Global Illumination Test Scenes

Brian Smits
University of Utah

Henrik Wann Jensen
Stanford University

University of Utah Technical Report UUCS-00-013
http://www.cs.utah.edu/~bes/papers/scenes

Abstract. The global illumination community has discussed having a database of scenes that could be used to compare and validate different global illumination algorithms. We present a set of test scenes for global illumination algorithms and propose that they be the beginning of such a database. These scenes are designed to be as simple as possible and still test a particular aspect of energy transport. The scenes are available on a web site in a variety of formats, along with images and pixel radiance data. We feel that the simplicity and availability of the models will make it easier for the community to use them, and that the field will benefit from a standard set of scenes. Additionally, we discuss classes of models with analytic solutions.

1 Introduction

In the late 1980's, research in ray tracing benefited from the adoption of a standard set of models (the SPD) created by Eric Haines [2]. By having a common set of models and the resulting images used by many different people, it was fairly easy to tell if a new acceleration scheme was reasonable, and have some idea on how it compared to others. This set of models was highly effective, and continues to be used today.

The global illumination community could benefit greatly by having a similar set of models for testing and comparing algorithms. Greg Ward Larson has made a number of models available in his Material and Geometry Format (MGF) as part of Radiance [4]. These models have proven very useful for testing performance and overall image appearance for different rendering algorithms. Peter Shirley developed a set of models for comparison purposes[1].

In this paper we propose another set of simple test scenes. We decided to focus only on global energy transfers between surfaces. Each test scene has been designed to test different modes of transport and as such stress different parts of a given renderer. We have deliberately decided to ignore color (or spectral) information, complex reflection models, complex luminaires, and participating media. These are quantities that are represented differently by just about everyone, and new test scenes will probably not be used if they force people to make significant changes to their rendering system. Additionally, the set of models should ideally satisfy several constraints:

- each should test a single aspect of global transport algorithms.
- they should require a minimal rendering framework to compute a solution.

[1] available at the MGF reposity `http://radsite.lbl.gov/mgf/`

- they should allow easy translation between formats and reasonable replication.
- models and benchmark data should be in an easily accessible repository[2].

The goal of simplicity directly conflicts with testing geometric complexity, which was left out for two reasons. First is that it is hard to parameterize complexity. What is complex for one algorithm is often handled easily by a slightly different algorithm, potentially requiring many different models to cover the space well. Additionally, complex models would in general test many different aspects of the transport problem, complicating the results. Another form of complexity is in the description of the scenes. Our most complicated scene presented has 9 blocks and a square light. That and the regular layout should make it possible to create close approximations of them from the text and figures. While this may not make it possible to compare images pixel by pixel, it does make qualitative comparisons possible.

The reader is assumed to be familiar with various rendering techniques. A good place to start is Glassner[1]. First we will describe a class of environments that have analytic solutions. In Section 3 the set of scenes is described.

2 Scenes with Analytic Solutions

One of the benefits of having benchmarking scenes and images is that it can help verify that an algorithm is computing correct solutions. One technique that can help with this task is to use environments that have easily computed analytic solutions. The simplest analytic test case is to compute the direct transfer from an emitting polygon to a receiving plane. This can be compared at each point with the analytic differential area to area form factor. As analytic unoccluded direct illumination is often built into global illumination algorithms, it does not always test anything useful.

Another simple type of test environment for global illumination is the *furnace*. The name for this class comes from a property of a furnace in thermal equilibrium. In this situation, at every point and direction the emission at each wavelength exactly cancels out the absorption at that wavelength, and the interior has a completely uniform appearance, causing all geometric features to disappear. We can create a version of this for global illumination with an enclosed environment where every surface has the same diffuse reflectivity ρ and emitted radiance L_e. Now the solution at any point is

$$L(x) = \frac{1}{1-\rho} L_e. \tag{1}$$

The simplest proof of this involves simulating a naive path tracer. At each bounce, L_e at the current point is weighted by the current ray weight, and the current ray weight is scaled by ρ. This gives the contribution for any infinite path as

$$L_{\text{path}} = \sum_{d=0}^{\infty} \rho^d L_e. \tag{2}$$

A slight modification on this is to include any object that absorbs and emits no energy into the scene (an perfect, but not necessarily specular, reflector). This object simply changes the direction of a path based on its BRDF (perfect importance sampling). This bounce can be ignored for the purpose of counting the depth, as it has

