test pattern is used to assess visual comfort. Direct sunlight and reflections on the monitor screen produced intolerable glare at many times in the daylit space.

Daylighting simulation on the GPU presents several unique challenges. The number of ambient bounces required for ray tracing is high; typically at least five bounces are needed to accurately capture the area that will be perceived as daylit by an occupant. To prevent the simulation time from growing exponentially, irradiance caching is employed to allow reuse of previously calculated diffuse lighting values. However, irradiance caching is not easily parallelizable. Instead, Accelerad performs iterative irradiance caching as a preprocessing step, employing both CUDA[®] and OptiX[™] (Parker, et al., 2010) kernels to select and compute irradiance cache entries. We test in this study whether Accelerad's ray tracing with irradiance caching and multiple ambient bounces performed on the GPU can be as accurate as traditional Radiance methods for predicting visual comfort conditions in a workplace.

BACKGROUND

Since the early days of computer graphics, rendering quality has been judged by comparison to real scenes. In the early 1980's, the Cornell box experiment demonstrated the accuracy of a rendered image by comparison to a photograph (Goral, et al., 1984). Within the computer graphics community, this test has become so ubiquitous that the scene is now frequently used to show the visual plausibility of rendering techniques even without photographic comparison. However, demonstrations of visual plausibility are insufficient to validate physically-based simulations of scene radiance.

Instead, validation work on the Radiance suite of programs has focused up to now on comparison to sensor data with few exceptions. Grynberg (1989) provides side-by-side comparison of Radiance *rpict* visualizations to photographs of a conference room,

but without a quantitative means of comparing the two. Jakubiec and Reinhart (2013) used image-based comfort metrics to assess a space, but compare their simulation results to self-reported occupant comfort rather than photographic evidence of glare. Other comparisons make use of Radiance *rtrace* to simulate sensor readings in buildings (Ng, et al., 2001; Galasiu & Atif, 2002) or controlled environments (Mardaljevic, 1995; Reinhart & Herkel, 2000; Mardaljevic, 2001; Reinhart & Walkenhorst, 2001; Reinhart & Breton, 2009). The use of HDR photography to capture quantitative luminance data has been tested (Inanici, 2006; Inanici, 2010; Van Den Wymelenberg, et al., 2010; Van Den Wymelenberg & Inanici, 2014). However, to the authors' knowledge, no previous study has made quantitative comparisons between simulated and photographed HDR imagery.

When performing a validation study, it is important to establish the acceptable magnitude of error. Daylighting studies rarely report a target accuracy, but a typical expected simulation error may be extrapolated from multiple studies. Ng, et al., (2001) report errors of up to 20% at individual sensors, though whether this error stems from measurement or simulation is unknown. Reinhart and Herkel (2000) compared six simulation engines and found root-mean-square errors (RMSEs) in global illumination ranging from 16% to 63%. Reinhart and Walkenhorst (2001) found errors under 20% and RMSE under 32%, which were later taken by Reinhart and Breton (2009) as acceptable maximums; however, the latter study produced higher errors in 15 of 80 data points. Reinhart and Andersen (2006) reduced error to 9% and RMSE to 19% using advanced modelling and measurement techniques; however, they allow for the possibility of 20% error in daylight simulation results when applying them to energy calculations. The expectation of 20% error in illuminance simulations appears to be typical in daylight studies. Because Radiance uses the same algorithms to calculate image and sensor data, we apply the same acceptable error criterion in our study.

METHOD

Lighting simulation validation studies typically make use of idealized scenes with simple, easily modelled geometry. While these studies are certainly useful for error diagnostic purposes, we take the position of Reinhart and Breton (2009) that simulation of and comparison to "real" spaces is necessary to demonstrate the reliability of simulation tools. Furthermore, simplified models tend to run at faster speeds, so speedups measured on idealized cases may not reflect the performance of software experienced by its end users. We therefore measured and modelled two spaces in the MIT Media Lab for this study. The first, the Perception Lab, is a small windowless room outfitted with dimmable, colour-mixing LED luminaires that allow both the intensity and colour temperature of the room's lighting to be tightly

controlled. The second, the 5th floor lounge, is a naturally lit open-plan space known to experience glare conditions. Desks and a display monitor were added to each space mimic an office environment.

