

A Non-parametric Guide for Radiance Sampling in Global Illumination

Pankaj Khanna¹ Mel Slater^{1,2} Jesper Mortensen¹ Insu Yu¹

¹Department of Computer Science, University College London, London, UK

²ICREA-Universitat Politècnica de Catalunya, Departament de LSI, Barcelona, Spain

p.khanna | m.slater | j.mortensen | i.yu @cs.ucl.ac.uk

Abstract

Global illumination as described by the radiance and potential equations is essentially a sampling problem. This sampling has two elements: ‘where’ to sample and ‘how’ to sample. In this paper we propose an alternative method to non-parametrically simplify the problem of ‘where’ in the general case. We present a solid angle based uniform sampling scheme with an efficient conservative cull of solid angles not affecting the transfer of radiance between two surfaces. This construction also provides a geometric area bounds within which a surface can receive energy from another along directions contained in the solid angle. The technique can be used to illuminate a scene from a general emitter, including incident light fields. The method has several applications and advantages, with results showing good sampling efficiency and speed.

1. Introduction

The illumination of a point depends on many factors. The spatial and directional distribution of light across the environment, visibility, geometric and material properties of surfaces influence its illumination as described by the radiance [14] and potential [21] equations. Finding a solution to this transport problem requires efficient sampling and reconstruction of radiance.

Parametric methods of importance sampling use *a priori* knowledge of the surfaces, materials and emission characteristics of a scene to efficiently place samples for faster convergence and low variance. Such approaches are however not always feasible – especially, when the domain does not easily lend itself to a parametric definition (for example, physically measured BRDFs or light fields). In such situations, one resorts to sampling schemes that amongst other approaches, attempt to progressively build a parametric model of the domain. In this paper, we consider uniform sampling without any parametric knowledge of the domain. Such blind sampling schemes are usually characterised by

long process times and high variance. We suggest an alternative approach that performs a uniform sampling of directions over a planar surface. With a stratification of directions into solid angles, we are able to efficiently conduct a conservative cull of non-contributing solid angles for significant speed-up and performance approaching that of alternative parametric methods. This process not only provides a set of potentially contributing solid angles, but for each solid angle also provides a geometric bounds within which to sample the ‘source’ surface with directions contained in the solid angle. We can thus benefit from visibility coherence in sampling radiance in an area. In this paper we have applied this approach to illumination computation from idealised emitters and incident light fields, however it has several applications beyond those presented here.

The remainder of this paper is organised as follows. In Section 2, we briefly touch upon some previous work related to our method. Next, in Section 3, we discuss the motivation and foundations for our approach, with additional discussions in Section 4. Implementation issues are highlighted in Section 5, and a simple implementation presented and analysed in Section 6. Finally, we present our conclusions and future directions of research in Section 7.

2 Previous work

Global illumination is a key component of realistic image synthesis; the aim being to compute all incident light at a point – not only directly from light-sources, but also by inter-reflection via other surfaces in the scene. Adequately sampling this incident radiance can be very computationally expensive, and simplifications are often necessary to contain this complexity. Ray-tracing [31] and radiosity [9] place limitations on the range of scene materials thereby simplifying sampling requirements. While several extensions have been suggested to extend these methods by the addition of a wider range of materials, they all have associated sampling costs. Amongst other methods, path-tracing [14] and photon-mapping [13] present more complete solutions to global illumination, allowing the representation and

sampling of a wide range of materials.

A popular and efficient approach to solving this sampling problem is by resorting to Monte Carlo techniques. A categorisation of these techniques can be made on the distinction of ‘blind’ versus ‘informed’ approaches. The former including methods of stratified and adaptive sampling by sample rejection. The main disadvantage of such methods is the significant computation needed to achieve convergence and low variance. These fall within the category of variance reduction techniques termed importance sampling; the goal of which is to direct sampling to regions of ‘high interest’. ‘Informed’ sampling techniques use *a priori* knowledge of the sampling domain to reduce the estimator’s variance. The most common approach to importance-sampling uses a probability distribution function (PDF) that closely approximates the domain to target sampling for low variance [3,22]. This however also necessitates the computation of the inverse of its cumulative density function (CDF) (for example [2, 17]) – which may not always be possible. Multiple importance sampling techniques [28] offer enhanced sampling performance by considering a combination of several estimators. Metropolis light transport [29] presents an alternate, robust approach to sampling general scenes with very good results. Importance re-sampling techniques require a PDF of the domain though they do not demand the invertibility of its CDF [18, 26]. We refer the reader to [8] for in-depth coverage of the subject. Our approach is an alternative to many of these methods; allowing stratified uniform sampling by sample rejection within a directional and spatial bounds by efficient, conservative solid angle culling.

