

# UC Davis

## IDAV Publications

### Title

Caustic Forecasting: Unbiased Estimation of Caustic Lighting for Global Illumination

### Permalink

<https://escholarship.org/uc/item/19h033dt>

### Authors

Budge, Brian C.  
Anderson, John C.  
Joy, Ken

### Publication Date

2008

Peer reviewed

# Caustic Forecasting: Unbiased Estimation of Caustic Lighting for Global Illumination

B. C. Budge, J. C. Anderson, and K. I. Joy

Institute for Data Analysis and Visualization, Department of Computer Science, University of California, Davis  
e-mail: {bcbudge, janderson, kijoy}@ucdavis.edu

---

## Abstract

*We present an unbiased method for generating caustic lighting using importance sampled Path Tracing with Caustic Forecasting. Our technique is part of a straightforward rendering scheme which extends the Illumination by Weak Singularities method to allow for fully unbiased global illumination with rapid convergence. A photon shooting preprocess, similar to that used in Photon Mapping, generates photons that interact with specular geometry. These photons are then clustered, effectively dividing the scene into regions which will contribute similar amounts of caustic lighting to the image. Finally, the photons are stored into spatial data structures associated with each cluster, and the clusters themselves are organized into a spatial data structure for fast searching. During rendering we use clusters to decide the caustic energy importance of a region, and use the local photons to aid in importance sampling, effectively reducing the number of samples required to capture caustic lighting.*

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism; I.3.6 [Computer Graphics]: Methodology and Techniques

---

## 1. Introduction

The term “global illumination” describes a class of algorithms that simulate complex light behavior within three-dimensional scenes, including soft shadows, color bleeding, illumination where no luminaire is visible, and caustic lighting. As part of the full global illumination solution, caustic lighting is important to capture accurately because it generates effects that dramatically impact rendering realism. Well-known caustic effects include distorted patterns on the bottom of a swimming pool, light streaming through stained glass onto a floor, and the cardioid-shaped light on the bottom of an empty coffee mug. These effects are generated by light paths that interact with one or more specular materials, such as water, glass, or metal.

The costly nature of global illumination algorithms can have a limiting effect on scene complexity. In order to accelerate the computation of complex illumination effects, biased illumination estimations have often been used. However, scene realism can be negatively affected by inaccurate illumination approximations. Thus, our work is motivated

by the need for fast, unbiased methods to capture complex global illumination effects.

Stochastic processes are often used to estimate global illumination solutions. It is common in these algorithms to handle caustic illumination in the same manner as indirect lighting – typically leading to areas with caustic effects that converge much slower than areas without caustics. Techniques that accelerate caustic illumination computation have been introduced, however, the price paid is often bias in the final image (bias is systematic error in an estimator, and is discussed further in Section 2).

We extend a highly efficient global illumination technique, Illumination by Weak Singularities (IWS), to account for unbiased caustics. The IWS algorithm is unbiased for non-caustic lighting, and very efficient. To calculate full global illumination, however, IWS must be paired with another algorithm to calculate caustics. To match IWS’s rapid convergence, Photon Mapping has typically been used in a combination we call IWS+PM. While Photon Mapping converges rapidly, it is a biased rendering method. In this paper, we pair IWS with a modified Path Tracing algorithm called

*Caustic Forecasting* that uses caustic power clustering to define splitting rules, and the  $n$ -nearest neighbor photons to build CDFs for importance sampling. The combination of IWS with Caustic Forecasting, IWS+CF, creates a global illumination algorithm that is unbiased and has rapid convergence for both caustic and non-caustic lighting.

## 2. Background

The rendering equation represents global illumination mathematically in terms of the radiance  $L$  transmitted from point  $\mathbf{y}$  to  $\mathbf{x}$  [Kaj86]:

$$L(\mathbf{y} \rightarrow \mathbf{x}) = L_e(\mathbf{y} \rightarrow \mathbf{x}) + \int_A f(\mathbf{a} \rightarrow \mathbf{y}, \mathbf{y} \rightarrow \mathbf{x}) \mathbf{G}(\mathbf{y} \leftrightarrow \mathbf{x}) L(\mathbf{a} \rightarrow \mathbf{y}) dA,$$