Data Acquisition

HDR photography provides the basis for comparison in this study. A Cannon 5D Mark II camera with a Sigma 4.5mm fisheye lens was set up in each space in front of the monitor at a typical seated head height (1.22 m (48 in.) in the Perception Lab and 1.09 m (43 in.) in the lounge). A scripted camera setting allowed multiple exposures to be captured, which were later composited into HDR images using the program *photosphere* (Anywhere Software, 2014). In the Perception Lab, a single set of images was captured under 4000 K lighting at full intensity. In the lounge, photographs were taken at 15-minute intervals during certain hours over three days to capture the scene under a variety of sky conditions. January 8 and 9, 2014, were clear days, and January 10, 2014, was overcast.

Additionally, luminance, illuminance, and reflectance were measured in each space. A Konica Minolta LS-110 luminance meter was used to measure luminance at marked points on the desk and wall at the time of each photograph; this later allowed each HDR image to be calibrated. The luminance meter was also used to record the luminance of white and black patches of a checkered test pattern on the monitor for determination of monitor contrast. A Konica Minolta TL-1 illuminance meter was used to record desk surface and vertical eye illuminances, as well as to determine the transmittance of windows. Finally, a Konica Minolta CM-2500d spectrophotometer was used to measure the reflectance and specularities of surfaces in both spaces. Post-processing of this data allowed custom Radiance materials to be created for all surfaces in the models of both spaces.

Weather data for the three days of lounge observation was obtained from a Hobo weather station located at the top of a tall building approximately 200 m (~600 ft) from the Media Lab. DAYSIM's *gen_reindl* and Radiance's *gendaylit* programs were used to convert the horizontal solar radiation measured from this station to sky definitions using the Perez all-weather sky model (Perez, et al., 1993).

Significant post-processing was performed on the HDR images to prepare them for comparison with simulated images. A circular crop was applied to remove the interior of the lens assembly visible in the fisheye projection. The cropped HDR photographs were scaled dimensionally to 512 × 512 pixels to match the simulation settings. A vintetting function was measured for the camera lens and applied to the images as in Inanici (2010). Finally, an intensity scale factor was applied to the image based on readings from the luminance meter in order to put the HDR pixel values in units of $W \cdot sr^{-1} \cdot m^{-2}$.

Modelling

Digital models of the Perception Lab and lounge were created in Trimble SketchUp and exported to the Radiance file format using Thomas Bleicher’s *su2rad* plug-in. Best practices were followed to the fullest extent possible in creating the models. Spatial dimensions were obtained from physical measurements and architectural plans of the two spaces. Material parameters for all opaque and transparent surfaces were extrapolated from spectrophotometer and illuminance meter readings, respectively. IES files for the Perception Lab’s LED luminaires at their reported colour temperature were obtained from the manufacturer’s website.

The monitor screens were modelled as Radiance glow materials positioned behind Radiance trans materials. The trans material properties were modelled based on spectrophotometer measurements of the screens when turned off, and the glow intensities were scaled so that transmitted light matched luminance readings of the screens when turned on in a dark room. This differs from the best practice described by Moghbel (2012); however, Accelerad currently lacks support for the Radiance mixture materials required by Moghbel’s model. For our purposes, the simplified monitor model is sufficient because the monitors are only viewed head-on.

Despite adhering to best modelling practices wherever possible, we are aware of numerous potential error sources resulting from the modelling process. Indeed, such errors seem inevitable when matching simulation to reality. For instance, the Perez all-weather sky model, while an accurate generalization, is not a model of the *particular* skies observed in photographs. We cannot therefore expect perfect correspondence between the actual and modelled scene illumination. Similarly, the geometric and material fidelity of the outside environment in the model is low compared to that of the interior. Five neighbouring buildings visible from the lounge were included in the model using data from SketchUp’s 3D Warehouse, but their materials were generalized from a few measurements. It is also worth noting the presence of some artificial light sources in the lounge that could not be turned off. While the illumination from these sources was considered low enough to have a negligible effect on measurements, they nonetheless represent light sources not accounted for in our model.

Simulation

Larson and Shakespeare (1998) recommend Radiance parameters for high accuracy simulation. However, these settings reflect the capabilities of an older generation of computer hardware and produced poor image quality in our Radiance renderings. Consultation with Greg Ward led to the development of settings tuned specifically to the models created for this study. The parameters used for Radiance and Accelerad and listed in Table 1.