[2] http://www.cs.utah.edu/~bes/graphics/scenes

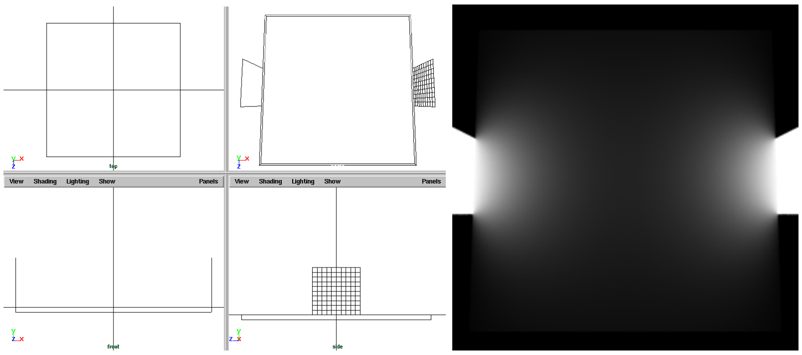

Fig. 1. Overview and solution for the emission test case.

no effect on the accumulated weight or returned luminance. Therefore the above path tracing proof still holds.

Although furnaces allow analytic solutions for complex environments, the path tracing proof shows a flaw for using them to verify a brute force path tracer. Since any ray in any direction contributes the same amount to the calculation, the path tracer may be distributing rays incorrectly without that showing up as an error in the solution. These environments do work very well for testing finite element approaches.

The property that allows the basic furnace to be evaluated analytically is that every point has the same incident radiance distribution at every bounce, therefore the solution simply involves calculating the infinite sum. This property can be generalized by looking at environments partitioned into n subsets P_i where all points in P_i have exactly the same view of the world on bounce d. This creates a simple finite element systems of low dimension. These are easiest to find when $n = 2$, such as a pair of infinite planes or concentric spheres each with different BRDF's and emissions.

3 Test Scenes

One of the primary goals in creating these models was simplicity. For that reason, the geometric primitives are all polygons, there are two triangles in one scene (that could be ignored if needed), and the rest are rectangles. The objects are all solid, except for the lights which are represented as single rectangles that emit only on the front facing side. There are a total of three different materials in the scenes. In the first four there are only two, a diffuse reflector with reflectivity of .5 and a diffuse (Lambertian) emitter with reflectivity of 0 and an emitted radiance of 1 (all results are for a single wavelength). In the final scene there is also a perfect mirror, with specular reflectivity of 1.

All of the reference images have been rendered using path tracing [3] with 100000 samples per pixel and with a resolution of 256x256.

3.1 Emission Test

This scene is composed of two rectangular lights illuminating a ground plane. The right rectangular lights is tessellated into 10 by 10 polygons, and the left is a single polygon (Figure 1). The scene tests accurate unoccluded transfer from lights to receivers. In particular, it checks sampling algorithms and that the effect of the 100 small lights is

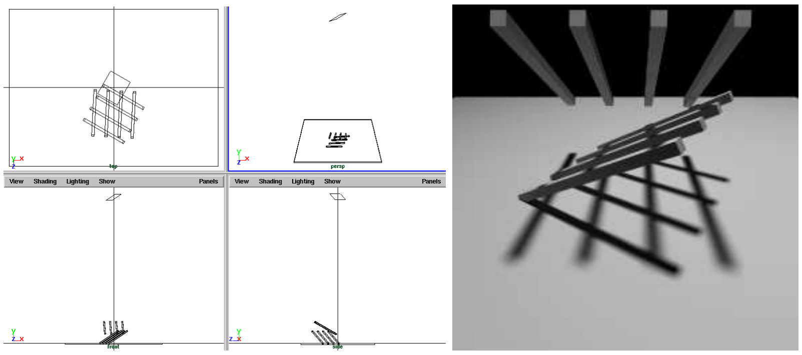

Fig. 2. Overview and solution for the shadow test case.