where  $A$  is the space of surfaces in the scene,  $f$  is the BRDF, and  $\mathbf{G}$  is the geometry term. Here, we have folded the visibility term  $V$  into  $\mathbf{G}$ :

$$\mathbf{G}(\mathbf{x} \leftrightarrow \mathbf{y}) = \begin{cases} V(\mathbf{x}, \mathbf{y}) \frac{(\vec{N}_x \cdot \vec{xy})(\vec{N}_y \cdot \vec{yx})}{|\vec{xy}|^4} & \text{if dot products } > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Numerical methods, such as Monte Carlo [HH64], are often preferred for global illumination, where analytic solutions would be difficult or impossible to derive. To numerically approximate the integral

$$I = \int_a^b f(x) dx$$

using Monte Carlo, we use the stochastic estimator

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^N \frac{f(x_i)}{p(x_i)},$$

where  $x_1, \dots, x_N$  are  $N$  random samples from the domain of integration, sampled with probabilities  $p(x_1), \dots, p(x_N)$ , respectively.

We can derive the bias of an estimator:

$$B = I - E[\langle I \rangle].$$

Unbiased Monte Carlo estimators, where  $B = 0$ , will converge to their true integral solutions with enough samples. Biased estimators,  $B \neq 0$ , are not guaranteed to converge to the correct solution, however, consistent biased estimators exist which can converge to the correct solution as algorithmic parameters are changed. Example sources of bias in image synthesis include artificial bounds, cutoffs, and inaccurate simplifying approximations. It has been suggested that in order for a rendering algorithm to be truly robust, it should be unbiased [Vea97].

It can be informative to classify stochastic global illumination methods based upon their bias. Unbiased methods for *full* global illumination include Path Tracing [Kaj86], Bidirectional Path Tracing [LW93, VG94], and Metropolis Light Transport [VG97]. Biased estimators have been

used in global illumination algorithms in order to render images more rapidly. Two such examples include Irradiance Caching [WRC88] and Photon Mapping [Jen01]. Instant Radiosity [Kel97] is not algorithmically biased, however most implementations place artificial bounds on point light computation in order to avoid blooming effects; this bounding is an example of bias.

The Illumination by Weak Singularities (IWS) method by Kollig and Keller [KK04] is an example of a method that splits the rendering equation into two parts: one for diffusive illumination  $L_d$  and another for caustic illumination  $L_c$ :

$$\begin{aligned} L(\mathbf{y} \rightarrow \mathbf{x}) &= L_d(\mathbf{y} \rightarrow \mathbf{x}) + L_c(\mathbf{y} \rightarrow \mathbf{x}) \\ L_d(\mathbf{y} \rightarrow \mathbf{x}) &= L_e(\mathbf{y} \rightarrow \mathbf{x}) \\ &\quad + \int_{A_D} f(\mathbf{d} \rightarrow \mathbf{y}, \mathbf{y} \rightarrow \mathbf{x}) \mathbf{G}(\mathbf{y} \leftrightarrow \mathbf{x}) L(\mathbf{d} \rightarrow \mathbf{y}) dA_D \\ L_c(\mathbf{y} \rightarrow \mathbf{x}) &= \int_{A_S} f(\mathbf{s} \rightarrow \mathbf{y}, \mathbf{y} \rightarrow \mathbf{x}) \mathbf{G}(\mathbf{y} \leftrightarrow \mathbf{x}) L(\mathbf{s} \rightarrow \mathbf{y}) dA_S \end{aligned}$$

The IWS algorithm (as described in Section 3.1) provides a rapid unbiased estimator for the non-caustic illumination part of the global illumination solution  $L_d$ . It does not, however, handle caustic lighting  $L_c$ . As a result, IWS is most often paired with Photon Mapping (IWS+PM) to capture caustics because Photon Mapping can rapidly capture caustic lighting [LC04]. This results in a technique that converges quickly, but with bias due to the use of Photon Mapping [Jen01].


We present a new, unbiased importance sampling method for caustic illumination estimation called *Caustic Forecasting*. This method influences the tracing of paths in two important ways: it detects regional caustic importance to control splitting and tracing, and it performs directional importance sampling via a method similar to that presented by Jensen [Jen95]. We replace Photon Mapping with Caustic Forecasting in the IWS framework to produce a rapidly converging, fully unbiased global illumination scheme called IWS+CF which handles all paths commonly handled by Path Tracing.