A further sampling issue is related to whether the global illumination solution emulates a continuous (such as [9, 13, 14]), or discrete [5, 24, 25] sampling of the environment. While methods employing a discretised ray-space offer several advantages, they can be prone to under-sampling and aliasing. We believe that several such approaches can benefit from a directional discretisation into solid angles rather than into discrete ray-sets.

As a scheme to represent globally illuminated scenes, the Light Field [19] and Lumigraph [10] provide an image-based representation of radiance exiting a real or synthetic scene. Subsequent to construction, the scene can be viewed from a large set of view-points by simple ray-queries in a 4D data-set. Similarly, an incident light field [11, 20, 27] represents incident illumination on a scene captured from the real world or computed from a synthetic one – an approach initially suggested in [12]. While [11] used a lumigraph, both [20] and [27] employ a sphere-plane parameterisation [6] to encode complex incident illumination and subsequently illuminate a synthetic scene or relight a light field. Our technique allows the use of similar sphere-plane incident light fields to efficiently illuminate a synthetic scene. Image-based lighting can also be performed efficiently via

the use of environment maps [1, 4], though we do not initially consider these methods due to their approximate representation of the flow of light.

3 Determining a bounds for sampling

Consider the local flow of potential [21] from a surface S along directions in a solid angle ω . This potential is defined over a subset of the visible environment of S corresponding to directions in ω . The 5-dimensional domain $\{x, y, z, \theta, \phi\}$ possibly in receipt of this potential is described by the convolution $S \otimes \omega$. To find the actual regions in receipt of this potential, we need to resolve visibility from S along ω .

Both these steps – constructing a practical realisation of the convolution $S \otimes \omega$, and sampling for visibility within it are complex, compute-intensive tasks. Rather than compute $S \otimes \omega$ in 5-dimensions, we approximate it in 3-dimensions $\{x, y, z\}$. The solution of visibility in $S \otimes \omega$ can also be avoided initially by considering the approximated $S \otimes \omega$ as a bounds for ‘possible potential’. This 3-dimensional volume is considered against a single surface R , reducing it to a 2-dimensional area on it that receives can ‘possibly’ receive potential from S along ω . We obtain this area from a *projected minkowski sum* as described in the following section. Given this bounds, we solve for visibility and sample radiance from S to the specified area of R along directions in ω . By considering a discretisation of the hemisphere over R into a set of solid angles $\{\omega_i\}$ we can thus gather radiance from S , discarding non-contributing solid-angles. In this sampling, S can have an arbitrary radiance representation such as any BRDF or incident light-field representation. A simple sampling scheme using this approach is discussed in section 6.

3.1 The Projected Minkowski Sum

In general, the minkowski sum [7] of convex sets A and B is a convex set defined as:

$$A \oplus B = \{a + b : a \in A, b \in B\} \quad (1)$$

Minkowski sums also have the following properties [7]:

$$A \oplus B = B \oplus A, \quad A \oplus (B \cup C) = (A \oplus B) \cup (A \oplus C) \quad (2)$$

With a polygonal source of potential S , and solid angle ω defined as a spherical triangle, we construct their minkowski sum $S \oplus \omega$. The vertices of spherical triangle ω lie on the unit sphere in $\{\theta, \phi\}$, but if we consider the canonical sphere centred at the origin, also have positions in world-coordinates $\{x, y, z\}$. Though we have described ω as a spherical triangle, we approximate it here by a planar triangular polygon B_ω that shares its vertices. The error in this approximation is proportional to the length of the

‘spherical edges’ of ω . With large solid angles, it may be necessary to construct B_ω with additional vertices for a better approximation. For small solid angles however, the error is negligible and can be ignored – we find this as the case for a subdivision of the hemisphere into 64 or more ‘uniform’ solid angles.

We initiate the construction of the minkowski sum in the $\{u, v, n\}$ coordinate system. This space is a simple rotation from the canonical $\{x, y, z\}$ space such that the y -axis is rotated into the n -axis lying along direction Θ_ω – the centroid direction in ω . Let this rotation be defined by matrix M_{Θ_ω} . Also, as we have assumed $\omega \approx B_\omega$, B_ω can also be transformed into $\{u, v, n\}$ coordinates by rotation M_{Θ_ω} . Let these rotated polygons be termed S' and B'_ω .

We now ignore (but do not discard) the n -coordinate of polygons S' and B'_ω , and consider their polygon-edges on the uv plane. If these edges are collectively sorted by increasing slope, with the terminating vertex of each edge stored successively, we obtain a polygon equivalent to the minkowski sum $S' \oplus B'_\omega$ on the uv -plane (see figure 1). For each vertex, we also store a flag to indicate if it originated from S' or B'_ω . Finally, we rotate $S' \oplus B'_\omega$ back into (x, y, z) coordinates by $M_{\Theta_\omega}^{-1}$ to obtain $S \oplus B_\omega$.

