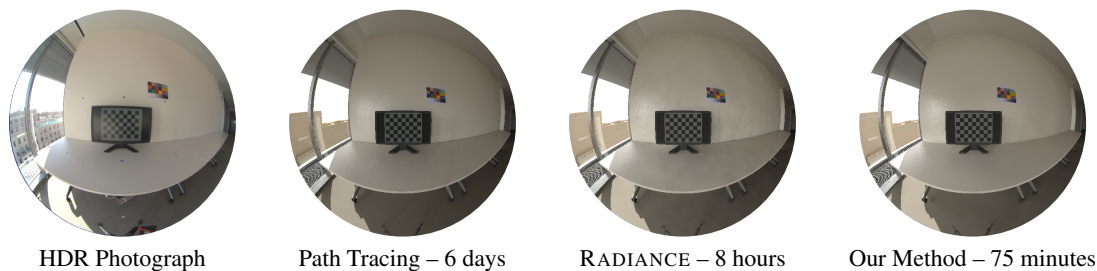


# Parallel Multiple-Bounce Irradiance Caching

Nathaniel L. Jones and Christoph F. Reinhart

Massachusetts Institute of Technology, USA



**Figure 1:** We use parallel irradiance caching to predict the luminance distribution in a daylit space of moderate architectural complexity and compare the result to calibrated high dynamic range photographs.

## Abstract

Building designers rely on predictive rendering techniques to design naturally and artificially lit environments. However, despite decades of work on the correctness of global illumination rendering techniques, our ability to accurately predict light levels in buildings—and to do so in a short time frame as part of an iterative design process—remains limited. In this paper, we present a novel approach to parallelizing construction of an irradiance cache over multiple-bounce paths. Relevant points for irradiance calculation based on one or multiple cameras are located by tracing rays through multiple-bounce paths. Irradiance values are then saved to a cache in reverse bounce order so that the irradiance calculation at each bounce samples from previously calculated values. We show by comparison to high-dynamic range photography of a moderately complex space that our method can predict luminance distribution as accurately as RADIANCE, the most widely validated tool used today for architectural predictive rendering of daylit spaces, and that it is faster by an order of magnitude.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Rendering, Global Illumination

## 1. Introduction

Architects and lighting designers use software tools to predict light levels to achieve qualitative design goals and to meet quantitative illumination requirements. Increasingly, illumination goals focus on the aesthetic [RA15] and functional role of daylight [IES12], which is highly time-dependent and in which indirect lighting plays a major role [JR14]. To this end, the design community depends on *predictive rendering*, image synthesis whose goal is not to look plausible but rather to verifiably match the physical scene once built.

Unfortunately, predictive rendering tools currently available to the design community fall short in terms of speed, accuracy, or both. The RADIANCE synthetic imaging software [War94] has become a staple of the architectural and lighting design communities due to its widespread adoption and well-validated results [RF06].

However, RADIANCE does not offer parallelism in many use cases, and diffuse lighting calculations computed by irradiance caching account for the majority of time spent in long simulations. While other rendering techniques produce plausible results, validation studies that test physical accuracy in realistically complex scenes are rare, and adoption by the architectural community has been slow. Most test scenes used in computer graphics validation do not include glazed windows, and many modern physically-based rendering algorithms produce lower-quality results when applied to naturally-lit urban indoor scenes. Our main contributions include:

- A scalable GPU-based method for calculating diffuse lighting using a multiple-bounce irradiance cache similar to RADIANCE.
- Validation by comparison to physical measurements that our method predicts global illumination of a daylit indoor scene as accurately as RADIANCE.

## 2. Related Work

The goals of predictive rendering vary substantially from one field of application to the next. Virtual prototyping, used for instance in the automotive industry, is concerned with accurate surface reflection properties [UWP06]. Visual psychophysics uses rendered images to assess human perception of color and shading [RB06]. Architects and lighting designers are concerned with producing the appropriate combination of artificial and natural lighting to provide a desired appearance for a space while maintaining comfortable and task-appropriate illumination and contrast levels. Increasingly, reducing energy use by increasing access to natural light is also a goal [Rei14]. In this paper, we focus on the field of *building performance simulation*, which includes the quantitative goals of architects and lighting designers.

Building performance simulation puts a unique set of demands on image synthesis tools. Design tools need to produce physically accurate representations and be reasonably fast so as not to interrupt the flow of the user's thought process [Csi96]. Unfortunately, verifying the accuracy of predictive rendering tools is a messy business, particularly in scenes of typical architectural complexity and under variable daylight conditions. Current lighting guidelines expect simulated luminance to fall within 20% of actual levels [Rea00], a goal that remains difficult to achieve. To meet design goals, focus is placed on producing correct illuminance levels over broad surfaces and work planes. This differs from the characteristics that contribute to the perception of photorealism because the human eye is more sensitive to relative luminance and higher spatial frequencies than are typically found in buildings [RWP\*95]. Animation is rarely a concern; objects in the scene are static, while the sun position and sky condition may change [RH00]. The accuracy of rendering glass is important because glazed windows both modulate incoming light and provide views to external geometry in the scene's urban context. However, caustic focusing [Nye99] is rarely encountered because architectural glass tends to have planar geometry. As a result of windows, image synthesis tools must be able to work with a variety of scales, including the detailed geometry of a room's interior and the far-away geometry of the urban surroundings which may participate in the diffuse lighting of the scene.

### 2.1. Validation Studies

A renderer produces photorealistic results if it meets three criteria: correct reflection models for surfaces, correct light transport paths, and physiologically correct display of the results [GTS\*97]. We are concerned mainly with the first two criteria, as the goal in building performance simulation is to obtain physically accurate light levels, not to display perceptually convincing images. Validation of the first two criteria is generally done by comparison to physically measured photometric values. The physical and virtual models must have matching geometry, materials, and light sources; however, some inaccuracy is inherent in any model.

Early validation studies compared rendered images to photographs [GTGB84, MRC\*86, Gry89]. McNamara et al. [MCTG00] used human subjects to compare RADIANCE and other rendering methods to a photograph and found that while RADIANCE performed well, human subjects were not suited to detecting certain physical inaccuracies. A similar comparison by Khodulev and

Kopylov [KK96] found RADIANCE to perform best in an analytic test case but lacked a ground truth for image comparison. Drago and Myszkowski [DM01] used highly accurate material models to approximate a photograph on a low-dynamic range display, but they simplified the luminance distribution in the scene by taking photographs at night.