## 3. Method

Our IWS+CF algorithm splits the rendering equation into two parts. We evaluate diffusive illumination  $L_d$  using IWS, while  $L_c$  is estimated by stochastically tracing paths with Caustic Forecasting.

### 3.1. Illumination by Weak Singularities (IWS)

The IWS algorithm is in the class of Bidirectional Path Tracing algorithms, and is based directly on the Instant Radiosity algorithm [Kel97, KK04]. IWS is a multi-pass algorithm. In each pass two steps are performed. In the first step, the scene is populated with point lights by particle tracing from the



**Figure 1:** Photon shooting. Direct point-lights are mustard colored, indirect point-lights are white, and caustic forecasters are black.

light source(s). These point lights are then used in the second step to estimate lighting along diffuse vertices of light paths from the eye.

The particle tracing preprocess is very similar to Photon Mapping, and is illustrated in Figure 1. *Direct* point-lights (mustard colored in the figure) are generated by sampling on the surface of all light sources in the scene. At each direct point-light location, a sample direction is chosen probabilistically, and a new particle is shot into the scene. If the particle intersects a diffusive surface, a new *indirect* point-light is stored (white), and the path is continued.

During the rendering phase of the algorithm, diffusive illumination  $L_d$  is calculated at a point  $\mathbf{p}$  by tracing rays to all of the direct and indirect point lights in the scene, computing the geometry term  $\mathbf{G}(\mathbf{p} \leftrightarrow \mathbf{x}_i)$  from Equation 1 (including visibility), and then bounding  $\mathbf{G}$  to  $b$ , a fixed bounding term, if necessary. All values are then summed as a biased estimate of the non-caustic lighting. To unbiased this estimate, a new ray is shot, typically in a direction determined by importance sampling the surface BRDF, and upon intersection, the path is terminated if  $\mathbf{G}(\mathbf{p} \leftrightarrow \mathbf{y}) < b$ , otherwise, the path is accepted, and any incident radiance is attenuated by  $\frac{\mathbf{G}-b}{\mathbf{G}}$ .

To ensure that all possible paths are considered, the direct and indirect point-light sources must be changed on occasion in order to avoid bias. Typically, a full image of the scene is computed with a single sample per pixel, followed by the re-computation of the direct and indirect point-lights with another particle trace.

The original IWS algorithm does not handle caustic light-

ing  $L_c$ . In order to approximate the full global illumination solution, IWS has typically been paired with Photon Mapping (IWS+PM). To use Photon Mapping with IWS, additional *caustic* photons are stored during the particle tracing stage when a particle interacts with a specular object and then intersects a non-specular surface. These caustic particles are shown as black dots in Figure 1. The estimation of caustic lighting is then performed as described in the work of Jensen [Jen01].

Due to the use of Photon Mapping for caustic lighting, IWS+PM is a biased estimator of the full global illumination solution. In the next part, we detail a novel, unbiased estimator of caustic lighting that couples well with diffusive illumination estimators like IWS.

### 3.2. Caustic Forecasting (CF)


In IWS+PM, caustic points are used to build a photon map, which is then used to estimate the caustic lighting term. We use caustic point lights differently than in Photon Mapping; in our algorithm, the caustic points are used to determine a rough estimate of the history of caustic particles. This history is used to aid in importance sampling, and as part of a cluster-based power heuristic to determine a splitting factor.

#### Preprocessing


The first step in the Caustic Forecasting preprocess is to trace particles through the scene starting from the light sources. Caustic points are stored when a particle interacts with a specular object and then intersects a non-specular surface (shown as black points in Figure 1). No interactions are stored at specular intersections. The goal of tracing particles and storing caustic points is to obtain the raw data needed to estimate the distribution of caustic lighting energy within the scene.

To refine our estimate of the distribution of caustic energy we perform clustering over the caustic particles. We cluster over the six dimensional space formed by caustic particle position and corresponding surface normal (i.e., two three-component vectors). Clustering simply groups nearby particles with similar surface normals into the same class, which allows for the calculation of aggregate statistics over many particles. By incorporating the surface normals – rather than just clustering based upon position – the classification is sensitive to the underlying scene geometry (e.g., creating cluster boundaries at  $90^\circ$  angles, multiple clusters on rounded objects, etc.). We have tried both  $k$ -Means [Mac67] and Vector Quantization (VQ) [AKCM90] clustering of the particle positions, and have found that VQ produces more intuitive clustering results – in particular, clusters were better distributed with particle density and more uniformly sized.