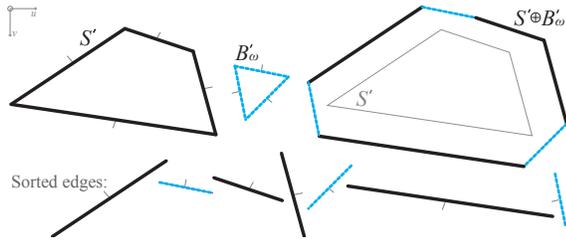


Figure 1. S' , B'_ω and their minkowski sum.

Given polygons S and $S \oplus B_\omega$ we can now construct $S \oplus \omega$ which is essentially a volume in $\{x, y, z\}$ coordinates. We iterate through the vertices of $S \oplus B_\omega$ till we have found a vertex that originally belonged to S and another that belonged to B_ω – let the most recent of these vertices encountered be referred to as v_S and v_ω respectively. The boundaries of volume $S \oplus \omega$ can then be constructed by the ray with origin at v_S and direction along $(v_S + v_\omega)$. Additional boundaries of the volume are found by updating v_S or v_ω as each subsequent vertex is encountered in $S \oplus B_\omega$. The total number of such boundary rays is equivalent to the number of vertices of S and B_ω . The rays create a pyramidal volume enclosed at its apex by S . Only regions of the scene within the volume $S \oplus \omega$ can possibly receive potential from S along directions in ω .

To further simplify computation, we consider a single surface R , and explore the possible transfer of energy between surfaces S and R along directions in ω . We intersect the bounding rays of $S \oplus \omega$ with the plane of R . This gives

us a 2-dimensional region on R , representing the only portion of its surface that can possibly receive energy directly from S along directions in ω . We refer to this region as the *Projected Minkowski Sum* of S on R through ω and denote it by $(S \oplus \omega)_R$ (see figure 2). Given its simple geometric construction, $(S \oplus \omega)_R$ can be computed quite efficiently. Depending on scene geometry, many regions $(S \oplus \omega)_R$ are not valid – either as a result of not allowing a valid clipped B_ω (see section 5), or by computing regions that do not overlap R . All these solid angles need not be sampled.

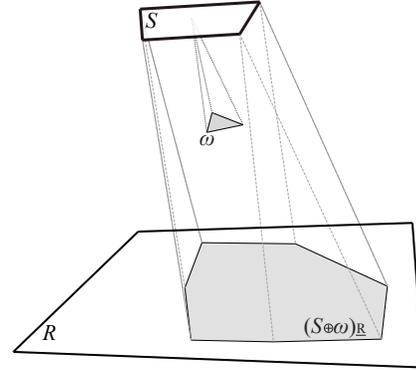


Figure 2. The *Projected Minkowski Sum* of S on R through solid angle ω .

Rather than limiting S to being a convex polygon, we can also consider more complex geometry. Consider the earlier defined (u, v, n) coordinate system. In transformation matrix M_{Θ_ω} , we also add a translation such that the n -axis is through the centre of S and the whole of S lies just above the uv -plane. We project the geometry of S orthographically along the n -axis onto this uv -plane. An impostor of S , \tilde{S} would then be constructed as the planar convex hull of this projected geometry. Using \tilde{S} in place of S , we can obtain a conservative approximation to $(S \oplus \omega)_R$. A similar impostor could also be constructed and used for R .

4 Hierarchical evaluation and sampling

We have considered a solid angle to be defined by a spherical triangle. A related problem then is of how to partition the canonical sphere into spherical triangles such that all triangles are of a uniform area and shape. This is a non-trivial problem for a partition comprising more than a small number of triangles. We instead consider a partitioning scheme which though non-uniform has a useful property of affording very fast, constant-time query for an arbitrary direction, returning the solid angle it lies in. This method unlike [23] does not require any additional lookup. The non-uniformity can be addressed during sampling by drawing additional samples or applying appropriate weights. This

scheme is presented in Appendix A, although any similar method could be considered.

Rather than have a fixed solid angle discretisation, it is possible to construct a hierarchy of solid angle partitions so that adaptive schemes may be considered. For example, with the suggested discretisation scheme (Appendix A), we can decompose a solid angle into any number of child solid angles by an appropriate choice of n (for a total of $8n^2$ solid angles). Starting with $n = 1$, we divide each solid angle into four child- solid angles by doubling n for each successive division. With $n = 1, 2, 4, 8, 16 \dots$ we obtain a partition of the sphere into $8, 32, 128, 512, 2048 \dots$ solid angles, with a solid angle at level l partitioned into four solid angles at level $l + 1$. While a direction Θ can be queried at any discretisation in constant time, we can also maintain a quad-tree data structure to explicitly store and process this hierarchical solid angle subdivision.