Table 1
Default simulation parameters

PARAMETER	CODE	P. LAB	LOUNGE
Ambient accuracy	-aa	0.2	0.05(0.2) ^a
Ambient bounces	-ab	varied	3(5) ^a
Ambient divisions	-ad	1024	2048
Ambient resolution	-ar	64	1024
Amb. super-samp.	-as	0	0
Direct jitter	-dj	0	0.6
Direct relays	-dr	1	3
Direct sampling	-ds	0.3	0.01
Max. ray reflections	-lr	6	6
Min. ray weight	-lw	0.002	0.002
Specular sampling	-ss	0.7	16
Specular threshold	-st	0.1	0.03
Pixel sampling	-ps	4	1
Pixel tolerance	-pt	0.08	0.08

^a Accelerad irradiance caching parameters in parentheses

Additionally, Accelerad introduces a parameter -ac to control the number of ambient values calculated at each bounce and stored in its irradiance cache (Jones & Reinhart, 2014b). This parameter was held constant at 8192 for lounge simulations and allowed to vary for simulations of the Perception Lab.

Simulations were run on two machines. Radiance’s serial implementation of *rpict* was tested on a workstation with a 3.4 GHz Intel® Core™ i7-4770 processor and an NVIDIA® Quadro® K4000 graphics card with 768 CUDA® cores. The Accelerad counterpart, *accelerad_rpict*, was tested on a workstation with a 2.27 GHz Intel® Xeon® E5520 processor and two NVIDIA® Tesla® K40 graphics accelerators with 2880 CUDA® cores each. These assignments were made so that the serial simulations had access to a faster CPU and the parallel simulations had access to more GPU cores.

Comparison Metrics

Many metrics for image comparison, both quantitative and qualitative, could be considered. The gold standard for accuracy might be pixel-per-pixel correlation between images, but this is impractical when comparing models to photographs because minute geometric inaccuracies create large errors. Furthermore, architects are generally not concerned with achieving this level of fidelity in their models. Instead, we consider three metrics that might be directly used by building designers: vertical eye illuminance (E_v), daylight glare probability (DGP), and monitor contrast ratio (CR).

E_v is a measure of the total illuminance reaching the camera sensor. Unlike DGP and CR, E_v is not sensitive to small geometric differences or localized differences in pixel brightness. However, comparison of E_v between images can reveal systematic error in the predicted brightness of a scene. We obtain E_v values from DAYSIM’s *evalglare* program. Because the camera was tilted toward the monitor in tests, these are

not truly “vertical” illuminance measurements, but they serve the same purpose for validation.

DGP assesses the probability that a person situated at the camera position will report glare in the scene (Wienold & Christoffersen, 2006). Previous studies have shown DGP to be a robust metric that is unlikely to yield false positives (Van Den Wymelenberg, et al., 2010; Jakubiec & Reinhart, 2012). DGP is calculated by *evalglare* as follows:

$$DGP = 0.16 + 5.87 \times 10^{-5} E_v + 0.0918 \times \log_{10} \left(1 + \sum_{i=1}^n \frac{L_{s,i}^2 \omega_{s,i}}{E_v^{1.87} P_i^2} \right) \quad (1)$$

where E_v is the vertical eye illuminance calculated from a 180° fisheye projection of the scene, L_s is the luminance of a glare source, ω_s is the solid angle size of that source, and P is the source’s Guth position index, which relates position within the field of view to human eye sensitivity. DGP values below 0.35 may be considered imperceptible, while those greater than 0.45 are considered intolerable.

An occupant’s ability to read from a monitor screen is determined by the contrast between bright and dark pixels. Light reflected by the screen from other sources illuminates bright and dark pixels equally, reducing the contrast from the viewer’s perspective. At the extreme, this causes veiling glare, where the image on the monitor is no longer intelligible. The monitor contrast ratio is represented by:

$$CR = \frac{L_H + L_r}{L_L + L_r} \quad (2)$$

where L_H is the high state luminance of a bright pixel, L_L is the low state luminance of a dark pixel, and L_r is the luminance contribution from reflected light. A minimum CR of 3 (ISO, 1992) or 4 (ISO, 2008) is required to preserve readability. In practice, the numerator and denominator sums of Equation 2 can be measured directly by a luminance meter reading or from a photographic or simulated HDR image with the *wxfalsecolor* or *pvalue* programs. The graphic user interface provided by *wxfalsecolor* is useful for measuring CR in HDR photographs, where small camera movements may shift the coordinates of the relevant pixels between measurements.