After clustering, we calculate the energy density of each cluster (i.e., total photon energy versus cluster area) in order to rank the clusters into levels of importance. The ranking of




(a) Rank 0 (50% Total Energy)



(b) Rank 2 (87.5% Total Energy)



(c) Rank 4 (96.88% Total Energy)



(d) Rank 8 (99.80% Total Energy)

**Figure 2:** Caustic photons are grouped into clusters using Vector Quantization (VQ). The clusters are then ranked based upon their energy contribution to the scene: the fraction of energy within rank  $[0, n]$  is  $1 - 0.5^n$ .

a cluster is calculated based upon its caustic energy contribution to the scene: low rank clusters have high power, while higher rank clusters are decreasingly powerful. Consider a list  $C$  of  $k$  clusters, sorted in descending order power based upon the cluster power (i.e., most powerful first). Mathematically, a cluster  $C_j$  is assigned rank  $n$  if all clusters  $(C_0, \dots, C_j)$  account for at most  $1 - 0.5^n$  of the total fractional energy of all clusters.

Our eventual goal is to use the energy estimation provided by the clustering to tune our sampling strategy for evaluating caustic lighting during rendering. If the cluster has very low energy, we perform a number of simple cosine samples determined by the diffusive ray depth. In high-energy clusters, we use information from the gathered photons to drive a local importance sampling strategy. Because our sampling strategy is more expensive than cosine sampling, we clamp the maximum rank of a cluster for which importance sampling will be considered *a priori*. Figure 2 provides insight into the energy distributions defined by the clusters: clusters drawn in color are considered during importance sampling, while black and white clusters have been culled.

Caustic points near centroids of the remaining clusters – which often account for over 99% of the scene’s total caustic energy – are added to a *kd-tree* called the *cluster map*. This structure allows for rapid cluster membership testing when calculating the caustic illumination. Additionally, we build per-cluster *kd-trees* called *forecast maps* containing the caustic points within the cluster and up to a surrounding epsilon.

Finally, the preprocess ends by using the remaining high-energy clusters to design a maximum search radius  $r$  for forecast map lookups. The radius is a heuristic based upon the average photon density over the clusters: we aim to find approximately  $n$  neighboring photons in regions of average density. This ensures that for very important clusters we will find at least  $n$  neighbors, while in less important clusters we will cull more distant photons.

### Caustic Forecast Importance Sampling

To calculate the caustic lighting contribution at point  $p$ , we begin by searching for the closest cluster to  $p$  in the cluster map. If no valid cluster is found, we simply take some number of samples from a cosine distribution. Otherwise,  $p$  is in a high-energy cluster and we will employ Caustic Forecasting importance sampling.

Importance sampling begins by gathering the  $n$ -nearest neighbor photons to point  $p$ . Here  $n$  is the maximum number of forecasters to gather, and we say  $k \leq n$  is the actual number gathered within the search radius  $r$ . For  $k$  gathered forecasters, the incoming power of each forecaster is modulated with the BRDF and the cosine to add weight to a bin in a 2D table (azimuth versus cosine of zenith) as done by Jensen [Jen95]. Next, we calculate the number  $S$  of splitting



samples to take over the table based on the importance of the cluster (i.e., its energy) and our current diffusive ray depth.

Due to the computational cost associated with large number of photons and large  $n$  gathers, the table size is generally much larger than  $n$ , and many of the table bins have zero energy deposited within them. Because some bins are empty, one cannot simply create a CDF based on this table for unbiased sampling (some valid regions of the hemisphere would be unreachable). Jensen’s solution was to insert a small amount of energy into the empty bins before summing the table. We found that by cosine-weighting the small unbiasing contribution, variance can be reduced. Another approach which works slightly better for caustics is to split the integral yet again. Instead of filling in the zero bins, we leave the bins zero, and create a dual table which has zeros where the original table has non-zero values (and only those bins). The values we fill into the new table are proportional to a representative cosine for the bin multiplied by the BRDF for the same representative outgoing vector.