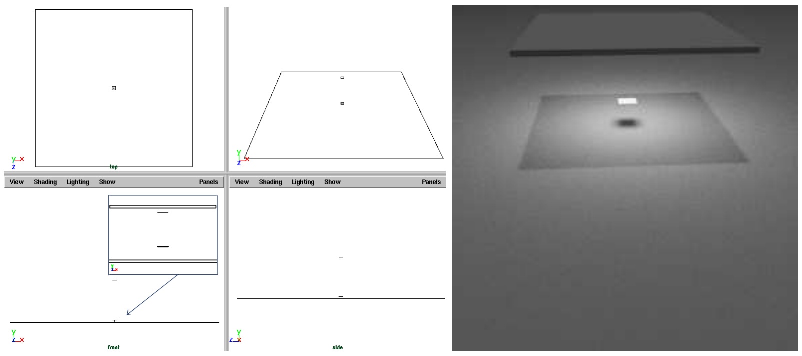

Fig. 3. Overview and solution for the huge test case.

equal to the the single large light. The symmetry of the image allows a simple sanity check by flipping the image across the vertical axis and comparing it with itself.

Most algorithm methods should be able to compute an arbitrarily accurate solution to this environment given enough time. Mesh-based approaches may have some trouble with the singularity at the corners, and algorithms that draw samples from the light may have trouble as the distance to the light goes to zero.

3.2 Shadow Test

This scene tests accurate transfer in the presence of occluders. It contains a single light source, 8 long beams, and a ground plane (Figure 2). The light is rotated so that it is parallel to the plane formed by the tops of the lower set of beams and so that its edges are parallel with the major edges of the lower set of beams. The height between the occluding objects and the floor varies from penetrating the ground plane to high enough to eliminate the umbra. The orientation of the light causes D1 shadow discontinuities on the ground plane.

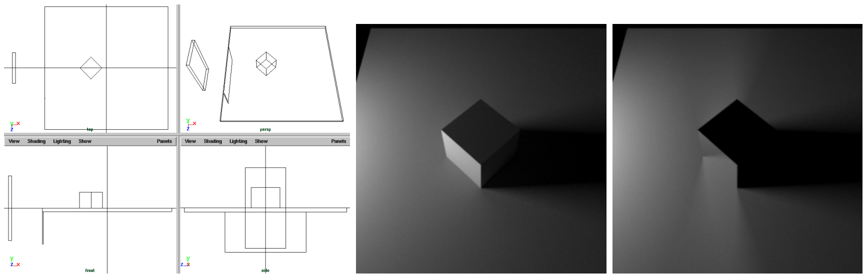

Fig. 4. Overview (left) and solution (middle) for the secondary test case. A simple modification creates a scene with a reflected caustic (right).

3.3 Large Spatial Extent

This scene has a large spatial extent, however the camera is positioned on a very small portion of the scene. The scene itself consists of a large ground plane, light, and blocker directly below and the same size as the light. Underneath the large blocker is a small light and blocker pair scaled down in width and depth by a factor of 10. The small light aims up. The camera looks under the large blocker (Figure 3). Algorithms that compute global solutions or sample lights based on the energy of each light may have trouble getting a high enough accuracy over the visible pixels, as 99% of the energy in the scene is due to the large light, but much of the visible energy is due to the small light.

3.4 Complex Illumination Paths

This scene checks complex illumination paths. It consists of a ground plane with a light underneath. The light shines on a diffuse surface to the side of the ground plane. All illumination on the visible objects is from this secondary luminaire (Figure 4). The secondary luminaire and the block on the ground plane create a very definite shadow and the light reflected off the block creates a visible indirect effect. Light can arrive in the umbra of this shadow only after 3 or more other interactions with diffuse surfaces.

A simple change to this scene can create a good test of the correct handling of specular objects. Replacing the diffuse cube with a mirrored cube creates a scene with a reflection of a caustic caused by a secondary source. It also eliminates the possibility of light reaching the umbra.