RESULTS

Artificially Lit Space

The Perception Lab was simulated under electrically lit conditions in both Radiance and Accelerad and compared to HDR photography of the space. Representative images appear in Figure 1.

Only one lighting state was modelled, but several simulation parameters were varied in order to determine their effect on time and accuracy. The number of ambient bounces (*-ab*) was varied in both Radiance and Accelerad. Additionally, the irradiance cache size (*-ac*) was varied in Accelerad. Radiance




Figure 1 Tone-mapped and false color images of the Perception Lab created by HDR photography (top), Radiance with 8 ambient bounces (middle), and Accelerad with 8 ambient bounces and 2048 cached ambient values (bottom).

sizes its irradiance cache dynamically, so the parameter was not applied to it. DGP results are not presented as lighting in the space did not cause glare.

All of the simulations underreported E_v . Figure 2 shows that greater numbers of ambient bounces and larger irradiance caches brought the result closer to the target set by the illuminance meter reading, which closely matched the value from the HDR photograph. However, there was little advantage to including more than five ambient bounces or 2048 ambient values.

Measurements from the HDR photograph and luminance meter disagree on the CR value. Figure 3 shows that after three ambient bounces, Radiance and most Accelerad simulations converge on a value for CR close to that of the HDR photograph. Accelerad simulations with small irradiance caches undersample ambient lighting on the monitor and thus overpredict CR.

Simulation times ranged from 43 to 169 seconds for Radiance and 30 to 197 seconds for Accelerad (Figure 4). Overly large or small irradiance caches caused particularly slow simulations; large caches took longer




Figure 2 Fraction of vertical eye illuminance achieved through simulation of the Perception Lab by varying ambient bounces and irradiance cache size relative to illuminance meter readings




Figure 3 CR for the monitor in the Perception Lab with varying ambient bounces and irradiance cache size

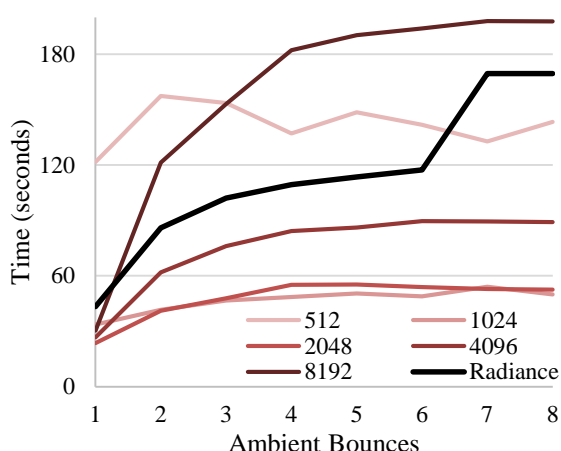


Figure 4 Simulation time for the Perception Lab with varying ambient bounces and irradiance cache size

to calculate, while small caches left portions of the scene unaccounted for, resulting in more work during Accelerad's final gather phase. Of note, an Accelerad simulation with 2048 cached ambient values achieved twice the speed of Radiance at five ambient bounces and 3.2 times the speed at eight ambient bounces while achieving accuracy comparable to Radiance for E_v and CR.

Daylit Space

The Media Lab's 5th floor lounge was simulated under 32 different natural lighting conditions in both Radiance and Accelerad and compared to HDR photography of the space. Figure 6 shows representative images under clear and overcast conditions.

Again, the simulations tended to underpredict E_v (Figure 5). On average, Accelerad simulation accounted for 75.5% of illuminance recorded by the meter ($\sigma = 17.3\%$), and Radiance simulations accounted for 74.0% of that illuminance ($\sigma = 16.9\%$). Accelerad's results had a smaller error than Radiance under 87.5% of observed sky conditions when compared to HDR photography and under 81.3% of observed sky conditions when compared to illuminance meter readings.

