Concrete Computation of Global Illumination Using Structured Sampling

George Drettakis
Eugene Fiume

Department of Computer Science
Unversity of Toronto
{dret|elf}@dgp.toronto.edu

ABSTRACT

A new methodology is presented for the computation of global illumination using structured sampling. Analytical/numerical solutions for illumination are developed for simple lighting configurations. These solutions are subsequently used to generate accurate reference images. The structured sampling solution for global illumination is then discussed, comprising sample placement for illumination calculation, reconstruction for light transfer and finally resampling and filtering of illumination samples for display. A first approximation to this technique is presented using a priori placement of samples, irregular polygon reflectors, grid resampling and a conical filter for display. The new algorithm is evaluated for image quality, and compared to the traditional radiosity-based approach. These first results show that the structured sampling solution yields significant computational savings while maintaining high image quality.

1. Goals of the Approach

The calculation of global illumination is inherently complex, even for environments that are simple in terms of geometry and reflectance properties. The physical laws governing illumination are described by integrals that can become extremely difficult to compute depending on geometrical configuration, light source type and surface properties. For reflection and refraction, the solution of a complex system of high-dimensional integral equations is required.

In computer graphics global illumination algorithms, a simulation of lighting phenomena is required so that visually satisfying realistic images can be synthesized. To achieve this goal, several simplifying assumptions have been made in previous work. The most common (e.g. [Zatz93, ChFe90]) is the subdivision of the environment into small diffuse-reflector pieces, in which constant luminosity (or radiosity) is assumed. This subdivision is done ‘‘blindly’’, or with adaptive methods that refine the level of subdivision based on the output of the already discretised environment [NiNa85]. The actual behaviour of the light-transfer functions, even for the simple special cases considered, is not studied, and not incorporated into the subdivision criteria. In addition the structure of the scene geometry, the viewing parameters, light source properties etc. are also ignored.

In most previous approaches, light is deposited onto samples of equal luminosity, in the form of luminosity values. To simulate the light transfer process, a subsequent *reconstruction* step is required to retrieve the light function from the segments of equal luminosity (elements). The most commonly used reconstruction is to shoot from an emitter (patch) using a straight average ([ChFe92, GrCT86]) of the samples on which illumination is placed (elements). Even though some work has been done to achieve intelligent reconstruction (e.g. [(null)]) again the adaptivity criteria are based on the discrete samples and not on light function behaviour or scene properties. In addition, most calculations are performed in object space without consideration of the viewing and imaging parameters.

In general, computing an exact solution for any environment of moderate complexity is not feasible. To be able to assess image quality of global illumination approximations, ‘‘converged’’ images are usually taken as the ‘‘correct’’ solution, and used as a quality measure. These images can sometimes contain significant errors, and therefore are not always reliable as a quality measure. To allow for more concrete measures of comparison, simpler environments are considered in which a combination of analytical solution and numerical quadrature can be used, to compute *exactly correct* luminance values, within reasonable time constraints. By carefully examining these simple cases, more concrete understanding of luminance function behaviours can be achieved. This understanding will be used in global illumination approximations as shall be shown later.

The ideas behind structured, or structure-directed, sampling were presented by the authors in [DrFi91]. An implementation of parts of the process will be presented here. In essence, the solution utilises the properties of illumination functions to guide sampling for display and light transfer. The experiments show that the new solution results in high-quality images at a lower cost than the corresponding radiosity-based approaches. The simple environments used allow careful comparisons with the exact solution, giving significant insight into the sampling rates used in the admittedly complex calculations of the radiosity-based approaches. A test containing a moderately complex environment is presented, indicating that the algorithm presented can be applied to real environments with significant benefit.

2. Analysis of Simple Diffuse Environments

To allow the computation of exact images for simple configurations it is necessary to formulate the equations for direct illumination and ‘‘one-bounce’’ reflection.

2.1. Direct Illumination

Direct or primary illumination is simpler to study since it involves only the calculation of the luminance due to light originating from the light source, reflecting off a visible surface and finally reaching to the eye. The *luminance* $l'(\bar{r}_1)$ in general direction $\bar{r}_1$ is the power leaving a surface per area, per solid angle. All subsequent equations are given in terms of luminance.