At this point, the valid directional sampling tables are normalized and summed into discrete CDFs that can be sampled. Because these CDFs are monotonic, they can be sampled by first choosing a uniform sample and then using a binary search to find the corresponding CDF bin. For full details on the directional sampling, we refer the reader to the paper by Jensen [Jen95]. Finally we sample both tables  $S$  times (based on our original cluster importance estimate and the diffusive depth), and *add* the contributions together. (Notice the distinction; when sampling  $S$  times over a whole domain, the estimate is the sum of the evaluated samples divided by  $S$ , but when splitting the domain, each part is independently estimated, and the estimates are summed).

Finally, we make two observations that can speed up rendering. The first is that, while direct and indirect light sources must be refreshed in IWS to avoid bias, this is not necessary for the caustic forecasters. In fact, as long as the entire hemisphere is accounted for (by filling in zero bins as in Jensen’s work, or splitting the domain as in this paper), the use of the CDF table(s) for importance sampling will produce unbiased results. This is of great benefit because the initial photon casting phase can be quite costly.

Second, because our use of CDF sampling is unbiased, we can shoot directly from the light source(s) toward specular objects in order to increase the number of caustic points gathered during the preprocess. As long as we correctly construct the sampling table(s) as described above, this will increase the quality of the forecast maps for importance sampling without introducing bias.

A good example of this can be seen in the *Glass Figurines* scene (in Figure 3). The scene itself is very large, however, we primarily focus on the two glass figurines and the metal tin. We can place a (non-visible) box around these objects, and start our particle paths from the light source toward the

box. It is important to note that the distribution of these particles will be biased! However, because sampling from our CDF will remain unbiased, and properly employed, splitting will always be unbiased, this causes no bias in the image. The direct point lights and indirect point lights in IWS must not be generated from this distribution, as bias would certainly be introduced.


#### 4. Results

Our test scenes can be seen in Figures 3, 4, and 5. Figure 3 shows our *Glass Figurines* scene which contains a glass horse and dragon next to a metallic box. Figure 4 is a *Cornell Box* scene with a metallic sphere (left side of the scene) and glass sphere (right side). Figure 5 is our *Ring* scene which is designed to create a partial cardioid-shaped caustic effect on the interior of a metallic ring.


All images were produced using 20 CPU cores via MPI on a small cluster consisting of 5 nodes, each having two, dual-core Opteron 2216 processors and 4 GB of RAM. The nodes run 64-bit Linux, and the software uses MVAPICH2 for interprocess communication. It should be noted that our renderer was built for maximal flexibility; because of this, our renderer is typically about 50 times slower than a state-of-the-art non-packet-based ray tracer. On the other hand, since flexibility is not compromised, operations such as photon map lookups and CDF creation are heavily optimized. A breakdown of algorithm runtime is shown in Table 1.

The *Glass Figurines* scene in Figure 3 consists of several hundred thousand triangles with multiple materials, including Dirac (the figurines), Lambertian, and Phong-like glossy material. The scene is very large when compared to the figurines, and is a difficult case for photon shooting algorithms when attempting to render caustics because the natural density of photons reaching the figurines is small. The images were rendered at 640x360 resolution for 60 minutes. In that 60 minute time frame, IWS+CF produced 1442 samples per pixel, Path Tracing produced 8891 samples per pixel, and IWS+PM with 500K caustic photons collected produced 1108 samples per pixel. Our method, IWS+CF, produces an unbiased rendering like Path Tracing, however the entire IWS+CF image is better converged. The caustics using Photon Maps (IWS+PM) are fully converged, but are clearly incorrect. More accurate, but still biased, caustics using Photon Mapping would have required the collection of many times more photons, which would have taken longer than the 60 minute rendering window allowed in our test.