Most validation studies rely on point illuminance measurements. When taken in real buildings, reference measurements tend to be heavily influenced by daylight, and thus accurate modeling of the sky is a key concern [PSM93]. RADIANCE produced the most accurate results in a comparison of four validated rendering packages to field measurements in one daylit building, although still off by up to 40% [UH98]. Other studies comparing physical daylit building interiors with glazed windows to RADIANCE predictions found relative mean bias error ( $MBE_{rel}$ ) under 20% and relative root mean square error ( $RMSE_{rel}$ ) under 32% [NPWN01, RW01, RA06]. RADIANCE and 3ds Max [Aut16] gave comparable results in a study of one building interior under a number of sky conditions [RB09]. However, RADIANCE offered superior results in a repetition of this study at four other geographic locations [BPF15].

High dynamic range (HDR) photography can also be used to capture accurately scaled luminance levels and provides greater spatial resolution than typical illuminance sensors [Ina06, Ina10, VIJ10, VI14]. Karner and Prantl [KP96] compared photographs of an artificially lit office to RADIANCE using gridded regions with  $MBE_{rel}$  44%–71% and  $RMSE_{rel}$  16.4%–18.5%. When corrected for image misalignment, the errors were lowered to  $MBE_{rel}$  21%–52% and  $RMSE_{rel}$  12.9%–17.8%. In cases where photographs and renderings have poor alignment, average luminance values in manually selected task-zones can also be compared [WC06].

### 2.2. Irradiance Caching

Diffuse lighting due to indirect irradiance typically has low-frequency variation, so a single irradiance calculation may be applied to all ray intersections with similar spatial coordinates and intersected surface normal directions. The *irradiance cache* is a collection of irradiance values that are saved for reuse, along with their associated point locations, normal directions, and validity radii. Ward's original method implemented in RADIANCE [WRC88] has been modified by the inclusion of gradients in the calculation of the validity radius [WH92] and by the use of second-order gradients (Hessians) to specify an ellipse in which an irradiance value is valid [JSKJ12, SJJ12].

RADIANCE takes a lazy approach to populating the irradiance cache [WRC88, WH92, SJJ12]. It calculates and saves an irradiance value only when no pre-existing value is found for a ray intersection. Those intersections may occur anywhere up to a user-specified number of ray bounces, so we refer to this as a *multiple-bounce irradiance cache*. Due to the lazy approach, irradiance values calculated at deeper bounces may sample from previously calculated irradiance values reached through fewer bounces, allowing them in effect to sample diffuse lighting from a greater portion of the scene. As a result, irradiance values calculated through a multiple-bounce irradiance cache converge toward an infinite bounce solution more quickly than those calculated from a single-bounce irradiance cache.

**Algorithm 1** Parallel multiple bounce irradiance caching algorithm

---

```

1: procedure BUILD IRRADIANCE CACHE
2:    $level = 0$ 
3:    $points[0] = eye$ 
4:    $cache[level_{max}] = \emptyset$ 
5:   while  $level < level_{max}$  do
6:      $temp = \text{sample geometry seen from } points[level]$  ▷ parallel coarse geometry sampling
7:     Sort  $temp$  by cell index
8:     Increment  $level$ 
9:      $points[level] = r_{coarse}$ -spaced points from  $temp$ 
10:  end while
11:  while  $level > 0$  do
12:    Decrement  $level$ 
13:     $cache[level] = \text{sample irradiance from } cache[level + 1] \text{ at } points[level + 1]$  ▷ parallel coarse irradiance sampling
14:     $temp = \text{sample geometry seen from } points[level]$  ▷ parallel fine geometry sampling
15:    Sort  $temp$  by cell index
16:     $points[level + 1] = r_{min}$ -spaced points from  $temp$ 
17:     $cache[level] += \text{sample irradiance from } cache[level + 1] \text{ at } points[level + 1]$  ▷ parallel fine irradiance sampling
18:  end while
19: end procedure

```

---

### 2.3. Parallel Irradiance Caching

Ward's lazy approach is well-suited to serial implementation. However, RADIANCE allows multiple processes to write to a single cache using file locks [LS03]. Improved synchronization methods use the Message-Passing Interface (MPI) [KMG99, DSC06] and wait-free synchronization [DDSC09]. Concurrent threads may calculate overlapping and redundant irradiance values, but this happens infrequently so long as the number of threads is small.

In massively parallel systems, however, a lazy approach to populating the irradiance cache is likely to produce many redundant entries. The cache can be prefilled with entries that are likely to be sampled; however, this requires a heuristic approach to determine where irradiance will need to be calculated [PH04]. Various strategies have been proposed to completely prefill the irradiance cache. Splatting [KG09] and pre-convolution [RZD14] provide view-dependent solutions computed in screen space. Neighbor clamping [KBpv08], Poisson-disk distribution [LZT\*08], micro-rendering [REG\*09], and dithering combined with z-curve clustering [FVKG13] can be used to place calculation points in world space, but consider only one diffuse bounce. An adaptive seeding method by Wang et al. [WWZ\*09] uses quadtrees and  $k$ -means clustering to choose irradiance calculation locations. It also considers only a single bounce within the irradiance cache, but the irradiance values themselves come from photon mapping [ZHWG08].

The adaptive seeding algorithm from Wang et al. [WWZ\*09] has been extended to cache irradiance at multiple bounces [JR14]. Rays are traced outward from a subset of the original point-normal pairs to locate new point-normal pairs for a second bounce irradiance cache. Iteration of this process produces locations for irradiance calculation for the desired number of bounces, with the irradiance values themselves calculated in the reverse order. The user remains responsible for predicting the necessary size of the irradiance cache, which is impractical in most real cases and prohibits widespread adoption of the method. Unfortunately, poor cache siz-

ing lead either to under- or over-coverage of the scene, which in either case can severely increase rendering times. Our method is similar in approach but automatically creates an irradiance cache that gives full scene coverage with near-optimal size.

### 3. Method

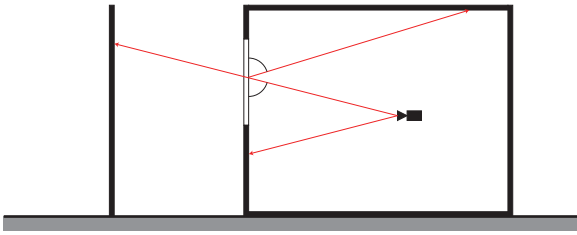
Our method allows any existing single-bounce or single-threaded irradiance caching algorithm to be used in a parallel multiple-bounce framework. The irradiance caching method must include an error metric that defines the region around the calculation point for which the irradiance value is valid, such as the split-sphere [WH92] or Hessian-based error metrics [JSKJ12]. For our implementation, we apply RADIANCE's irradiance calculation method, which uses Hessian-based error control [SJJ12]. In the lazy approach, the analytical error metric determines the spacing between irradiance calculation points, and in combination with ray traversal order, it determines the position of each irradiance calculation location. For parallel computation, rays at the same bounce depth are effectively traversed simultaneously, so we must introduce an alternate placement method to pick irradiance calculation points without having previously calculated their neighbors.