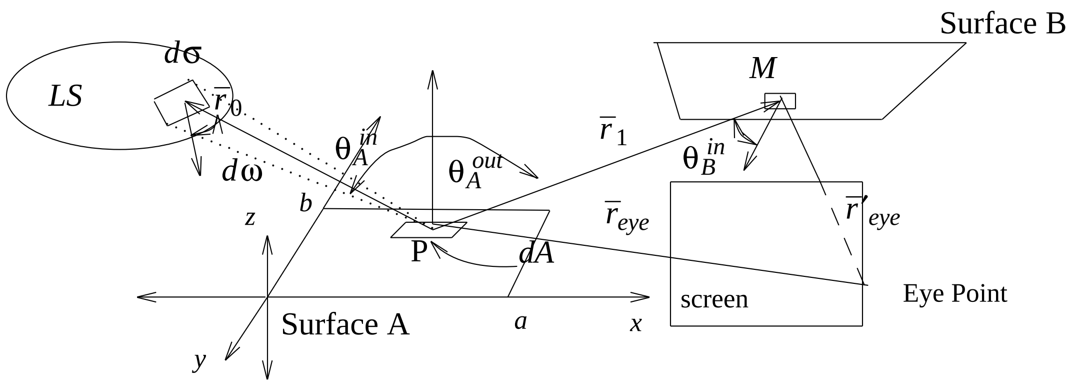

Figure 1

For the configuration shown in Figure 1, where light source LS has luminance $l_{LS}(\bar{r}_0)$, and the surface S has reflectivity $\rho_S''(\bar{r}_0, \bar{r}_1)$, the luminance in direction $\bar{r}_{eye}$, from the surface differential area dA, due to the light source LS, is given:

$$l_A'(\bar{r}_{eye}) = \int_{LS} l_{LS}'(\bar{r}_0) \rho_A''(\bar{r}_0, \bar{r}_1) \cos\theta_A^{in} \, d\omega \tag{1}$$

where $d\omega$ is as shown in Figure 1. This is a general equation that makes no limiting assumptions about the nature of the reflectance (encapsulated in $\rho_A''(\bar{r}_0, \bar{r}_1)$) or the light source behaviour (described by $l_{LS}'(\bar{r}_0)$). For a perfectly diffuse light source LS and a perfectly diffuse reflector S, source luminance and surface reflectivity become simple constants l_{LS} and ρ_S, and if R_0 is the distance between dA and $d\sigma$, Eq. (1) simplifies to:

$$l_A'(\bar{r}_{eye}) = l_{LS} \frac{\rho_A}{\pi} \int_{LS} \cos\theta_A^{in} \, d\omega = l_{LS} \, \rho_A \int_{LS} \frac{\cos\theta_A^{in} \cos\theta_A^{out}}{\pi R_0^2} d\sigma$$

$$= l_{LS} \, \rho_A \, F_{dA-LS} \tag{2}$$

The quantity F_{dA-LS} is termed the *form-factor* between the differential surface element dA and the source LS.

2.1.1. Solving for Simple Light Source Geometries

In the general case, Eq. (2) is a complicated function that does not usually have an closed form solution. This complexity is due to the inherent difficulty of finding the domain of integration, and applying a suitable transformation to allow computation of the cosines. For certain specific geometries and types of light sources, analytical solutions to F_{dA-LS} exist (see [Moon36, SiHo72, Dret91]).

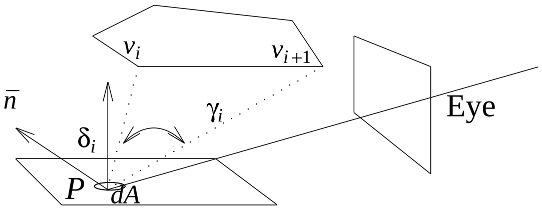

Figure 2

One such example is an arbitrary planar polygonal light source with n sides. For the polygon shown in Figure 2, the analytical solution for F_{dA-LS} is:

$$F_{dA-LS} = \frac{1}{2} \sum_{i=0}^{n} \gamma_i \cos\delta_i \tag{3}$$

The angle γ_i is defined by the lines connecting the point P and the vertices v_i, v_{i+1} of the polygon, and δ_i is the angle between the normal $\bar{n}$ to the plane defined by these lines and the normal at P, at the centre of dA. Because the derivation is based on Stokes theorem, the only limitation is that the polygonal light sources have no holes.