If we consider a subdivision of solid angle ω into a set of non-overlapping sub solid angles $\{\omega_i\}$, we have from (2):

$$S \oplus \omega = S \oplus \left(\bigcup \omega_i \right) = \bigcup (S \oplus \omega_i) \quad (3)$$

Further, considering the intercept of convex regions $S \oplus \omega$ and $\bigcup (S \oplus \omega_i)$ on the plane of surface R , we have:

$$(S \oplus \omega)_R = \bigcup (S \oplus \omega_i)_R \quad \text{where } \{\omega_i\} \text{ covers } \omega \quad (4)$$

This is illustrated in figure 3, showing the projected minkowski sum for a solid angle ω and its four sub- solid angles ω_i . Equation (4) allows for elegant adaptive sam-

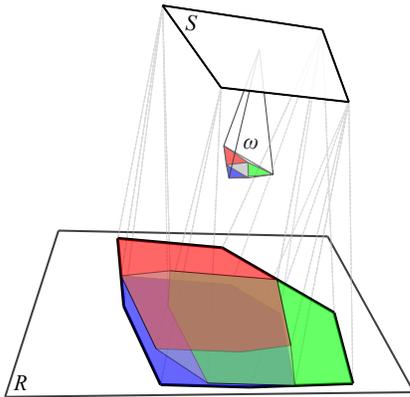


Figure 3. Projected Minkowski sums of S on R through ω and its four sub- solid angles.

pling schemes. Radiance could be sampled initially from a large solid angle, and if further refinement is deemed necessary, its sub- solid angles can be sampled in turn. This not only enforces stratified sampling, but also saves computation as the projected minkowski sums for the sub- solid angles define smaller regions on R for possible transfer. In



(a) ‘beams’ - lit by an incident light field; 2 light bounces (42 polygons – 107,984 radiance elements)
 (b) ‘classroom’ - 4 diffuse emitters; locally illuminated (3024 polygons – 1,564,342 radiance elements)

Figure 4. Test scenes used for comparison.

this approach, if the original samples were chosen equally distributed amongst the solid angle’s four child- solid angles, they can be reused later without causing bias.

5 Implementation issues

When computing the projected minkowski sum, care needs to be taken if the solid angle ω is partially back-facing either S or R . Clipping ω ’s polygonal approximation B_ω against the planes of S and R such that we only consider a valid solid angle for the transfer solves this issue. A further problem arises when projecting the minkowski sum onto R in cases where the boundary rays of $S \oplus \omega$ are parallel to R ’s plane. Deviating the ray into the plane by an infinitesimal amount removes this singularity.

If a valid projected minkowski sum can be computed – that is, if a valid transfer of potential possibly exists between S and R via ω , we determine the region on R we need to sample from. We do this in two-steps – first with an axis-aligned bounding box overlap check, and if necessary a test against the polygonal boundary of $(S \oplus \omega)_R$. It should also be noted that a significant portion of the algorithms allow for very efficient SIMD implementation – we are however yet to aggressively pursue this.

6 Radiance sampling results

We now consider our method in two scenarios – illuminating a synthetic scene with an incident light field, and comparing its performance for an idealised emitter against parametric importance sampling with a PDF.

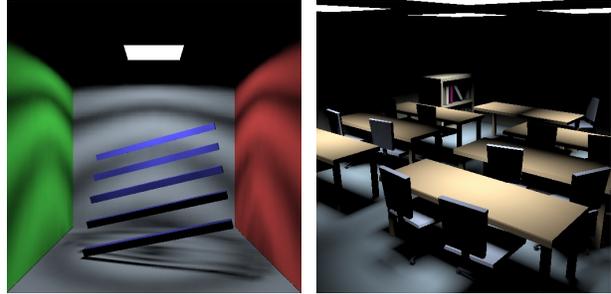
Similar to [20], we test our scene with a synthetically generated 4D incident light field, though we could easily substitute this for a real one captured using existing methods [20, 27]. For our tests, we consider two scenes, a ‘beams’ scene comprising of five diagonal beams in a ‘Cornell-type’

room, and a ‘classroom’ scene with more complex geometry and four light sources (see figure 4). We use a light field parameterisation of $\mathbf{L}_i(\theta, \phi, x, y)$, with incident radiance stored at each 4-tuple. In our incident light fields, we consider discrete θ and ϕ similar to that used for the centroids of our hemispherical partition scheme (see Appendix A). When querying irradiance along a ray, we determine the solid angle the ray’s direction lies in and its 3-connected neighbours and query an interpolated radiance from the corresponding (x, y) points. Using an incident light field captured from the ‘real’ world would call for a pre-process stage of depth-correction [10] to re-project the incident light field onto the scene geometry used as its imposter – this is similar to the method applied in [27]. All propagation times were gathered on an Intel Core2 2.66Ghz processor in Windows XP with the single-threaded application given ‘high’ scheduling priority.