Our method breaks the irradiance caching algorithm into two phases: first, a geometry sampling phase identifies point-normal pairs that require irradiance calculation, and second, an irradiance sampling phase collects indirect lighting contributions to each point. These two phases repeat for both coarse and fine spacing at each level of bounce recursion. Cell-based greedy selection applied to the output of the coarse geometry sampling phase prevents the number of irradiance samples from growing exponentially with the number of bounces. We calculate irradiance at each selected point along with an irradiance Hessian defining its validity area. Gaps may occur between the coarsely spaced ellipses, so we use the fine sampling phases to locate and compute additional irradiance values for the cache. We show pseudocode for our method in Algorithm 1.

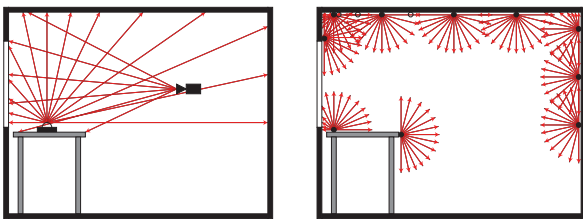
### 3.1. Coarse Geometry Sampling

The first step is to find point-normal pairs at which irradiance will later be calculated. Unlike the lazy approach, we identify point-normal pairs in advance rather than on an as-needed basis. As a result, the validity radius of each irradiance value and thus the appropriate spacing between point-normal pairs is not known. We assume a user-defined lower limit on validity radii  $r_{min}$ ; in practice, RADIANCE users choose this limit based on rules of thumb. We could space calculation points by a distance of  $r_{min}$  in order to guarantee complete coverage, but this would create an excessively large irradiance cache. Instead, we use an intermediate spacing  $r_{coarse}$  equal to the geometric mean of  $r_{min}$  and  $r_{max}$ . This results in a coarse sampling of the scene, which will be refined later.

Geometry sampling begins with the creation of a candidate list of all potential point-normal pairs. The candidate list contains the first-hit points found by ray tracing, which can be performed in parallel. A tree of rays originate from the position of the virtual camera or sensor (first bounce) or from the locations of previously identified point-normal pairs (subsequent bounces) and extend to the first diffuse surface (Figure 2). This scheme allows us to simulate multiple cameras or sensors simultaneously with minimal overhead (Figure 3). Each ray returns a payload containing a hit point and associated normal, along with an index representing the cell containing the point-normal pair.



**Figure 2:** A tree of geometry sampling rays branch out until intersecting diffuse surfaces.



**Figure 3:** In the first round of geometry sampling, rays originate from any cameras or sensors in the scene (left). On subsequent bounces, rays originate from selected hit points of the previous round (right).

Cells divide the scene both spatially and with respect to normal direction. The dimensions of each cell are larger than  $r_{coarse}$ ; in practice, a dimension about thirty to sixty times  $r_{coarse}$  works well. Each cell is also associated with one of six axis-aligned directions,

such that a point-normal pair will be assigned to the cell whose dominant direction lies closest to its own normal direction. Intuitively, irradiance values calculated at point-normal pairs located near each other on the same surface are likely to have overlapping validity ellipses, and if geometry sampling is dense enough, some ellipses will be completely overlapped and therefore unnecessary. To remove point-normal pairs that are likely to be redundant, we sort and group the candidate pairs by cell.

With the candidate point-normal pairs sorted by cell, we can pick a near-optimal subset of these pairs at which irradiance will be calculated. We scan the list of candidate point-normal pairs and use a greedy approach within each cell to choose pairs for irradiance calculation. We automatically accept the first point-normal pair in each cell. We accept additional point-normal pairs if they are located at least  $r_{coarse}$  away from all previously accepted pairs in the same cell or if their normal direction differs by more than a user-specified deviation, defaulting to 0.2 radians. While this directional deviation test is not optimal, it has produced good results previously [SJJ12] and does not cause misses because we use it at other times that we search the irradiance cache as well. We save the selected point-normal pairs as origins for later irradiance calculation and also more immediately as origins for geometry sampling at the next bounce.

### 3.2. Coarse Irradiance Sampling

Once we complete coarse geometry sampling for all bounces, we begin irradiance calculation starting with the deepest bounce. Any irradiance calculation method may be used; our implementation uses jittered Shirley-Chiu radiance sampling [SC97] with Hessian-based validity ellipses as implemented in RADIANCE [SJJ12]. We trace irradiance sampling rays only to their first hit point, where we sample either direct irradiation only (at the deepest bounce) or direct irradiation and cached irradiance values from the next deeper bounce (at all other bounces).

Hessian-based irradiance calculation determines an elliptical area for which the irradiance value holds with major and minor radii  $R_i^{\lambda_1}$  and  $R_i^{\lambda_2}$  defined as follows [JSKJ12]:

$$\left( R_i^{\lambda_1}, R_i^{\lambda_2} \right) \approx \sqrt[4]{\frac{4\epsilon^t}{\pi}} \left( \sqrt[4]{\frac{1}{\lambda_1}}, \sqrt[4]{\frac{1}{\lambda_2}} \right) \quad (1)$$

where  $\lambda_1$  and  $\lambda_2$  are the eigenvalues of the irradiance Hessian matrix and  $\epsilon^t$  is the user-defined total allowable error. We bound the radii between  $r_{min}$  and  $r_{max}$  to prevent under- or over-sampling, respectively. If the minor radius of the ellipse is less than  $r_{coarse}$ , then some relevant portions of the scene may not be covered by any cached irradiance value, and irradiance sampling at the next bounce closer to the eye could encounter holes. To prevent this, we insert a fine sampling phase before proceeding to the next bounce.

### 3.3. Fine Geometry Sampling

The fine geometry sampling phase proceeds much like the coarse version with a few exceptions. First, during the parallel ray casting, we ignore first hit points if they fall within the validity ellipse of an irradiance value calculated during coarse irradiance sampling at the

same bounce. We can still use these points as locations for specular reflections, however. Second, we use a smaller spacing to pick point-normal pairs within each cell. In principle, the spacing should be  $r_{min}$  to guarantee complete coverage of the scene; however, in many cases, we can achieve complete coverage with a larger spacing, as the validity radius will only reach its lower limit near edges under low light levels. Finally, because of the smaller spacing, we also use a smaller cell size. Typically, a cell edge length thirty to sixty times  $r_{min}$  works well.