2.2. Mathematical Formulation of One-Bounce Reflection

The configuration now considered is that of a single, constant luminosity, diffuse light source LS, placed on a plane parallel and directly opposing receiver A, and a parallel receiver B that receives no light directly from LS (see again Figure 1). From the previous section, and ignoring the interaction from B to A, the luminance on visible points on B will be determined.

One way of viewing this calculation for each visible point on B is as if the receiver A is a light source. In this case the value of the luminance function $l_A(\bar{r}_1)$ *varies* over the surface A, in the manner described by Eq. (2). The luminance $l_B(\bar{r}'_{eye})$ at point M on surface B is given by the following expression (note that R_1, the distance between dA and M, is a function of the position of dA):

$$l_B(\bar{r}'_{eye}) = \rho_B \int_A l_A(\bar{r}_1) \frac{\cos\theta_B^{in} \cos\theta_A^{out}}{\pi R_1^2}\, dA \tag{4}$$

Using Eq. (2), Eq. (4) becomes:

$$l_B(\bar{r}'_{eye}) = l_S\ \rho_B\ \rho_A\ \frac{1}{\pi^2} \int_A F_{dA-LS}\ \frac{\cos\theta_B^{in} \cos\theta_A^{out}}{R_1^2}\, dA$$

The above equation entails a single integration over the surface of receiver A, assuming that there is an analytical expression for the term F_{dA-LS}. To give an indication of the complexity involved for such a computation, Eq. (5) is the final expression for a rectangular light source. Briefly, $\bar{n}_A, \bar{b}_B, \bar{n}_s$ are the normals at A, B and the light source, $\bar{r}_{A-B}, \bar{r}_{B-A}$ are unit vectors between the visible point on B and the point (x,y) on A, and $\bar{s}_i$ is a vector from vertex i of the source to (x,y). For details and the derivation, see [Dret91]. The luminance at a visible point on B is given as follows:

$$l_B(\bar{r}'_{eye}) = l_S\ \frac{\rho_B\ \rho_A}{\pi^2} \int_{y=0}^{b} \int_{x=0}^{a} \frac{(\bar{r}_{B-A} \cdot \bar{n}_B)\,(\bar{r}_{A-B} \cdot \bar{n}_A)}{(x_B - x)^2 + (y_B - y)^2 + (z_B)^2} \sum_{i=0}^{3} \cos^{-1}(\,\bar{s}_i \cdot \bar{s}_{i+1}) \left[(\bar{s}_i \times \bar{s}_{i+1}) \cdot \bar{n}_S \right] dxdy \tag{5}$$

2.3. Accurate Numerical Solution

To allow the use of a numerical solution for the computation, the integral in Eq. (5) is given to an adaptive quadrature routine, written in FORTRAN, taken from the netlib library. Even though it is a complicated function, it is tractable for medium resolution images†.

Test scenes have the configuration of Figure 1. The renderer applies the direct illumination equations (Eq. (2) and (3)) for the visible pixels of A, resulting in exact analytical values. For the tertiary reflector B the scan-converter first finds the appropriate light-source and the appropriate secondary reflector. All values are transformed into the suitable space so that Eq. (5) can be applied, and the numerical integrator (a Gaussian quadrature) is invoked to any desired accuracy. This calculation is performed for each visible pixel on B.

An example of an image computed with this method is shown in Image 1, which corresponds to the environment presented in Figure 1. From now on this will be referred to the ‘‘simple’’ environment.

3. Structured Sampling: a First Approach

As discussed in detail in [DrFi91] the basis of *structured sampling* is the recognition of the properties (structure) of light in a scene and their subsequent utilisation in computing illumination.

---

† An average 200x200 resolution image takes 10 minutes to compute. All computation times reported are on an SGI Iris 4D/35 workstation.

3.1. Fundamental Concepts

The basic tenet of the approach is that in diffuse environments, luminance functions across surfaces are very smooth (with the exception of shadow boundaries), and therefore the number of samples required to reconstruct the luminance is small. Taking this assumption as a starting point, there are three main actions that need to be performed to calculate illumination in a scene and generate a realistic synthesized image.

Selection and tagging of illumination samples. Illumination samples are associated with the original given geometry. Three things are worth noting here. First, scene structure, such as geometry variation, shadow boundaries and viewing parameters can be used to determine the density and positioning of these samples. Second, this sampling can change adaptively as the requirements of illumination become more apparent. Third, the samples are directly related to the resampling for display described below. This ties in with the concept of ‘‘tagged samples’’ introduced in the earlier paper [DrFi91]. A *tagged* sample contains information indicating whether it is in shadow, or it is in an area with smoothly varying illumination, or it was placed due to viewing parameter constraints etc.