3.5 Geometric Accuracy

This model consists of 3 cubes and a wedge on a ground plane (Figure 5). It contains a few of the more common geometric problems in models. The leftmost cube penetrates the ground plane by 1% of its height. The rightmost cube is raised above the ground plane by 1%. A corner of the center cube sticks through the ground plane as well. The center object is a wedge viewed from the thin side. All objects are rotated slightly from alignment with the ground plane axes. Many rendering methods that store energy on surfaces or volumes will have trouble with light or shadow leaks near the point of the wedge. Some ray tracing based approaches using tolerances may exhibit a similar problem here as well, as the shadow rays may start on the wrong side of the wedge. The cubes check the robustness of meshing algorithms.

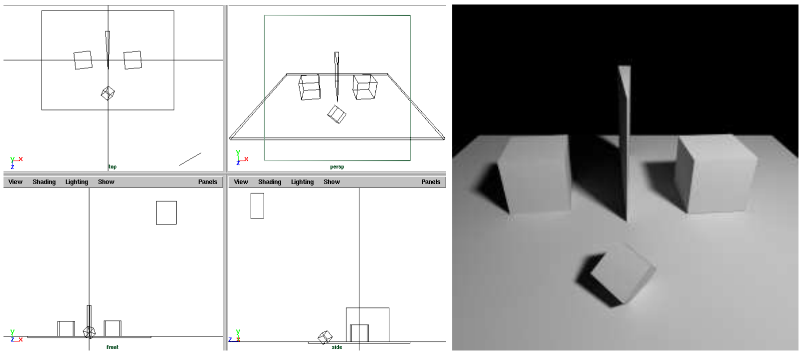

Fig. 5. Overview and solution for the geometric test case.

4 Discussion

The purpose of these scenes is not to break as many algorithms as possible, it is to exercise these algorithms. Passing these tests doesn't mean all real environments are easily and accurately rendered. All of these scenes could be made more difficult. As the area of the diffuse secondary luminaire in the caustic scene gets smaller, many algorithms get much less efficient. If the mirror cube in that scene is replaced by a sphere, some approaches break. These simple changes show where different algorithms break in a quick and easy manner.

One of the obvious things missing from the description of the test scenes is the run-time it took to generate the images. The scenes we present are not realistic. Algorithms should probably not be optimized for these cases. Some of the scenes create situations where efficiency issues show up, such as the 101 lights in the emission test. This can serve as a good reminder of some of the common assumptions on scenes.

There are a range of useful levels of adoption for a database of test scenes. These can be characterized as:

- Everyone uses these models, and agreed upon benchmark images are regularly used for testing and comparing accuracy.
- The repository grows with the addition of test scenes from other work, allowing comparisons between algorithms, but no single set of scenes becomes standard.
- The community uses the test names and similar models to discuss strengths and weaknesses of existing and new algorithms.

The first is pretty optimistic, both in assuming a universal standard and in assuming that this set of models captures every aspect worth testing in current and future algorithms. The second seems much more likely, and becomes especially useful with a maintained repository of models and data. Even the last is beneficial, however, as it allows a much less ambiguous way of describing cases where algorithms do or do not work.

5 Acknowledgments

Thanks to the reviewers from the Rendering Workshop 2000 for their helpful comments. Thanks to Peter Shirley for the simple path tracing proof for the furnace. All models were created using Maya, generously donated by Alias|Wavefront. This work was funded by NSF grant 97-20192 and DARPA DABT63-95-C0085.

References

1. GLASSNER, A. S. *Principles of Digital Image Synthesis*. Morgan-Kaufman, San Francisco, 1995.
2. HAINES, E. A. A proposal for standard graphics environments. *IEEE Computer Graphics and Applications 7*, 11 (Nov. 1987), 3–5.
3. KAJIYA, J. T. The rendering equation. In *Computer Graphics (SIGGRAPH '86 Proceedings)* (Aug. 1986), D. C. Evans and R. J. Athay, Eds., vol. 20, pp. 143–150.
4. LARSON, G. W., AND SHAKESPEARE, R. *Rendering with Radiance*. Morgan Kaufmann, Los Altos, California, 1998. Explains Radiance rendering system, comes with CD of software.