### 3.4. Fine Irradiance Sampling

Irradiance calculation for fine sampling proceeds the same as for coarse irradiance sampling with validity areas defined by Equation 1. The origin points are those identified by the fine geometry sampling phase. We skip this phase if the fine geometry sampling phase identified no new points, which may happen if the irradiance values calculated in the coarse phase all have minor elliptical radii greater than  $r_{coarse}$ .

We add new irradiance values found in this phase to the cache created previously in the coarse irradiance sampling phase. Only the coarse and fine irradiance calculation at the next bounce closer to the eye will then sample from this irradiance cache, or, for the bounce closest to the eye, this irradiance cache will be used for the synthesis of the final images or sensor readings.

Because  $r_{min}$  is the lower limit on the radius of a validity ellipse, we expect the irradiance cache to completely cover relevant areas of the scene after the fine irradiance sampling phase. In practice, a larger spacing is sometimes preferable in the fine geometry sampling phase to prevent over-sampling. We can repeat geometry and irradiance sampling phases again with even smaller spacing to ensure complete scene coverage in this case. Alternately, we can perform a final gather in the event that a pixel is not covered by any cached value when a larger spacing is chosen. In our experience, misses occur in these cases at between zero and 0.04% of pixels using well-chosen parameters, typically at edges of very narrow surfaces.

## 4. Validation

We tested our method and several others by modeling a daylit interior room of a campus building. The room is of typical complexity for architectural models, and the space was mocked-up to resemble a typical office environment. We measured diffuse and specular reflectance for all materials with a Konica Minolta CM-2500d spectrophotometer, and we determined the transmissivity of the glazed surfaces with a Konica Minolta TL-1 illuminance meter. Unfortunately, no equipment was available to measure bidirectional reflectance *in situ*, so Lambertian reflectance modified by the measured specular component was generally assumed. We estimated material roughness to match appearance in our renderings, as is typical practice in RADIANCE.

We observed the scene under a variety of sky conditions over several days. For each observation, we captured HDR imagery from a typical seated head height in front of a monitor using a Canon EOS 5D Mark II camera with a Sigma 4.5mm fisheye lens. We

also took spot luminance measurements with a Konica Minolta LS-110 luminance meter for image calibration. To correct vignetting, we measured light fall-off as an angular function for the camera and applied the inverse function to the images to remove lens effects [Ina10]. The resulting HDR photographs capture the luminance distribution of a hemisphere approximating the field of view of a seated individual.

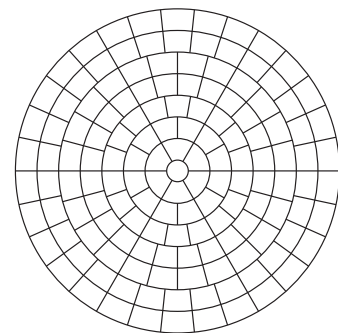
### 4.1. Illumination Sources

The largest source of uncertainty in our validation is the luminance distribution from the sun and sky. Solar irradiance data was measured at the time of each photograph with an Onset S-LIB-M003 silicon pyranometer connected to a HOBO weather station located atop a tall building approximately 200 m from the test scene. The station gives a single value for global horizontal solar irradiance, which is separated into direct and diffuse components using the Reindl separation model [RBD90]. These values were used to generate sun and sky definitions with the Perez all-weather sky model [PSM93]. These definitions are directly readable by RADIANCE and by our implementation. While the limitations of separation models are well known, their use in building performance simulation is fairly standard due to the expense of equipment that can directly measure direct sunlight.

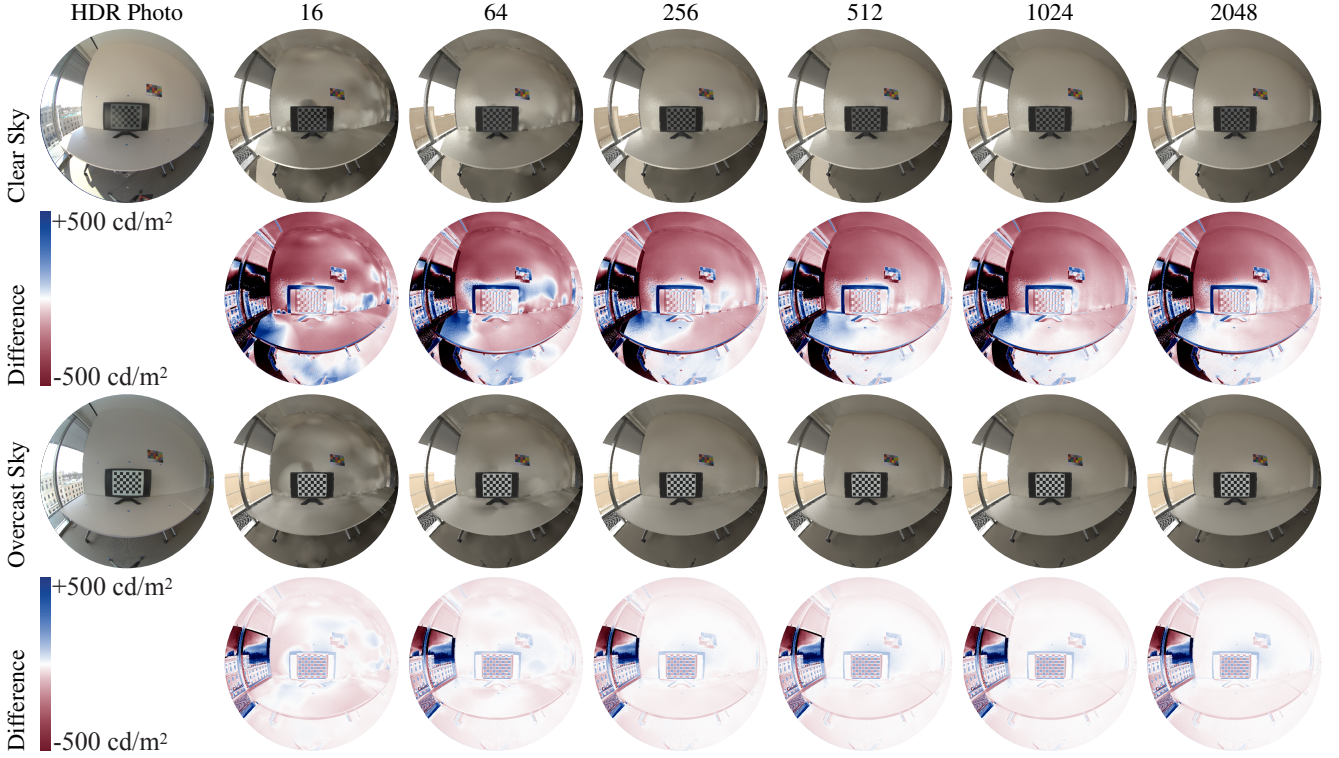
We recorded 32 sets of measurements over several days under a variety of sky conditions. As might be expected, the agreement between the simulated and measured images is imperfect, but  $MBE_{rel}$  was within the 20% margin specified by the Illuminating Engineering Society of North America [Rea00]. Based on our initial study [JR15], we chose two measurement times that generated the closest agreement in vertical eye illuminance between the HDR photographs, RADIANCE simulations, and illuminance meter readings. These were 14:00 on January 8, 2014, which was a clear day, and 10:00 on January 10, 2014, which was overcast. Since these measurements agree within 8%, we have higher confidence that the Perez model accurately depicted the real sky at these times.