Reconstruction for light-transfer. Some of the samples that have been previously selected are used to propagate light. A *shooting paradigm* is used, similar to that of the progressive refinement radiosity-based technique [ChSD93]. Depending on the requirements and variation of the illumination functions, varying degrees and quality of interpolation can be used to achieve the desired result. If for example a collection of samples have extremely smooth variation, and the geometry is suitable, simple averaging may be sufficient. For more complex variation, linear or higher order approximations may be required.

Resampling and filtering for display. The samples that have been selected contain values, exact or approximate, of the luminance at a certain point in space. The original geometric primitives must then be rendered using this information. In the radiosity-based algorithms the radiosity is computed at vertices [GrCT86] or at centres [Zatz93] of polygons and therefore there is a direct relationship between illumination samples and geometry. For structured sampling, the samples are chosen to match the variation of illumination, they are irregular and sparse and do not correspond directly to geometry. To render the primitives, a *resampling* is required, mapping the illumination back to the original geometric primitives. This method is especially beneficial when graphics hardware is used and therefore displaying numerous primitives is cheap compared to the illumination calculations.

The essence of the approach is to direct sampling and reconstruction based on the study and analysis of illumination function behaviour. Using isolux contours to guide sample placement and light transfer reconstruction, sampling rates are significantly reduced. In addition, using resampling and filtering for display enhances image quality without adding light-transfer calculations. Further work on selecting appropriate filters will also be guided by the study of light function variation.

3.2. A First Approach to Structured Sampling

In what follows a first structured sampling solution is presented. This is not yet a complete new algorithm: the main purpose of this presentation is to discuss the encouraging first results of experiments. The use of controlled and simple environments facilitates careful evaluation of the proposed solutions.

Sample selection. In this paper samples are placed in a preprocessing step by the user, using an interactive *sample editor* built into a global illumination testbed. Even though automatic placement has not yet been achieved, the guidelines used to interactively place samples will be presented, since they will become the basis for such an algorithm.

The illumination samples associated with a surface can be of two kinds. *Display samples* are used only for display, and are collectors of illumination. *Light transfer samples* also serve as collectors, but are grouped with several other samples to form secondary emitters. The purpose of the grouping is to form surfaces that correspond to areas of equal or almost equal luminosity. The placement of these samples is based on the properties of the light sources and shadow regions.

To assist in this selection lower resolution colour images can be used to identify the areas with almost equal illumination. Consider the environment shown in Image 1, that corresponds to Figure 1. This image has been computed using the analytical/numerical technique previously described. Figures 3(a) and 3(b) show the variation of illumination with 8-bit colour resolution instead of 24-bit, across surfaces *A* and *B*. In each of these Figures the viewpoint is exactly above the centre of each surface. The regions of almost equal illumination (‘‘isolux’’ from now on) can easily be seen.

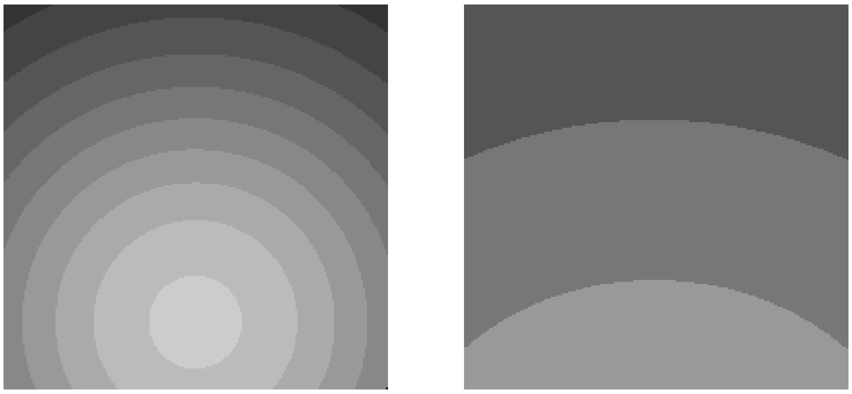

(a) **(b)**

Figure 3: Isolux Contouring

Samples are now placed to correspond with the curvature of these contours. Depending on the requirements for further propagation, some of these samples are connected together to form (usually coarse) approximations of the ‘‘isolux’’ areas. Display samples can be selected in a ‘‘view-dependent’’ manner, so that the regions that are larger in the final image have higher sample density. The samples selected and the secondary emitter polygons created for the scene shown in Image 1, are shown in Fig 4(a). These are coarse approximations of the contours in Figures 3(a) and 3(b).