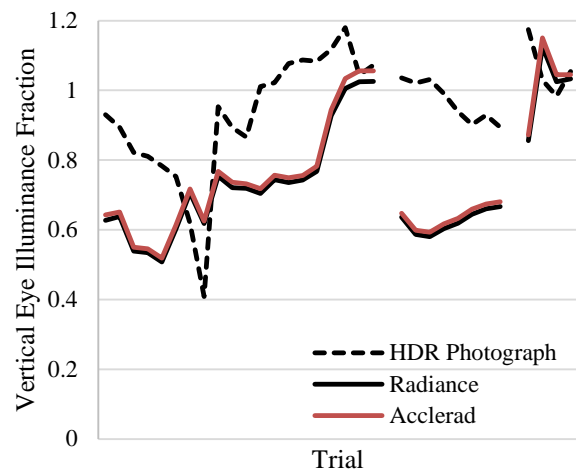


Figure 5 Fraction of vertical eye illuminance achieved through simulation and HDR photography of the lounge at each measurement time taken over three periods relative to illuminance meter readings

The two simulation tools also tended to underreport DGP compared to HDR photography. Figure 7 shows that DGP results were simulated most accurately at times when the sun did not directly illuminate the scene – toward the end of the first day of observations and on the third day, which was overcast. DGP predictions by Radiance and Accelerad were generally similar to each other and less than the observed values. However, during a one-hour period on the first day when the sun was directly in the field of view, they rise well above the DGP value recorded by HDR photography. The low DGP in photographs is the result of luminous overflow, where the camera's




Figure 6 Tone-mapped and false color images of the 5th floor lounge created by HDR photography (top), Radiance with 3 ambient bounces (middle), and Accelerad with 5 ambient bounces and 8192 cached ambient values (bottom). Clear sky conditions were recorded at 9:15 AM on 9 January 2014 (left), and overcast sky conditions were recorded at 9:15 AM on 10 January 2014 (right).

sensor is saturated by light even at its shortest exposure (Jakubiec, 2014). Of the 29 unaffected values, Accelerad tended to underpredict DGP by 0.101 ($\sigma = 0.065$), and Radiance underpredicted it by 0.106 ($\sigma = 0.066$).

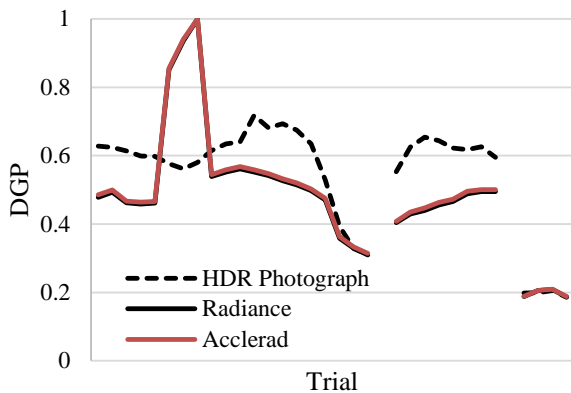


Figure 7 DGP for the lounge scene at each measurement time taken over three periods

There is close agreement on CR between the two simulation engines and the two measurement

techniques (Figure 8). The simulations again tended to predict higher CR values than were measured, but in this case, the CR values measured by the luminance meter were generally higher than those calculated from HDR photography were. Accelerad's results differed on average from the luminance meter by 0.14 ($\sigma = 3.85$) and from HDR photography by 0.93 ($\sigma = 1.50$). Radiance's results differed on average from the luminance meter by 0.01 ($\sigma = 3.78$) and from HDR photography by 1.09 ($\sigma = 1.52$). While these errors are small, it is worth noting that the actual CR was frequently very close to 4, the minimum required by the current ISO standard. For this particular scene, even small errors can result in incorrect assessments of lighting quality; Radiance and Accelerad correctly predicted whether the scene met the ISO standard only 69% and 66% of the time, respectively.

Simulation times ranged from 13.1 to 16.8 minutes for Accelerad and from 341 to 378 minutes for Radiance (Figure 9). The mean simulation time for Accelerad was 15.1 minutes ($\sigma = 1.1$ minutes), while for Radiance it was 361 minutes ($\sigma = 10.1$ minutes). On average, Accelerad performed each simulation 24 times faster than Radiance ($\sigma = 2.1$).