### 4.2. Comparison Metrics

A simple and obvious comparison method would be pixel-by-pixel comparison of the HDR photograph to the rendered image. How-



**Figure 4:** Stereographic projection of Tregenza's division of a hemisphere into 145 patches.



**Figure 5:** The scene captured by HDR photography (left side) and rendered by our method with varying numbers of radiance samples per irradiance calculation seen in 180° fisheye projection.

ever, such a comparison is not terribly useful and will tend to highlight small misalignments between objects in the physical and modeled scenes. Furthermore, this comparison is difficult across images with differing resolutions or projections. As building performance simulation is more concerned with light levels across a work area, an alternative approach is to compare the averaged luminance values of manually selected broad regions such as task-zones. In order to produce a more complete comparison, we use an automated zone division process. We choose the Tregenza subdivision, which divides the hemisphere into 145 patches of approximately equal solid angle (see Figure 4) [TW83]. The patches are large enough that minor misalignments between real and modeled geometry are unlikely to significantly affect comparisons, yet they still localize areas of extreme brightness and so are likely to detect differences in sources of glare [WC06]. Regardless of the image resolution or projection used, the luminance  $L_{t,i}$  of Tregenza patch  $i$  can be found by summing the solid angle-weighted pixel luminance values in that region:

$$L_{t,i} = \frac{\sum_{p \in t,i} L_p \omega_p}{\sum_{p \in t,i} \omega_p} \quad (2)$$

where  $L_p$  is the luminance of pixel  $p$  and  $\omega_p$  is the solid angle occupied by that pixel. From this, we find the relative error in the luminance computed for each patch using the HDR photograph as

a reference.

$$\eta_{t,i} = \frac{L_{t,i,HDR} - L_{t,i,simulation}}{L_{t,i,HDR}} \quad (3)$$

We then compute  $MBE_{rel}$  and  $RMSE_{rel}$  for each image based on  $n$  fully-visible Tregenza patches.

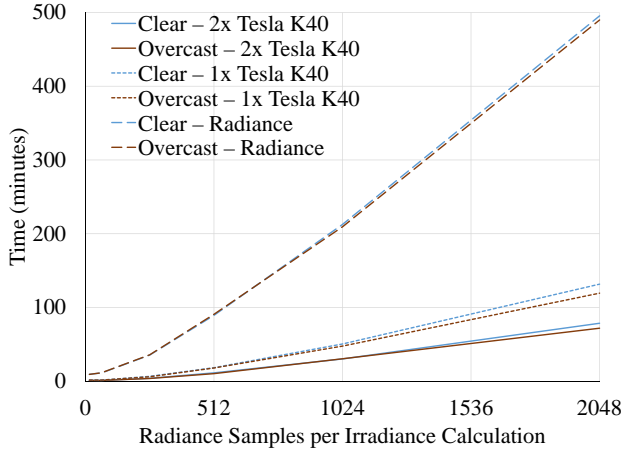
$$MBE_{rel} = \left| \frac{1}{n} \sum_{i=1}^n \eta_{t,i} \right| \quad (4)$$

$$RMSE_{rel} = \sqrt{\frac{1}{n} \sum_{i=1}^n \eta_{t,i}^2} \quad (5)$$

We also observe the fraction of patches that fall within an error threshold as a function of that threshold. More accurate simulations will yield curves that rise more steeply initially. This method of comparison is particularly useful for determining how well a simulation complies with modeling accuracy requirements.

## 5. Results and Discussion

We implemented our method on top of RADIANCE, using OptiX™ version 3.9.0 for parallel ray tracing [PBD\*10]. We ran simulations of the scene described in the previous section under observed clear and overcast sky conditions using our implementation as well as classic RADIANCE and naïve path tracing in OptiX™ (Figure 1). We allowed five diffuse bounces in our method and in RA-



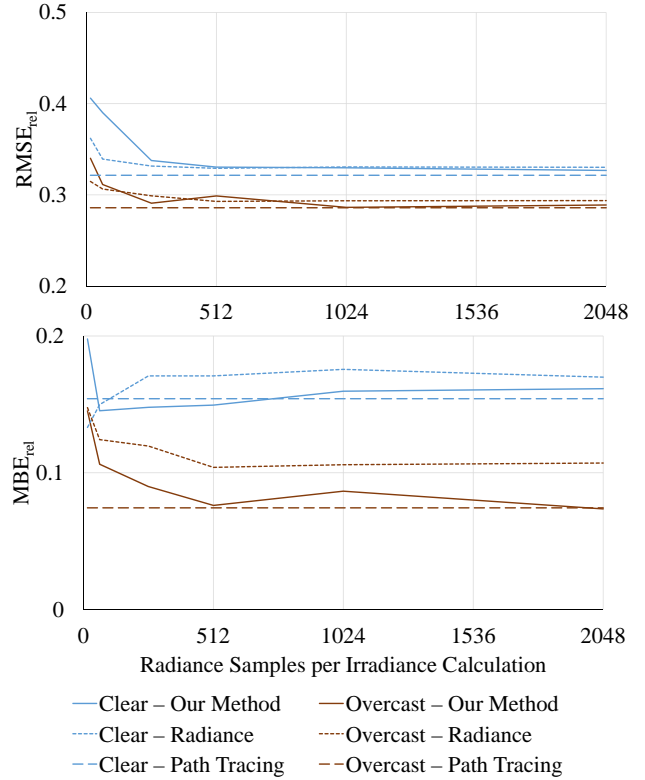
**Figure 6:** Simulation time as a function of sampling rate. Our method ran on one or two NVIDIA® Tesla® K40 GPUs. RADIANCE ran on one 3.4 GHz core.

DIANCE as recommended by Jones and Reinhart [JR14]. Based on the scene’s scale,  $r_{min}$  was approximately 0.05 m, and  $r_{max}$  was approximately 3 m.