We gather radiance values on radiance-maps applied to all surfaces of the scene in a breadth-first computation. Radiance-maps comprised a total of 107,984 radiance cells for the ‘beams’ scene and 1,564,342 radiance cells for the ‘classroom’ scene. The method uses a fixed non-hierarchical spherical discretisation into solid angles $\{\omega_i\}$ to sample each emitter (S) in turn. Each ω_i is considered and $(S \oplus \omega_i)_R$ computed. Uniformly distributed directions within $-\omega_i$ are generated with low-discrepancy Halton sequences [15] and sampled from the region of R demarcated by $(S \oplus \omega_i)_R$. Rays that do not subsequently intersect S (as $(S \oplus \omega_i)_R$ is conservative) are discarded. Similar to [30], we use SSE 4-ray bundles – both with and without a BSP tree for the ‘classroom’ and ‘beams’ scenes respectively. The decision not to use a BSP tree for the latter was due to its low polygon count causing the BSP tree to form a bottleneck. Figure 5 shows the results obtained on illuminating the scenes with synthetic incident light fields. The time needed to locally illuminate the scenes from the incident light fields was near constant (at 36.85 seconds and 51.83 minutes for figures 5(a) and 5(b)), irrespective of the resolution of light-field used. This was mainly due to the coherent memory access and area based sampling of our method.

We consider a more intelligent scheme that first considers the radiance from the light field for a given (θ, ϕ) and disables radiance transfer for the corresponding sets of solid angles if the radiance is near-zero. We get significant improvement in speed depending on the complexity of the light field radiance (20.39 seconds and 10.01 minutes for figures 5(a) and 5(b) respectively), with identical results. As part of ongoing research, we are considered storing the sphere-plane parameterised light field in a hierarchical representation that can be sampled adaptively.

We now compare our method against a PDF based importance-sampling scheme and naïve uniform sampling of diffuse emitters. This scenario is perhaps the simplest for

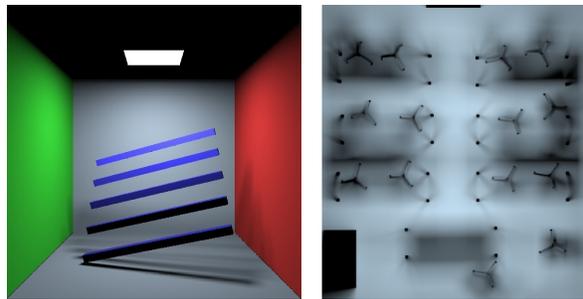


(a) Incident light field with a ‘radial’ emittance – 20.39 seconds (b) 4 incident light fields, each encoding a spotlight – 10.01 minutes

Figure 5. Illumination from incident light fields (4096 solid angles, $\times 4$ rays samples)

importance-sampling with a very easy to sample and compute PDF – and the worst case for our method. However, we are interested in seeing how our technique fares as a general sampling strategy against importance-sampling and naïve, uniform sampling. This naïve scheme selects uniformly distributed directions in a spherical triangulation of the hemisphere above a point as per [16]. For each radiance cell on R , rays are cast to find and intersect emitters, with appropriate radiance added if the ray is unoccluded to any emitter. As expected, this approach is very compute-intensive and is considered mainly as a ‘sanity-check’.

Next, we consider PDF based importance-sampling of emitter(s) S from R . Uniformly distributed points within a stratification of S are sampled from uniformly distributed points in radiance cells on R . These uniform samples are drawn from four Halton sequences [15]. Stratification of S produced only minor sampling improvement, though computation time saved when using a BSP-tree for visibility acceleration was significant (due to ray-bundle coherence [30]). In the ‘classroom’ scene, each of the light sources was sampled in turn to keep the comparison like-for-like.



(a) ‘beams’ (b) ‘classroom’ floor

Figure 6. ‘True’ results with diffuse emitters.

Finally, as before, a brute-force (non-adaptive) imple-

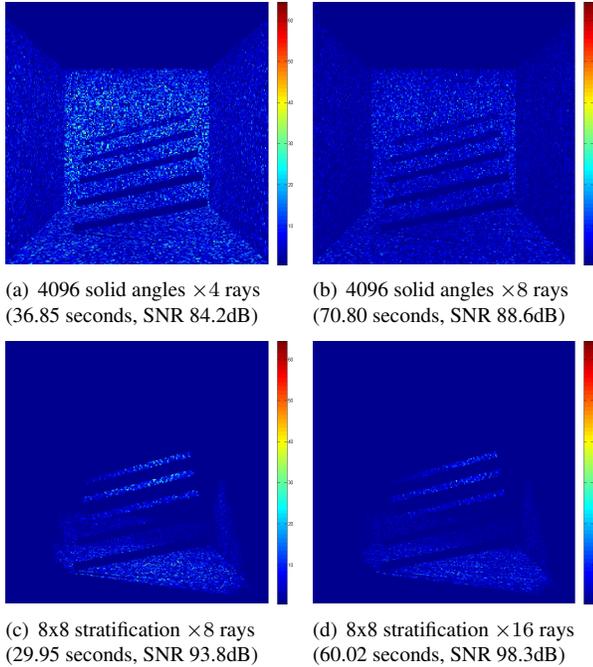


Figure 7. ‘beams’ error ($\times 10^4$) – top-row our method, bottom-row importance-sampling.