Shadow boundaries are treated in a similar manner. Samples are placed across the umbra and penumbra regions with varying density. The emitters need not contain many samples, especially in the case of the shadowed regions, since these do not contribute much to scene illumination. Figure 5 shows structured samples corresponding to Image 4.

Even though this selection process assumes knowledge that does not exist prior to a global illumination solution, it offers significant insight into how automatic sample placement algorithms will be created. By carefully studying common light-source types and configuration, isolux contours can be computed exactly or approximately. Shadow boundary calculations, at varying levels of granularity, can also be used. As mentioned previously, samples will maintain information about the reason for their selection, becoming intelligent or ‘‘tagged’’ samples. This information can be used both for light-transfer as well as resampling and filtering.

Light-transfer reconstruction. As mentioned before, some samples have been collected into secondary emitter polygons representing ‘‘isolux’’ regions (these are the polygons shown in Figure 4(a) and Figure 5). For *direct* illumination, light is shot from the light sources to the samples, and exact luminance values are stored using the analytic formulas (Eq. (2) and (3)). Thus *exact*

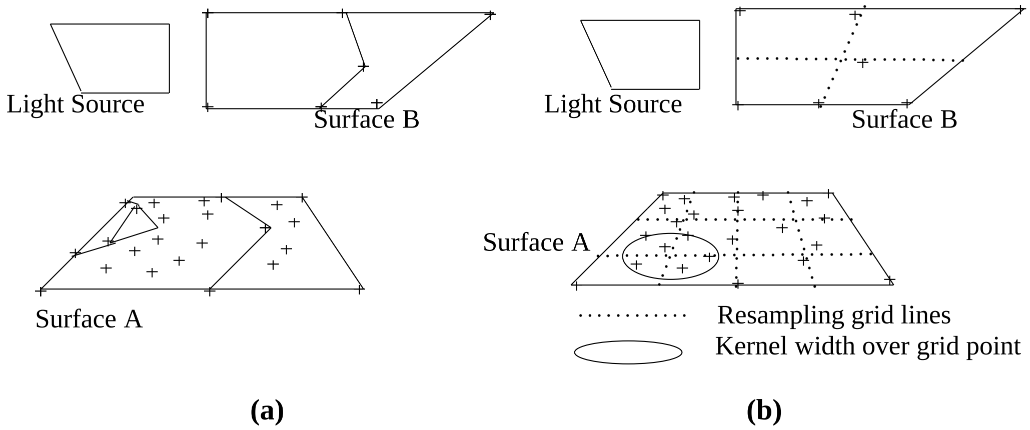

Figure 4: Structured Sampling Samples and Resampling Grid

samples of the luminance function defined over a surface are stored at points carefully selected to correspond well to the illumination variation. For *secondary light*, the values at the vertices of each polygon are averaged, and the light is shot from the polygon corresponding to the region to the display samples, again using Eqs. (2) and (3). The shape of these polygons is based on the isolux contours. As a result this approximation is much more accurate than the regular patch subdivision, since the constant luminosity polygons correspond closely to regions that have almost constant luminosity as a result of the light function properties.

For more complex environments, shooting will continue until unshot light power has dropped below a preset threshold. At each step, the polygon with the highest power is selected. In future work, occlusion will be computed during sample placement (since shadowed regioned will have been identified), and samples will carry a tag indicating their occluded or visible status. For now, a simple ray-casting scheme is used, for demonstration purposes. When a secondary emitter polygon with vertices $v_{0,} \cdots, v_n$ shoots light, a ray is cast from each v_i to every display sample in the environment, and the result at each display sample is weighted based on the ray intersection count.

Resampling and filtering for display. To incorporate the information contained in the display and light transfer samples for the final image generation, a grid resampling technique is used. A grid is placed over the surface to be rendered that contains the samples, the sparse irregular samples are then filtered to the grid vertices, and the polygons subsequently rendered. An example grid is shown in Figure 4(b), used for the samples of Figure 4(a).