The number of radiance samples used to calculate each irradiance value can be varied to affect both simulation time and image quality. Our renderings use an angular fisheye projection, which closely matches that of the Sigma 4.5mm fisheye lens used in the reference HDR photographs and gives complete visibility to all 145 Tregenza patches (Figure 5). Simulation time varies roughly linearly with the number of samples used to generate each irradiance value and scales with the number of GPU cores (Figure 6). Scaling with the number of GPU cores is not linear, mainly because OptiX™ creates the computation kernel for each GPU in serial.

Low numbers of radiance samples per irradiance calculation result in noticeably blotchy images, but Monte Carlo sampling results in surprisingly little variation in patch-level error; except for the lowest quality image, all the renderings detect roughly equal fractions of the total luminance from the HDR photographs. In general, our method produces less error than RADIANCE both in comparison to the photographed scene and to path tracing (Figure 7). We used an overture pass to smooth irradiance caching artifacts in the RADIANCE images, which had a negligible effect on patch-level errors. The path traced images were allowed to render until diffuse noise artifacts disappeared, although the variation in patch-level error over the duration of the run was minimal. Error in the path traced images is attributable to discrepancies between actual and modeled geometry, materials, and sky luminance distributions.

Higher numbers of radiance samples per irradiance calculation result do increase the number of Tregenza patches within an error threshold (Figure 8). For both sky conditions with any given error tolerance, simulations that use more radiance samples produce more patches within the tolerance. The advantage is especially apparent for the clear sky simulation at very low tolerances. However, the additional samples yield diminishing returns, and it seems un-



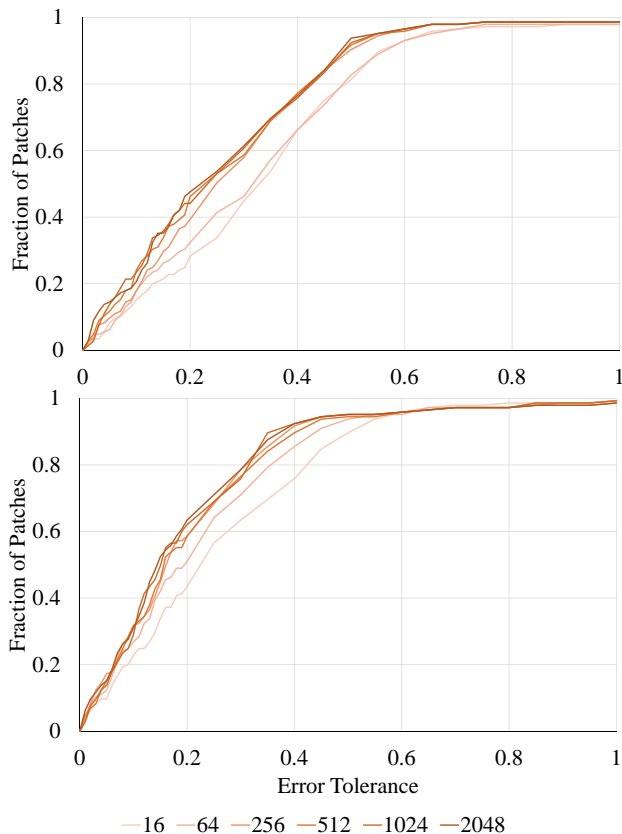
**Figure 7:**  $MBE_{rel}$  and  $RMSE_{rel}$  in luminance relative to measured values as a function of sampling rate.

likely that increasing beyond 2048 samples per irradiance calculation will yield much benefit.

## 6. Conclusion

We have presented a method for parallel construction of a multiple-bounce irradiance cache. Our framework can be adapted to a variety of irradiance sampling strategies, and through the use of coarse and fine passes, it can achieve a close to optimally sized irradiance cache without the individual validity radii being known in advance. Under models of real sky conditions, our method predicts the luminance distribution in a physical space as well as RADIANCE.

Additionally, we have introduced validation metrics appropriate for comparing simulations of complex spaces to ground truth measurements. We anticipate that further validation studies will be necessary to demonstrate the predictive capabilities of our method, and that validation of other global illumination simulation methods may also benefit from these metrics. Because of the complexity involved in modeling daylight spaces, and also because of their societal importance, validation should make use of a greater variety of physical locations and more varied sky conditions. Elimination of certain error sources, for example by capturing exact sky luminance distributions through photographic means instead of relying on recorded weather data, will improve the reliability of future studies. We stress the importance of validation by comparison



**Figure 8:** Fraction of Tregenza patches within a given error tolerance for clear sky (top) and overcast sky (bottom) rendered with our method depending on sampling rate.

to realistically complex scenes as opposed to analytically solvable simple cases.

The primary goal of our work is to give building designers the ability to evaluate their designs' performance in early design stages. The scalability of our method allows it to generate useful results an order of magnitude faster than RADIANCE using commodity graphics hardware. At early design stages, building plans change quickly, and designers may wish to compare many options in a short amount of time. The scalability of a parallel computation method for accurate global illumination may one day lead to interactive feedback, which will greatly advance the cause of early design building performance simulation.

### Acknowledgements

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. Frédo Durand provided indispensable technical and editorial guidance in the development of this paper. Matt Aldrich provided valuable assistance coordinating the physical validation study.