Figure 8 CR for the monitor in the lounge at each measurement time taken over three periods

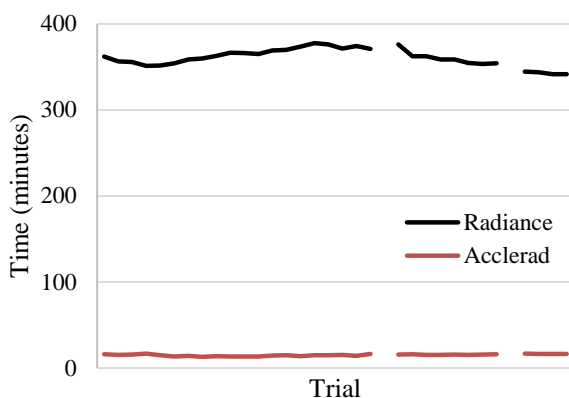


Figure 9 Simulation time for the lounge at each measurement time taken over three periods

CONCLUSION

The two global illumination simulation engines tested in this study, Radiance and Accelerad, have roughly equivalent accuracy for simulating naturally and artificially lit spaces. Both tools tended to underpredict the total luminance of the scenes studied, and as a result, tended to overpredict CR and underpredict E_v and DGP. While the magnitude of error for visual comfort metrics was generally within the accepted 20% bounds, the error is still sufficient to create problems in the use of visual comfort metrics.

The error in predicting CR was small enough that it may be disregarded for practical applications. However, in the lounge scene, the actual CR was typically very close to 4, the minimum acceptable values by ISO standards. The fact that such small numerical errors can lead to incorrect assessment of glare points to problems with the enforceability of current standards.

The error found in predicting DGP indicates a more serious problem with using this metric for the design of daylight spaces. Accelerad's underprediction by 0.101 and Radiance's underprediction by 0.106 are both within the typical range of error seen in validations of daylighting simulation tools. However, as the difference between imperceptible and intolerable glare is only 0.1 on this scale, it is questionable whether

today's best modelling practices can provide useful glare predictions in cases that are "on the edge" of glare conditions.

The largest error found in this study was a systematic underprediction of E_v by both simulation tools in comparison to illuminance meter readings and HDR photography. Such systematic errors are generally the result of discrepancies between the model and the real space. In this case, it could be the result of incorrect source data from weather and IES files or of measurement errors in material reflectance and transmittance data. Although E_v is not frequently used as a metric by building designers, further study of these issues is still warranted, as inaccurate calculation may produce errors in other metrics.

The studies of both spaces showed a clear advantage to performing simulation on the GPU. For both spaces, Accelerad settings were found that generated equivalent results to Radiance with notable speedups. The maximum speedup achieved in the Perception Lab was 3.2, while a speedup of 24 times was achieved in the lounge scene. Both the speedup factor and the settings required to achieve it are highly dependent on the scene in question. Larger scenes offer more potential for speedup on the GPU, but also require higher accuracy settings.

These speedups represent a first step toward fast visual comfort feedback during building design. With continued research and development, we hope that this feedback may someday be available to designers at interactive rates. Interactive feedback will require not only faster software tools, but also new generations of graphics hardware that will provide more compute cores and faster memory access than those available today. The benefit of this will be a new depth of information available to architects through all stages of design.

ACKNOWLEDGEMENT

This research was funded through the Kuwait-MIT Center for Natural Resources and the Environment by the Kuwait Foundation for the Advancement of Sciences. The Tesla K40 accelerators used for this research were donated by the NVIDIA Corporation. Matt Aldrich provided valuable assistance coordinating studies in the Media Lab, and Susanne Seitingner provided data on the lighting of the Perception Lab.

REFERENCES

- Anywhere Software, 2014. *Photosphere*.
- Galasiu, A. D. & Atif, M. R., 2002. Applicability of daylighting computer modeling in real case studies: Comparison between measured and simulated daylight availability and lighting consumption. *Building and Environment*, 37(4), pp. 363-377.
- Goral, C. M., Torrance, K. E., Greenberg, D. P. & Battaile, B., 1984. Modeling the interaction of