The *Cornell Box* scene in Figure 4 was rendered at 640x640 resolution for 4 minutes. This simple scene contains only a handful of objects, but nicely shows primary as well as secondary caustic lighting. In 4 minutes, IWS+CF completed 60 samples per pixel, Path Tracing completed 405 samples per pixel, and IWS+PM completed 99 samples per pixel. IWS+CF captures the secondary caustic on the green



(a) IWS+CF



(b) Path Tracing



(c) IWS+PM (200K caustic photons)

**Figure 3:** Glass Figurines scene rendered in 60 minutes using (a) our method, IWS+CF, (b) Path Tracing, and (c) IWS+PM.

wall with greater fidelity than Path Tracing. To get comparable fidelity of the secondary caustic using IWS+PM, 1 million caustic photons had to be collected. This also has the unfortunate effect of causing some bright blooming effects in the image, which could be removed by further biasing the image with more aggressive filtering for our density estimation. The other, lighter blooming effects could likely be removed by casting photons for each IWS+PM pass, but this would be computationally prohibitive. A close up comparison of IWS+CF and Path Tracing is shown in Figure 6.

The *Ring* scene in Figure 5 emphasizes a partial cardioid-shaped caustic within the ring. The ring is composed of just over 12,000 triangles. This tessellation is what leads to the

|                            | Cornell Box | Ring    | Figurines |
|----------------------------|-------------|---------|-----------|
| Total Without Caustics     | 306 s       | 869 s   | 27600 s   |
| Total With Caustics:       | 2290 s      | 2610 s  | 71600 s   |
| <i>Tracing Photons</i>     | 5.4 s       | 15.7 s  | 117.3 s   |
| <i>Clustering Caustics</i> | 9.4 s       | 47.1 s  | 3.5 s     |
| <i>Building kd-trees</i>   | 4.9 s       | 21.7 s  | 1.28 s    |
| <i>Creating CDFs</i>       | 190 s       | 172.1 s | 2321.6 s  |
| Forecast Photons           | 200 K       | 1 M     | 75 K      |

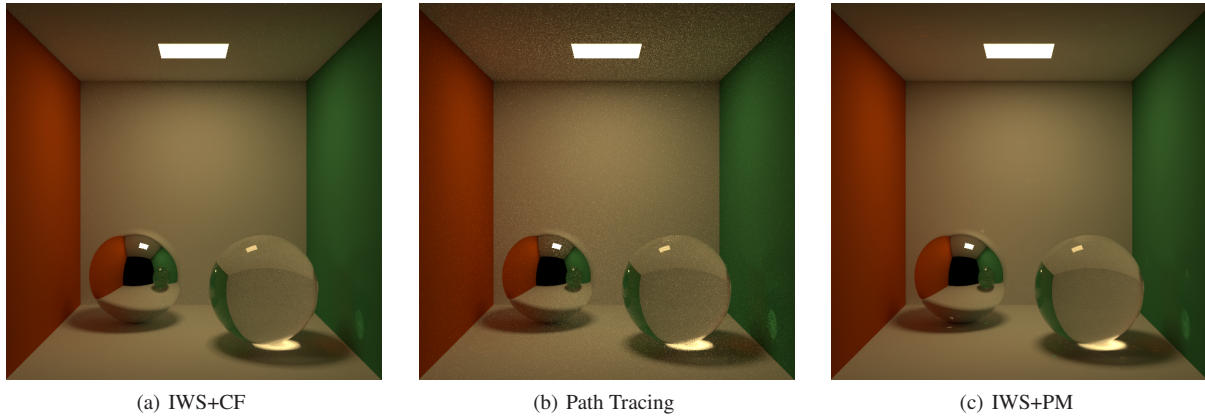
**Table 1:** We show the breakdown of timings for different components within the algorithm for images presented in this paper (the ring image corresponds to the eight minute MPI timing). The total time without caustics was measured separately by turning off caustic calculations. For ease of timing, these measurements were taken in a single-processor environment.

bright lines within the cardioid. We rendered images using IWS+CF, Path Tracing, and IWS+PM across four run times: 4, 8, 16, and 32 minutes. IWS+CF produces nearly converged results after only 4 minutes. The caustic effects within the Path Traced images have noticeable grain through 32 minutes. IWS+PM with one million photons converges quickly, but exhibits visual artifacts including light bleeding from the wrong side of the ring, and the caustic “fingers” in the middle of the ring that are too dark and blurred. Additionally, it is not possible to render the IWS+PM image with one million photons in 4 minutes, as the photon tracing phase barely finishes within that time frame. A faster ray tracer would have allowed the tracing phase to finish more quickly, however, results from the other two techniques would also benefit from faster ray tracing. Figure 6 shows a close up comparison of IWS+CF and Path Tracing. Notice how quickly the high frequency features converge when compared to strict Path Tracing.