### References

- [Aut16] AUTODESK: 3ds Max. <http://www.autodesk.com/3dsMax>, 2016. 2
- [BPF15] BELLIA L., PEDACE A., FRAGLIASSO F.: The impact of the software's choice on dynamic daylight simulations' results: A comparison between Daysim and 3ds Max Design®. *Solar Energy* 122 (2015), 249–263. doi:10.1016/j.solener.2015.08.027. 2
- [Csi96] CSIKSZENTMIHALYI M.: *Creativity: Flow and the Psychology of Discovery and Invention*. Modern classics. HarperCollinsPublishers, 1996. 2
- [DDSC09] DUBLA P., DEBATTISTA K., SANTOS L. P., CHALMERS A.: Wait-free shared-memory irradiance cache. In *Eurographics Symposium on Parallel and Visualization* (2009), Eurographics, pp. 1–8. URL: <http://hdl.handle.net/1822/16441>. 3
- [DM01] DRAGO F., MYZKOWSKI K.: Validation proposal for global illumination and rendering techniques. *Computers & Graphics* 25, 3 (2001), 511–518. doi:10.1016/S0097-8493(01)00072-3. 2
- [DSC06] DEBATTISTA K., SANTOS L. P., CHALMERS A.: Accelerating the irradiance cache through parallel component-based rendering. In *Eurographics Symposium on Parallel and Visualization* (2006). URL: <http://hdl.handle.net/1822/5447>. 3
- [FVKG13] FROLOV V., VOSTRYAKOV K., KHARLAMOV A., GALAKTIONOV V.: Implementing irradiance cache in a GPU realistic renderer. In *Transactions on Computational Science XIX*, Gavrilova M., Tan C., Konushin A., (Eds.), vol. 7870 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013, pp. 17–32. doi:10.1007/978-3-642-39759-2\_2. 3
- [Gry89] GRYNBERG A.: Validation of Radiance. *Lawrence Berkeley Laboratory, Applied Sciences Division, Lighting Systems Research Group, Report # LBID 1575* (1989). 2
- [GTGB84] GORAL C. M., TORRANCE K. E., GREENBERG D. P., BATAILLE B.: Modeling the interaction of light between diffuse surfaces. *SIGGRAPH Comput. Graph.* 18, 3 (Jan. 1984), 213–222. doi:10.1145/964965.808601. 2
- [GTS\*97] GREENBERG D. P., TORRANCE K. E., SHIRLEY P., ARVO J., LAFORTUNE E., FERWERDA J. A., WALTER B., TRUMBORE B., PATTANAİK S., FOO S.-C.: A framework for realistic image synthesis. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1997), SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., pp. 477–494. doi:10.1145/258734.258914. 2
- [IES12] IESNA DAYLIGHTING METRICS COMMITTEE: *Lighting Measurement #83, Spatial Daylight Autonomy (sDA) and Annual Sunlight Exposure (ASE)*. IESNA Lighting Measurement, 2012. 1
- [Ina06] INANICI M.: Evaluation of high dynamic range photography as a luminance data acquisition system. *Lighting Research and Technology* 38, 2 (2006), 123–134. doi:10.1191/1365782806111640a. 2
- [Ina10] INANICI M.: Evaluation of high dynamic range image-based sky models in lighting simulation. *LEUKOS* 7, 2 (2010), 69–84. doi:10.1582/LEUKOS.2010.07.02001. 2, 5
- [JR14] JONES N. L., REINHART C. F.: Irradiance caching for global illumination calculation on graphics hardware. In *2014 ASHRAE/IBPSA-USA Building Simulation Conference, Atlanta, GA, September 10-12, 2014* (2014), American Society of Heating, Refrigerating, and Air-Conditioning Engineers, pp. 111–120. 1, 3, 7
- [JR15] JONES N. L., REINHART C. F.: Validation of GPU lighting simulation in naturally and artificially lit spaces. In *Proceedings of BS2015: 14th Conference of the International Building Performance Simulation Association* (2015), International Building Performance Simulation Association, pp. 1229–1236. 5
- [JSKJ12] JAROSZ W., SCHÖNEFELD V., KOBELT L., JENSEN H. W.: Theory, analysis and applications of 2d global illumination. *ACM Trans. Graph.* 31, 5 (Sept. 2012), 125:1–125:21. doi:10.1145/2231816.2231823. 2, 3, 4