To assign a luminance value to each vertex on the grid, the sparse irregular samples within a certain radius r of each grid point are filtered using a simple conical filter. Assume that the grid point is at (u,v), and there are n samples $S_i, i=1, \cdots n$ at corresponding positions (u_i, v_i), and $d_i = \sqrt{(u-u_{i)}^2+(v-v_i)^2} < r$ is the distance from the grid point to each sample S_i. If the luminance at each sample S_i is L_i then the luminance value, L_{uv}, at grid point (u,v) is given as follows:

$$L_{uv} = \frac{\sum_{i=0}^{n} (1-\frac{d_i}{r}) L_i}{\sum_{i=0}^{n} d_i}$$

It is evident that when ‘‘tagged samples’’ are used, different weights or kernels can be used, depending on whether the samples represent discontinuities such as shadows, or areas in which smoothing is appropriate. For example if a set of samples are tagged as shadow boundaries, non-linear or adaptive filters can be used to produce the desired effect. The current conical filter is a starting point, simply showing the feasibility of the resampling approach. Better filters, with radial characteristics are being designed that are better adapted to the properties of isolux contours. In particular, the shadow boundaries resulting from the conical filter are not of very high quality, since the filtering results in excessive blurring.

The algorithm just described is not a full solution. However, the implementation and experiments presented in the following section show very encouraging first results that confirm the choices made.

4. Experiments and Comparative Analysis

The goal of the experiments is to determine the quality of the new algorithm and show that the sampling rates required are significantly reduced. In addition a comparison with traditional radiosity-based solutions is performed.

Image 1: Simple Environment Computed with Analytical/Numerical

Image 2: Structured Solution

Image 3: Radiosity-based Solution

Simple Case Scene Parameters			
Resolution	200x200	Light Source Area	9
Aspect Ratio	2.2	Light Source Power (r=g=b)	12
Offset of LS	1.2	Area of A, B	25

Table 1: Simple Case Scene Parameters

4.1. Simple Case Experiments

The first experiments are of the environment previously described in Section 2.3. The parameters of the scene are listed in Table 1. In Image 1, the exact solution was shown, calculated with the analytical/numerical solution. The structured solution was computed using the structured samples shown in Figure 4(a), based on the isolux contouring in Figure 3(a) and (b). The structured solution image is shown in Image 2. Image 3 shows the image computed using the ray-traced form-factor solution (Scene 1 in Table 2). Results from four runs of the ray-traced form-factor solution (Scenes 1-4) are presented in which different parameters are varied.

Image quality can be measured in a number of ways, and the error metric can be weighted to be sensitive to different types of errors. The metric used here is an pixel-by-pixel absolute difference of the the approximate image from the reference image. Only pixels corresponding to visible points are taken into consideration. The percentage of pixels with an absolute difference greater than a certain tolerance τ is taken as the error of the approximation image.

Structured Sampling		Ray-Traced Form-Factor Solution				
			Scene 1	Scene 2	Scene 3	Scene 4
Polys (A)	4	# of Patches (A)	4	4	4	4
Samples on A	48	Total Samps. (A)	100	64	64	64
Samples (B)	10	# Elem. Vert. (B)	36	36	16	9
		Samps. on Source	400	400	400	400
		Samps./Patch (A)	16	16	16	16
% > 10% err (B)	0.02		0.00	0.02	0.92	2.92
% > 10% err (A)	1.52		0.53	2.54	2.54	2.54
Time (LS $\rightarrow$ A)	0.02		0.97	0.69	0.69	0.69
Time (A $\rightarrow$ B)	0.01		0.06	0.05	0.04	0.03
Time Total (sec)	0.03		1.03	0.74	0.88	0.68

Table 2: Results for Rectangular Light Source

In Table 2, the results of the experiments in quality and speed are listed, in terms of the relative sampling rates. These results are indicators of general behaviour of diffuse functions and lighting situations. The controlled environment, and the benefit of an analytical solution allows trustworthy and concrete conclusions to be reached. The structured solution results in three main benefits:

Reduction in the number of display samples. To achieve less than 2% error on surface *A*, (τ = 10%), 48 structured samples are required for the structured solution, while a total of 100 element vertices are needed for the radiosity-based approach (Scene 1). If the ray-traced solution uses only 64 elements, the error is somewhat larger (2.54% in Scene 2), than the corresponding structured approach result. This reduction in required number of samples is due to the fact that display samples are placed *based on variation of illumination*, following the isolux contours.