mentation of projected minkowski sums is considered. This method again uses a fixed spherical discretisation into solid angles, evaluating and sampling all solid angles sequentially as necessary. Results are compared against a ‘true’ converged solution (figure 6). Sampling parameters for the schemes were selected such that results were suitably close to the ‘true’ solutions to allow meaningful comparison, all other parameters were unchanged. Results obtained were subtracted from those of the ‘true’ solution as raw float RGB images, with the error scaled $\times 10^4$ for display to lie in a $0 \dots 65^\circ$ ‘temperature’ range (coloured navy-blue to brown) in Matlab. These error images and their signal-to-noise ratios are presented in figure 7.

Details of propagation time for different settings of importance sampling and our approach are listed in figure 7. As expected with PDF based importance sampling, surfaces unoccluded from the light sources converge faster, and have negligible variance and bias with the importance sampling scheme. However, partially occluded regions would continue to have error which may be considered as under-sampling or bias. Our method has quite different results; there is under-sampling noise on all surfaces. This is however quite low for the given computation time for a blind sampling scheme (note the scale of $\times 10^4$ applied for display purposes to the error in figure 7). Visually, albeit subjective, we could not differentiate between the four original images corresponding to figure 7. Both the schemes em-

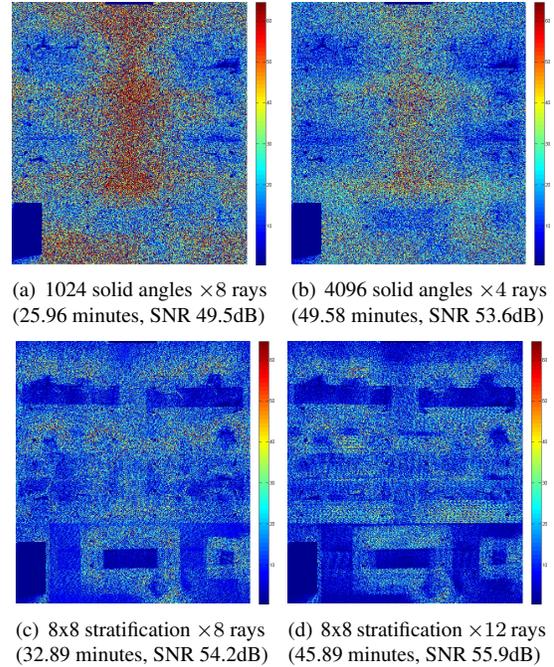


Figure 8. ‘classroom’ floor error ($\times 10^4$) – top-row our method, bottom-row importance-sampling.

ploy rejection-sampling – PDF based importance sampling rejecting rays occluded to S , while our approach also rejected rays within ω that did not terminate on S (again, as $(S \oplus \omega_i)_R$ is conservative). The naïve scheme for error close to that of figure 7(a) took approximately $115\times$ the computation time of our approach.

Given the results from the ‘beams’ scene, while gathering radiance for the whole of the ‘classroom’ scene, we decided to concentrate only on the floor of the ‘classroom’ (see figures 4(b) and 6(b)) as it presents significant variation in emitter occlusion. The propagation time from the four emitters, along with error and SNR of the floor are detailed in figure 8. We once again note a similar distribution of noise for both approaches as we had encountered in the ‘beams’ scene. For similar propagation times, our scheme had more overall under-sampling noise whereas this was concentrated in the penumbras with importance sampling. Naïve sampling equivalent to that of figure 8(b) took 1341 minutes – roughly $27\times$ the computation time of our method. We believe our approach works well for a general, non-parametric approach – with noise though slightly higher than that for PDF based importance sampling, being distributed equally throughout the radiance map rather than being concentrated within penumbras.

7 Conclusions

In this paper we have presented an alternative to parametric methods to sample a general surface. The primary aim of our research was to reconsider the flow of radiance and potential between surfaces in a novel manner, combining the processes of ‘shoot’ and ‘gather’ to develop alternative sampling strategies. As illustrated by the results in this paper, our method performs strongly in sampling radiance from incident light fields and also for diffuse emitters – albeit with increased variance. Whereas other sampling techniques may out-perform our method in specific instances, we believe our approach has advantage in the general case. Further, we do not claim our approach is applicable in all scenarios – especially since many existing methods follow a per-point depth-first particle tracing. It is possible to reconsider some of these with our approach, and we believe our method has significant potential in these techniques.

The technique can be extended to sample a large set of materials, BRDFs and image-based representations for a full global-illumination solution. As mentioned in Section 4, we are extending the approach to adaptive sampling and alternative sampling strategies. These are being addressed as part of ongoing research and current results from that work show significant promise. Our approach also offers possibilities in other areas of computer graphics such as from-region visibility culling and ambient-occlusion shading – we are investigating several of these.