Finally, we have compared the use of a modified single-table method to the two-table method described in Section 3.2. Although the extra cost of sampling two tables led to a 16% reduction in the number of image samples per second, the image root mean squared error measured against path tracing with 20,000 samples, was reduced by just over 1% when compared to the single table method. We additionally tried the single table method with a gaussian smoothing of the energy into neighboring table cells. The convergence of this method was very poor compared to the other two methods. We attribute this to the fact that, although our tables are larger than Jensen’s, our tables are still relatively coarse (typically we use  $32 \times 16$  bins). This means that each bin covers a large portion of the hemisphere, and a naive bin-based smoothing scheme can place energy into bins where none is justified.

## Conclusion

We have presented Caustic Forecasting, a technique which can be paired with Illumination by Weak Singularities (IWS)



**Figure 4:** Cornell Box scene rendered in just over 4 minutes with (a) our method, IWS+CF, (b) Path Tracing, and (c) the original IWS+PM algorithm. IWS+CF has converged better than both methods; IWS+PM even exhibits a number of light blooming artifacts caused by the large number of caustic photons required to capture the secondary caustic on the green wall.

to create a totally unbiased rendering algorithm for rendering images with global illumination. In this paper, we have referred to this combined algorithm as IWS+CF. The IWS+CF method allows for convergence at nearly the same rate as the original IWS algorithm with Photon Mapping (IWS+PM), but with correct, unbiased caustic lighting.

IWS+CF has proven to be fairly robust and competitive in terms of speed and quality. IWS+CF has a comparable rate of convergence to IWS+PM, with the benefit of being unbiased, and typically converges much faster than Path Tracing. Even with the strengths of our method, there is room for improvement; computing caustics is still the bottleneck in our global illumination calculations. IWS *without caustics*, for example, converges very quickly – several times faster than when paired with our technique (or with Photon Mapping). It is possible that through tuning of our heuristics even more efficiency might be derived from the algorithm.

Finally, a potential limitation of our method is one that is shared by Path Tracing: we cannot compute caustic lighting if the original light sources are point lights. Although point light sources are not completely realistic, in computer graphics they are widely used, and it would be useful to be able to compute “unbiased” caustics in scenes containing point lights.

### Acknowledgments

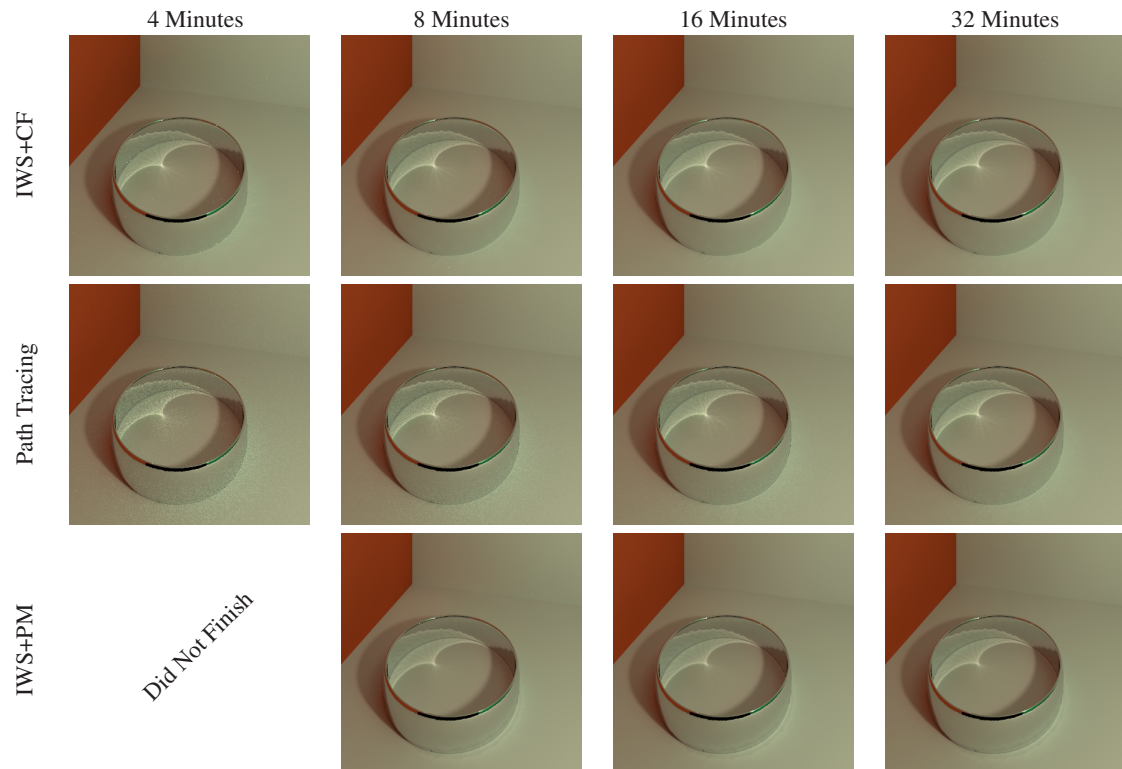
We would like to thank our colleagues in the Institute for Data Analysis and Visualization (IDAV). Our work was supported by Lawrence Livermore National Laboratory, by a United States Department of Education Government Assistance in Areas of National Need (DOE-GAANN) grant #P200A980307, and by the Office of Science, U.S. Department of Energy under Contract No. DE-AC02-05CH11231 with the Scientific Discovery through Advanced Computing (SciDAC) Visualization and Analytics Center for Enabling Technologies (VACET).