Significant speedup. As a result of the reduction in the required number of samples, the structured solution outperforms the radiosity-based approach in terms of speed. The secondary reflection is 5-6 faster than the corresponding radiosity-based solutions, while the total varies from 24-34 times, although this is dependent on the sampling rate of the light source†.

Reduction in samples required for light transfer. The use of isolux region approximations as constant luminosity polygons, allows equivalent quality light-transfer to be performed at a lower cost. The total required number of samples on surface *B* for less than 1% error, is 10 for the structured solution. Keeping all other parameters constant, in Scene 3, the radiosity-based solution requires 16 element vertices (samples) on *B* to achieve the same quality. If this is reduced to 9 element vertices the error reaches 2.2% (Scene 4).

4.2. Moderate-Complexity Scenes

As discussed above, simple scenes give valuable insight due to their simplicity and the existence of exact solutions. However, it is also beneficial to examine moderate complexity scenes to evaluate possible performance gains from new solutions. The advantage of having the exact solution is lost yielding results that are less concrete, since numerous other parameters have been introduced. They do however give a sense of feasibility of the structured method.

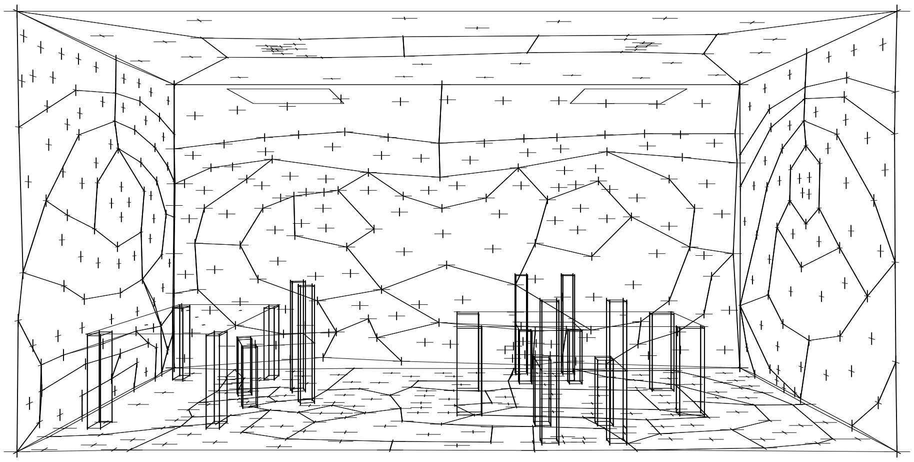

Figure 5: Structured Samples, Secondary Emitters for Moderate Geometry Scene

Such a scene (computed using the parameters listed in Table 3 as ‘‘Converged’’, with the ray-traced form-factor solution), is shown in Image 4. This image is used for evaluation purposes. A relatively good quality structured solution is shown in Image 5, while the structured samples are shown in Figure 5. A image that has similar computational requirements using the ray-traced form-factors is shown in Image 6. The results for this run are listed in Table 3 in the ‘‘Same speed’’ column. A second ray-traced form-factor solution is run that results in equivalent image quality, listed in the ‘‘Same quality’’ column. The error tolerance τ for all tests was set to 15%.

	Structured Approach	Same speed	Same quality	Converged
Error (>15%)	23.00%	41.48%	23.30%	-
Computation Time	02:30	03:20	08:11	1:10:00
# Samples (floor)	194	225	400	4225
# Samples (walls)	493	720	900	12500

Table 3: Moderate Complexity Scene Results

† In this particular unoccluded case, the exact solution could also have been used for the light source shot in the ray-traced form-factor solution.

Table 3 shows that the ray-traced form-factor solution requiring the same computation time results in approximately double error rate. The illumination levels in Image 5 for the structured solution are much closer to those in the converged solution (Image 4), than the radiosity-based solution (Image 6). This is particularly evident on the ceiling. The ray-traced form factor solution that achieves the same quality requires almost quadruple computation time, and significantly more sample points on walls and floor. In addition, despite the lower number of samples, element gridding can be seen in the radiosity-based solution, while the filtering in the structured solution results in a much smoother overall image.

These results must be taken only as a first indication, since a full structured solution will not require ray-casting for occlusion, and will incorporate different filters for display. When completed, much improved results are expected from the structured approach.