Acknowledgments

This research has been funded by EPSRC grants EP/C511824/1 and GR/R13685/01. Many thanks to Jan Kautz and the anonymous reviewers for their suggestions and input on this paper.

References

- [1] S. Agarwal, R. Ramamoorthi, S. Belongie, and H. W. Jensen. Structured importance sampling of environment maps. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, pages 605–612, New York, NY, USA, 2003. ACM Press.
- [2] J. Arvo. Stratified sampling of spherical triangles. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 437–438, New York, NY, USA, 1995. ACM Press.
- [3] J. Arvo and D. Kirk. Particle transport and image synthesis. *Computer Graphics*, 24(4):53–66, 1990.
- [4] J. F. Blinn and M. E. Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19(10):542–547, 1976.
- [5] C. Buchalew and D. Fussell. Illumination networks: fast realistic rendering with general reflectance functions. In *SIGGRAPH '89: Proceedings of the 16th annual conference on Computer graphics and interactive techniques*, pages 89–98, New York, NY, USA, 1989. ACM Press.
- [6] E. Camahort, A. Leros, and D. Fussell. Uniformly sampled light fields. In G. Drettakis and N. Max, editors, *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)*, pages 117–130, New York, NY, 1998. Springer Wien.
- [7] M. de Berg, O. Schwarzkopf, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2000.
- [8] P. Dutré, H. W. Jensen, J. Arvo, K. Bala, P. Bekaert, S. Marschner, and M. Pharr. State of the art in monte carlo global illumination. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Course Notes*, New York, NY, USA, 2004. ACM Press.
- [9] C. M. Goral, K. E. Torrance, D. P. Greenberg, and B. Bataille. Modeling the interaction of light between diffuse surfaces. In *SIGGRAPH '84: Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 213–222, New York, NY, USA, 1984. ACM Press.
- [10] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54, New York, NY, USA, 1996. ACM Press.
- [11] W. Heidrich, J. Kautz, P. Slusallek, and H.-P. Seidel. Canned light sources. In *Rendering Techniques*, pages 293–300, 1998.
- [12] W. Heidrich, H. P. A. Lensch, M. Cohen, and H.-P. Seidel. Light field techniques for reflections and refractions. In *Rendering Techniques '99: Proceedings of the 10th Eurographics Workshop on Rendering (EGRW-99)*, pages 187–196, Granada, Spain, 1999. Springer.
- [13] H. W. Jensen. Global illumination using photon maps. In *Proceedings of the eurographics workshop on Rendering techniques '96*, pages 21–30, London, UK, 1996. Springer-Verlag.
- [14] J. T. Kajiya. The rendering equation. In *SIGGRAPH '86: Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, New York, NY, USA, 1986. ACM Press.
- [15] A. Keller. The fast calculation of form factors using low discrepancy sequences. In *Proceedings of the Spring Conference on Computer Graphics (SCCG '96)*, pages 195–204, Bratislava, Slovakia, 1996. Comenius University Press.
- [16] D. Kirk and J. Arvo. Unbiased sampling techniques for image synthesis. In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 153–156, New York, NY, USA, 1991. ACM Press.
- [17] E. P. Lafortune and Y. D. Willems. Using the modified phong BRDF for physically based rendering. Technical report, Katholieke Universiteit Leuven, Department of Computer Science, K.U.Leuven, 1994.
- [18] E. P. Lafortune and Y. D. Willems. Reducing the Number of Shadow Rays in Bidirectional Path Tracing. In V. Skala, editor, *Proceedings of the Winter School of Computer Graphics and CAD Systems '95*, pages 384–392, Plzen, Czech Republic, 1995. University of West Bohemia.

- [19] M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, New York, NY, USA, 1996. ACM Press.
- [20] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems. Relighting with 4d incident light fields. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, pages 613–620, New York, NY, USA, 2003. ACM Press.
- [21] S. N. Pattanaik and S. P. Mudur. The potential equation and importance in illumination computations. *Computer Graphics Forum*, 12(2):131–136, June 1993.
- [22] P. Shirley, C. Wang, and K. Zimmerman. Monte carlo techniques for direct lighting calculations. volume 15, pages 1–36, New York, NY, USA, 1996. ACM Press.
- [23] M. Slater. Constant time queries on uniformly distributed points on a hemisphere. *Journal of Graphics Tools*, 7(1):33–44, 2002.
- [24] M. Slater, J. Mortensen, P. Khanna, and I. Yu. A virtual light field approach to global illumination. In *Computer Graphics International, 2004. Proceedings*, pages 102–109, June 2004.
- [25] L. Szirmay-Kalos. Global ray-bundle tracing. Technical Report TR-186-2-98-18, Institute of Computer Graphics and Algorithms, Vienna University of Technology, July 1998.
- [26] J. Talbot, D. Cline, and P. Egbert. Importance resampling for global illumination. In *Rendering Techniques*, pages 139–146, 2005.
- [27] J. Unger, A. Wenger, T. Hawkins, A. Gardner, and P. Debevec. Capturing and rendering with incident light fields. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 141–149, Switzerland, 2003.
- [28] E. Veach and L. J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 419–428, New York, NY, USA, 1995. ACM Press.
- [29] E. Veach and L. J. Guibas. Metropolis light transport. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 65–76, New York, NY, USA, 1997. ACM Press.
- [30] I. Wald, C. Benthin, M. Wagner, and P. Slusallek. Interactive rendering with coherent ray tracing. In *Computer Graphics Forum (Proceedings of EUROGRAPHICS 2001)*, volume 20, pages 153–164, 2001.
- [31] T. Whitted. An improved illumination model for shaded display. *Communications of the ACM*, 23(6):343–349, 1980.