submitted to *Pacific Graphics* (2008)


### References

- [AKCM90] AHALT S. C., KRISHNAMURTHY A. K., CHEN P., MELTON D. E.: Competitive learning algorithms for vector quantization. *Neural Networks* 3, 3 (1990), 277–290.
- [HH64] HAMMERSLEY J. M., HANDSCOMB D. C.: *Monte Carlo Methods*. Methuen & Co, London, 1964.
- [Jen95] JENSEN H. W.: Importance driven path tracing using the Photon Map. In *Rendering Techniques* (1995), pp. 326–335.
- [Jen01] JENSEN H. W.: *Realistic image synthesis using Photon Mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001.
- [Kaj86] KAJIYA J. T.: The rendering equation. In *Proc. of SIGGRAPH* (1986), pp. 143–150.
- [Kel97] KELLER A.: Instant radiosity. In *Proc. of SIGGRAPH* (1997), pp. 49–56.
- [KK04] KOLLIG T., KELLER A.: Illumination in the presence of weak singularities. *Monte Carlo and quasi-Monte Carlo Methods* (2004).
- [LC04] LARSEN B. D., CHRISTENSEN N.: Simulating Photon Mapping for real-time applications. In *Eurographics Symposium on Rendering* (June 2004).
- [LW93] LAFORTUNE E. P., WILLEMS Y. D.: Bi-directional Path Tracing. In *Proc. of Compugraphics* (1993), pp. 145–153.
- [Mac67] MACQUEEN J. B.: Some methods of classification and analysis of multivariate observations. In *Proc. of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* (1967), pp. 281–297.
- [Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, Dec. 1997.
- [VG94] VEACH E., GUIBAS L.: Bidirectional estimators for light transport. In *Fifth Eurographics Workshop on Rendering* (June 1994), pp. 147–162.
- [VG97] VEACH E., GUIBAS L. J.: Metropolis light transport. In *Proc. of SIGGRAPH* (1997), pp. 65–76.
- [WRC88] WARD G. J., RUBINSTEIN F. M., CLEAR R. D.: A ray tracing solution for diffuse interreflection. In *Proc. of SIGGRAPH* (1988), pp. 85–92.





**Figure 5:** The Ring scene emphasizes a partial cardioid-shaped caustic within the ring. IWS+CF, Path Tracing, and IWS+PM across four run times. IWS+CF produces nearly converged, unbiased results after only 4 minutes. The caustic effects within the Path Traced images have more noticeable grain. IWS+PM converges quickly, but with an incorrect, biased result. (The photon shooting phase for IWS+PM with one million photons runs for nearly the entire allotted 4 minutes, leaving no time for image rendering.)



**Figure 6:** Closeups of (a) the Cornell Box after 4 minutes and (b) Ring scenes after 8 minutes. For each scene, the results of IWS+CF are to the left of Path Tracing.