- light between diffuse surfaces. *Computer Graphics*, 18(3), pp. 213-222.
- Grynberg, A., 1989. *Validation of Radiance*, Berkeley, CA: Lawrence Berkeley Laboratories. Document ID 1575.
- Inanici, M., 2006. Evaluation of high dynamic range photography as a luminance data acquisition system. *Lighting Research and Technology*, 38(2), pp. 123-134.
- Inanici, M., 2010. Evaluation of high dynamic range image-based sky models in lighting simulation. *LEUKOS: The Journal of the Illuminating Engineering Society of North America*, 7(2), pp. 69-84.
- ISO 9241-303, 2008. *Ergonomics of human-system interaction – Part 303: Requirements for electronic visual displays*, Geneva: International Organization for Standardization.
- ISO 9241-3, 1992. *Ergonomic requirements for office work with visual display terminals (VDTs) – Part 3: Visual display requirements*, Geneva: International Organization for Standardization.
- Jakubiec, A. J. & Reinhart, C. F., 2013. Predicting visual comfort conditions in a large daylit space based on long-term occupant evaluations: A field study. *Proceedings of BS2013: 13th Conference of International Building Performance Simulation Association, Chambéry, France, August 26-28*, pp. 3408-3415.
- Jakubiec, J. A., 2014. *The use of visual comfort metrics in the design of daylit spaces*, PhD dissertation: Massachusetts Institute of Technology.
- Jakubiec, J. A. & Reinhart, C. F., 2012. The 'adaptive zone' - A concept for assessing discomfort glare throughout daylit spaces. *Lighting Research and Technology*, 44(2), pp. 149-170.
- Jones, N. L. & Reinhart, C. F., 2014a. Physically based global illumination calculation using graphics hardware. *Proceedings of eSim 2014: The Canadian Conference on Building Simulation*, pp. 474-487.
- Jones, N. L. & Reinhart, C. F., 2014b. Irradiance caching for global illumination calculation on graphics hardware. *2014 ASHRAE/IBPSA-USA Building Simulation Conference, Atlanta, GA, September 10-12*, pp. 111-120.
- Larson, G. W. & Shakespeare, R., 1998. *Rendering with Radiance: The Art and Science of Lighting Visualization*. San Francisco: Morgan Kaufmann Publishers, Inc.
- Mardaljevic, J., 1995. Validation of a lighting simulation program under real sky conditions. *Lighting Research and Technology*, 27(4), pp. 181-188.
- Mardaljevic, J., 2001. The BRE-IDMP dataset: a new benchmark for the validation of illuminance prediction techniques. *Lighting Research and Technology*, 33(2), pp. 117-136.
- Moghbel, N., 2012. *New model for VDT associated visual comfort in office spaces*, PhD dissertation: Karlsruhe Institute of Technology.
- Ng, E. Y.-Y., Poh, L. K., Wei, W. & Nagakura, T., 2001. Advanced lighting simulation in architectural design in the tropics. *Automation in Construction*, 10(3), pp. 365-379.
- Parker, S. G. et al., 2010. OptiX: A general purpose ray tracing engine. *ACM Transactions on Graphics – Proceedings of ACM SIGGRAPH 2010*, 29(4).
- Perez, R., Seals, R. & Michalsky, J., 1993. All-weather model for sky luminance distribution—Preliminary configuration and validation. *Solar Energy*, 50(3), pp. 235-245.
- Reinhart, C. & Breton, P.-F., 2009. Experimental validation of Autodesk® 3ds Max® Design 2009 and Daysim 3.0. *LEUKOS: The Journal of the Illuminating Engineering Society of North America*, 6(1), pp. 7-35.
- Reinhart, C. F. & Andersen, M., 2006. Development and validation of a radiance model for a translucent panel. *Energy and Buildings*, 38(7), pp. 890-904.
- Reinhart, C. F. & Herkel, S., 2000. The simulation of annual daylight illuminance distributions - a state-of-the-art comparison of six RADIANCE-based methods. *Energy and Buildings*, 32(2), pp. 167-187.
- Reinhart, C. F. & Walkenhorst, O., 2001. Validation of dynamic RADIANCE-based daylight simulations for a test office with external blinds. *Energy and Buildings*, 33(7), pp. 683-697.
- Van Den Wymelenberg, K. & Inanici, M., 2014. A critical investigation of common lighting design metrics for predicting human visual comfort in offices with daylight. *LEUKOS: The Journal of the Illuminating Engineering Society of North America*, 10(3), pp. 145-164.
- Van Den Wymelenberg, K., Inanici, M. & Johnson, P., 2010. The effect of luminance distribution patterns on occupant preference in a daylit office environment. *LEUKOS: The Journal of the Illuminating Engineering Society of North America*, 7(2), pp. 103-122.
- Wienold, J. & Christoffersen, J., 2006. Evaluation methods and development of a new glare prediction model for daylight environments with the use of CCD cameras. *Energy and Buildings*, 38(7), p. 743-757.