Appendix A Partitioning the sphere

Consider the positive octant of the canonical unit sphere. In this octant, the sphere lies entirely above the $x+y+z = 1$ plane. The equilateral triangle with vertices $v_0 = (1, 0, 0)$, $v_1 = (0, 1, 0)$ and $v_2 = (0, 0, 1)$ touches the sphere and lies on this plane (see figure A). We divide each of its edges ($\overrightarrow{v_0v_1}$, $\overrightarrow{v_0v_2}$, $\overrightarrow{v_1v_2}$) into n equal units. Consider edge $\overrightarrow{v_1v_2}$; we connect subdivisions on the other two edges ($\overrightarrow{v_0v_1}$ and $\overrightarrow{v_0v_2}$) forming line segments parallel to $\overrightarrow{v_1v_2}$. If we repeat

this process for the remaining two edges, we get a tessellation of the original triangle into n^2 triangles. Projecting vertices of this triangulation onto the surface of the sphere (by normalisation), we get a partition of the positive octant of the sphere. Each spherical triangle then represents a solid angle ω , with its centroid being its associated (central) direction Θ_ω . Individual triangles in the triangulation are assigned unique identifiers in the order depicted in figure A. To query an arbitrary direction Θ in a grid of $n \times n$ solid

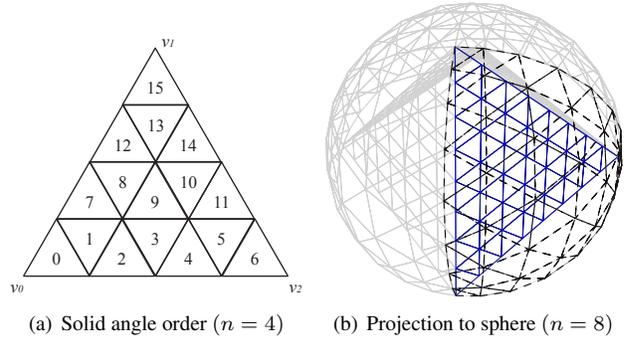


Figure A. Discretisation of the sphere into solid angles.

angles in the positive octant, we find p_Θ , the intersect of the unit vector along Θ on plane $x + y + z = 1$. We then find the intervals of $\overrightarrow{v_0v_1}$ and $\overrightarrow{v_1v_2}$ that p_Θ lies in. This is simply identifying which of the n x, z intervals p_Θ lies in. p_Θ is then limited to a parallelogram defined by these intercepts – this is the region we would have been obtained if parallel divisions of the original triangle were only constructed along $\overrightarrow{v_1v_2}$ and $\overrightarrow{v_0v_1}$. Isolating which triangle p_Θ belongs to requires testing if it lies above or below the diagonal of this parallelogram. Θ is thus associated to ξ_Θ , a unique solid angle identifier – see (Pseudocode A) below.

```

 $p_\Theta = \text{intersect of } \Theta \text{ on } x + y + z = 1$  (Pseudocode A)
 $x = n p_{\Theta_x}, z = n p_{\Theta_z}$  // scale from  $[0, 1]$  to  $[0, n]$ 
 $x_i = \lfloor x \rfloor, z_i = \lfloor z \rfloor$  // get floor() of  $x, z$  intercepts
 $\xi_\Theta = z_i(2n - z_i) + 2x_i$  // parallelogram cell  $x, z$ 
 $diag = (x - x_i)(z - z_i)$  // parallelogram diagonal offset
if ( $diag > 1$ )  $\xi_\Theta++$  // increment if above diagonal

```

The discretisation of the positive octant is mirrored across to the other octants. This produces $8n^2$ solid angles over the sphere – the centroid of each solid angle is stored as its representative direction. By appropriately mapping an arbitrary direction Θ to the positive octant along with a suitable offset for ξ_Θ , we use the above lookup to determine which solid angle it lies in within constant time without any additional lookup.