- [KBPv08] KRÍVÁNEK J., BOUATOUCH K., PATTANAİK S., ŽÁRA J.: Making radiance and irradiance caching practical: Adaptive caching and neighbor clamping. In *ACM SIGGRAPH 2008 Classes* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 77:1–77:12. doi:10.1145/1401132.1401230. 3
- [KG09] KRÍVÁNEK J., GAUTRON P.: Practical global illumination with irradiance caching. *Synthesis lectures on computer graphics and animation 4*, 1 (2009), 1–148. 3
- [KH96] KHODULEV A., KOPYLOV E.: Physically accurate lighting simulation in computer graphics software. In *Sixth International Conference on Computer Graphics and Visualization* (1996). 2
- [KMG99] KOHOLKA R., MAYER H., GOLLER A.: MPI-parallelized radiance on SGI CoW and SMP. In *Parallel Computation*, Zinterhof P., Vajteršic M., Uhl A., (Eds.), vol. 1557 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1999, pp. 549–558. doi:10.1007/3-540-49164-3\_53. 3
- [KP96] KARNER K. F., PRANTL M.: A concept for evaluating the accuracy of computer generated images. In *Proceedings of the Twelfth Spring Conference on Computer Graphics (SCCG'96)* (1996). 2
- [LS03] LARSON G., SHAKESPEARE R.: *Rendering with Radiance: The Art and Science of Lighting Visualization*. Space & Light, 2003. 3
- [LZT\*08] LEHTINEN J., ZWICKER M., TURQUIN E., KONTKANEN J., DURAND F., SILLION F. X., AILA T.: A meshless hierarchical representation for light transport. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 37:1–37:9. doi:10.1145/1360612.1360636. 3
- [MCTG00] MCNAMARA A., CHALMERS A., TROSCIANKO T., GILCHRIST I.: Comparing real & synthetic scenes using human judgments of lightness. In *Rendering Techniques 2000*, Péroche B., Rushmeier H., (Eds.), Eurographics. Springer Vienna, 2000, pp. 207–218. doi:10.1007/978-3-7091-6303-0\_19. 2
- [MRC\*86] MEYER G. W., RUSHMEIER H. E., COHEN M. F., GREENBERG D. P., TORRANCE K. E.: An experimental evaluation of computer graphics imagery. *ACM Trans. Graph.* 5, 1 (Jan. 1986), 30–50. doi:10.1145/7529.7920. 2
- [NPWN01] NG E. Y.-Y., POH L. K., WEI W., NAGAKURA T.: Advanced lighting simulation in architectural design in the tropics. *Automation in Construction* 10, 3 (2001), 365 – 379. doi:10.1016/S0926-5805(00)00053-4. 2
- [Nye99] NYE J.: *Natural Focusing and Fine Structure of Light: Caustics and Wave Dislocations*. Taylor & Francis, 1999. 2
- [PBD\*10] PARKER S. G., BIGLER J., DIETRICH A., FRIEDRICH H., HOBEROCK J., LUEBKE D., MCALLISTER D., MCGUIRE M., MORLEY K., ROBISON A., STICH M.: OptiX: A general purpose ray tracing engine. In *ACM SIGGRAPH 2010 Papers* (New York, NY, USA, 2010), SIGGRAPH '10, ACM, pp. 66:1–66:13. doi:10.1145/1833349.1778803. 6
- [PH04] PHARR M., HUMPHREYS G.: *Physically Based Rendering: From Theory To Implementation*. Morgan Kaufmann series in interactive 3D technology. Elsevier Science, 2004. 3
- [PSM93] PEREZ R., SEALS R., MICHALSKY J.: All-weather model for sky luminance distribution – preliminary configuration and validation. *Solar Energy* 50, 3 (1993), 235 – 245. doi:10.1016/0038-092X(93)90017-I. 2, 5
- [RA06] REINHART C. F., ANDERSEN M.: Development and validation of a Radiance model for a translucent panel. *Energy and Buildings* 38, 7 (2006), 890 – 904. Special Issue on Daylighting Buildings. doi:10.1016/j.enbuild.2006.03.006. 2
- [RA15] ROCKCASTLE S., ANDERSEN M.: Human perceptions of daylight composition in architecture: A preliminary study to compare quantitative contrast measures with subjective user assessments in hdr renderings. In *Proceedings of BS2015: 14th Conference of the International Building Performance Simulation Association* (2015), International Building Performance Simulation Association, pp. 1205 – 1212. 1
- [RB06] RUPPERTSBERG A. I., BLOJ M.: Rendering complex scenes for psychophysics using RADIANCE: How accurate can you get? *J. Opt. Soc. Am. A* 23, 4 (Apr 2006), 759–768. doi:10.1364/JOSAA.23.000759. 2
- [RB09] REINHART C., BRETON P.-F.: Experimental validation of Autodesk® 3ds Max® Design 2009 and Daysim 3.0. *LEUKOS* 6, 1 (2009), 7–35. doi:10.1582/LEUKOS.2009.06.01001. 2
- [RBD90] REINDL D., BECKMAN W., DUFFIE J.: Diffuse fraction correlations. *Solar Energy* 45, 1 (1990), 1 – 7. doi:10.1016/0038-092X(90)90060-P. 5
- [Rea00] REA M. (Ed.): *The IESNA Lighting Handbook: Reference & Application*. Illuminating Engineering Society of North America, 2000. 2, 5
- [REG\*09] RITSCHEL T., ENGELHARDT T., GROSCHE T., SEIDEL H.-P., KAUTZ J., DACHSBACHER C.: Micro-rendering for scalable, parallel final gathering. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 132:1–132:8. doi:10.1145/1618452.1618478. 3
- [Rei14] REINHART C.: *Daylighting Handbook I*. 2014. URL: <http://www.daylightinghandbook.com/>. 2
- [RF06] REINHART C., FITZ A.: Findings from a survey on the current use of daylight simulations in building design. *Energy and Buildings* 38, 7 (2006), 824 – 835. Special Issue on Daylighting Buildings. doi:10.1016/j.enbuild.2006.03.012. 1
- [RH00] REINHART C. F., HERKEL S.: The simulation of annual daylight illuminance distributions – a state-of-the-art comparison of six RADIANCE-based methods. *Energy and Buildings* 32, 2 (2000), 167 – 187. doi:10.1016/S0378-7788(00)00042-6. 2
- [RW01] REINHART C. F., WALKENHORST O.: Validation of dynamic RADIANCE-based daylight simulations for a test office with external blinds. *Energy and Buildings* 33, 7 (2001), 683 – 697. doi:10.1016/S0378-7788(01)00058-5. 2
- [RWP\*95] RUSHMEIER H., WARD G., PIATKO C., SANDERS P., RUST B.: Comparing real and synthetic images: Some ideas about metrics. In *Rendering Techniques '95*, Hanrahan P. M., Purgathofer W., (Eds.), Eurographics. Springer Vienna, 1995, pp. 82–91. doi:10.1007/978-3-7091-9430-0\_9. 2
- [RZD14] REHFELD H., ZIRR T., DACHSBACHER C.: Clustered pre-convolved radiance caching. In *Proceedings of the 14th Eurographics Symposium on Parallel Graphics and Visualization* (Aire-la-Ville, Switzerland, Switzerland, 2014), PGV '14, Eurographics Association, pp. 25–32. doi:10.2312/pgv.20141081. 3
- [SC97] SHIRLEY P., CHIU K.: A low distortion map between disk and square. *J. Graph. Tools* 2, 3 (Dec. 1997), 45–52. doi:10.1080/10867651.1997.10487479. 4
- [SJJ12] SCHWARZHAUPT J., JENSEN H. W., JAROSZ W.: Practical hessian-based error control for irradiance caching. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31, 6 (Nov. 2012). doi:10.1145/2366145.2366212. 2, 3, 4
- [TW83] TREGENZA P., WATERS I.: Daylight coefficients. *Lighting Research and Technology* 15, 2 (1983), 65–71. doi:10.1177/096032718301500201. 6
- [UH98] UBBELOHDE M., HUMANN C.: *Comparative evaluation of four daylighting software programs [Lumen Micro, SuperLite, Radiance, and Lightscape]*. American Council for an Energy-Efficient Economy, Washington, DC (US), Jul 1998. 2
- [UWP06] ULBRICHT C., WILKIE A., PURGATHOFER W.: Verification of physically based rendering algorithms. *Computer Graphics Forum* 25, 2 (2006), 237–255. doi:10.1111/j.1467-8659.2006.00938.x. 2
- [VI14] VAN DEN WYMELENBERG K., INANICI M.: A critical investigation of common lighting design metrics for predicting human visual comfort in offices with daylight. *LEUKOS* 10, 3 (2014), 145–164. doi:10.1080/15502724.2014.881720. 2

- [VIJ10] VAN DEN WYMELENBERG K., INANICI M., JOHNSON P.: The effect of luminance distribution patterns on occupant preference in a daylit office environment. *LEUKOS* 7, 2 (2010), 103–122. doi:10.1582/LEUKOS.2010.07.02003. 2
- [War94] WARD G. J.: The RADIANCE lighting simulation and rendering system. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1994), SIGGRAPH '94, ACM, pp. 459–472. doi:10.1145/192161.192286. 1
- [WC06] WIENOLD J., CHRISTOFFERSEN J.: Evaluation methods and development of a new glare prediction model for daylight environments with the use of CCD cameras. *Energy and Buildings* 38, 7 (2006), 743–757. Special Issue on Daylighting Buildings. doi:10.1016/j.enbuild.2006.03.017. 2, 6
- [WH92] WARD G. J., HECKBERT P.: Irradiance gradients. In *Third Eurographics Workshop on Rendering* (1992), vol. 8598. 2, 3
- [WRC88] WARD G. J., RUBINSTEIN F. M., CLEAR R. D.: A ray tracing solution for diffuse interreflection. *SIGGRAPH Comput. Graph.* 22, 4 (June 1988), 85–92. doi:10.1145/378456.378490. 2
- [WWZ\*09] WANG R., WANG R., ZHOU K., PAN M., BAO H.: An efficient GPU-based approach for interactive global illumination. *ACM Trans. Graph.* 28, 3 (July 2009), 91:1–91:8. doi:10.1145/1531326.1531397. 3
- [ZHWG08] ZHOU K., HOU Q., WANG R., GUO B.: Real-time KD-tree construction on graphics hardware. In *ACM SIGGRAPH Asia 2008 Papers* (New York, NY, USA, 2008), SIGGRAPH Asia '08, ACM, pp. 126:1–126:11. doi:10.1145/1457515.1409079. 3