5. Summary, Conclusions and Future Prospects.

In this paper a first concrete approach to structured sampling is presented. The use of simple environments in which exact solutions can be used for evaluation purposes is advocated. Such cases are presented and exact solutions derived, using analytical solutions for light sources and numerical quadrature for reflection. The structured sampling concepts are presented and a first partial implementation of the methods proposed is described. Experiments involving simple environments that allow careful evaluation, and moderately complex scenes that indicate feasibility are presented, with images and quality/computation time statistics.

The first results presented here are very encouraging. Computation time and sample requirements are significantly reduced. Careful evaluation of simple environments show the importance of the following new techniques: *separation of samples* into those used for display and those used for light-transfer, *use of better reconstruction* for light transfer, in this case by using isolux region and shadow boundary information and *resampling and filtering for display* that allows some of the cost of generating high-quality images to be performed by image-processing techniques. In addition, the moderate complexity environments, even without a full occlusion solution, show encouraging sampling, quality and speed benefits.

The possibilities for future work are numerous. For *sample placement*, shadow boundary techniques and corresponding sample ‘‘tagging’’ as well as automatic calculation of ‘‘isolux’’ contours are currently being investigated. The use of view-specific information (for display samples), and adaptive techniques for secondary lighting will also be investigated. For *light-transfer reconstruction*, further research will include evaluation of the quality required for reconstruction. For *resampling and filtering for display* the emphasis will be on design and evaluation of appropriate filters and sampling rates, as well as use of other filter types when displaying tagged samples. Occlusion is an important issue and involves all three stages. Adaptive techniques must be used, and the use of different reconstruction in light-transfer may alleviate part of the problem. Finally the use of special filters for display to enhance desired effects can be used to in the place of accurate occlusion calculation.

Acknowledgements

The authors gratefully acknowledge the financial support of National Science and Engineering Research Council of Canada, the Information Technology Research Centre of Ontario, and the University of Toronto, for this research and for travel to the workshop. Thanks also to Pierre Poulin for fruitful discussions, Tom Milligan for shooting the final pictures and Rob Lansdale for providing the software base, on which most of the global illumination testbed was built.

Image 4: Moderate Complexity Converged Image

Image 5: Structured Sampling Solution for Moderate Complexity Scene

Image 6: Radiosity-Based Solution for Moderate Complexity Scene

References

ChFe92.
Chin, Norman and Steven Feiner, ‘‘Fact Object Precision Shadow Generation for Area Light Source using BSP Trees,’’ *ACM Computer Graphics (SIGGRAPH Symp. on Interactive 3D Graphics).*

ChFe90.
Chin, Norman and Steven Feiner, ‘‘Near Real-Time Shadow Generation for Global Diffuse Illumination,’’ *ACM Computer Graphics (SIGGRAPH ’90 Proceedings)*, vol. 24, no. 4, pp. 99-106.

ChSD93.
Christensen, Per H., David H. Salesin, and Tony D. DeRose, and Larry Aupperle, ‘‘A Continuous Adjoint Formulation for Radiance Transport,’’ *4th Eurographics Workshop on Rendering*, Paris, France June 1993.

DrFi91.
Drettakis, George and Eugene L. Fiume, ‘‘Structured-directed Sampling, Reconstruction and Data Representation for Global Illumination,’’ *2nd Eurographics Workshop on Rendering*, Barcelona, Spain May 1991.

Dret91.
Drettakis, George, ‘‘A Framework for Direct Lighting Calculations,’’ *DGP Technical Memo*, Dept. of Computer Science, Un. of Toronto, November 1991.

GrCT86.
Greenberg, Donald P., Michael F. Cohen, and Kenneth E. Torrance, *Radiosity: A Method for Computing Global Illumination*, 18, July 1984.

Moon36.
Moon, Parry, ‘‘*The Scientific Basis of Illuminating Engineering*’’, McGraw-Hill, 1936.

NiNa85.
Nishita, T. and E. Nakamae, *Continuous Tome Representations of Three-Dimensional Objects Taking Account of Shadows and Interreflections*.

SiHo72.
Siegel, Robert and John R. Howell, ‘‘*Thermal Radiation Heat Transfer*’’, McGraw-Hill Book Co, 1972.

(null).Torrance, Kenneth E., Donald P. Greenberg, and Bennet Battaile, *Modelling The Interaction Of Light Between Diffuse Surfaces*.

Zatz93.
Zatz, H. R., ‘‘Galerkin Radiosity: A Higher Order Solution Method for Global Illumination,’’ *ACM Computer Graphics (SIGGRAPH ’93 Proceedings)*, August 1993.