

Adaptive Splatting for Specular to Diffuse Light Transport

Steven Collins*

Trinity College Dublin, Ireland.

1. Abstract

We present an extension to existing techniques to provide for more accurate resolution of specular to diffuse transfer within a global illumination framework. In particular this new model is adaptive with a view to capturing high frequency phenomena such as caustic curves in sharp detail and yet allowing for low frequency detail without compromising noise levels and aliasing artefacts.

A 2-pass ray-tracing algorithm is used, with an adaptive light-pass followed by a standard eye-pass. During the light-pass, rays are traced from the light sources (essentially sampling the wavefront radiating from the sources), each carrying a fraction of the total power per wavelength of the source. The interactions of these rays with diffuse surfaces are recorded in ‘illumination-maps’, as first proposed by Arvo [Arvo86]. The key to reconstructing the intensity gradients due to this light-pass lies in the construction of the illumination maps. We record the power carried by the ray as a ‘splat’ of energy flux, deposited on the surface using a gaussian distribution kernel. The kernel of the splat is adaptively scaled according to an estimation of the wavefront divergence or convergence, thus resolving sharp intensity gradients in regions of high wavefront convergence and smooth gradients in areas of divergence.

The 2nd pass eye-trace modulates the surface’s radiance according to the power stored in the illumination map in order to include the specular to diffuse light modelled during the first pass.

2. Introduction

Ultimately, to produce realistic renderings of a scene, we wish to determine the radiance at every point in the scene visible to the viewer. The radiance at any point in the scene may be expressed using a recursive integral:

$$L(x, \theta, \phi, \lambda) = L_c(x, \theta, \phi, \lambda) + \int_{\Omega} \rho_{bd}(\theta, \phi, \theta', \phi', \lambda) L(x, \theta', \phi', \lambda) \cos\theta d\omega \quad (2.1)$$

* The author can be reached at Steven.Collins@cs.tcd.ie. HTTP reference: <http://vangogh.cs.tcd.ie/scollins/scollins.html>.

$L(x, \theta, \phi, \lambda)$ is the radiance (energy per unit time per unit projected area per unit solid angle per wavelength) leaving point x in direction (θ, ϕ) and is defined in terms of the radiance emitted by the surface in that direction, $L_e(x, \theta, \phi, \lambda)$ and the radiance impinging on point x from all other directions (θ', ϕ') scaled by the *bi-directional reflectance distribution function* (BRDF) which is a function of the incoming and outgoing angles and the wavelength. Existing algorithms, ray tracing and radiosity solve for many of the modes of light transport but not all. The most notable omission is that of specular to diffuse light transport, or light that reaches a surface through one or more specular interactions (either reflection or refraction).

Classic ray-tracing simulates only the **LDS*E** | **LS*E** paths (using Heckbert's notation [Heckbert90]) whereas the radiosity technique simulates **LD*E** (although extensions to the radiosity techniques to include specular transfer [Immel86][Wallace87] have extended this path classification). The technique presented here implements a 2-pass method first proposed by Arvo [Arvo86] and later augmented by Heckbert [Heckbert90] to model the **LS+DE** transport as a first pass, tracing rays from the light sources as they interact with specular surfaces and depositing power on these surfaces if they exhibit diffuse characteristics and terminating at a purely diffuse surface (or when the power falls below a pre-defined threshold). A subsequent eye-pass backwards traces paths of the form **DS*E** and extracts the power deposited by the first pass onto the diffuse surfaces.

3. Tracing LS*D paths

The task of tracing these paths is a 3-fold operation:

1. Trace rays from the light sources, assigning a fraction of the total light source power Φ_s to each ray.
2. Record the rays collisions with diffuse surfaces by depositing this power, Φ_r , on the surface.
3. Reconstruct the power distribution during the eye-pass by modulating the radiance from the surface according to the power stored on the surface.

We will now discuss the details of each of these phases and their implementation in the current system.

3.1 Light Ray Tracing

We trace rays originating at the light sources in the scene and record the interactions of these rays and the surfaces in the scene, depositing power on diffuse surfaces along the rays' paths. In order to record this power, illumination maps in the spirit of Arvo [Arvo86] are used.

3.1.1 Illumination Map Resolution Determination

We chose not to use adaptive illumination maps such as the approach adopted in [Heckbert90] as we wish to use these maps directly as texture maps in the visualisation phase and due to the large memory requirements of high resolution quadtrees. We must however determine the resolution of the illumination map on a per surface basis. Heckbert proposes tracing rays to the surfaces and recording the distance between these rays with the surface, using this information in the later light-pass in order to determine the threshold for quadtree subdivision during illumination map adaptation. Chen et al. [Chen91] propose a similar scheme, where the distance between intersections during this initial *size pass* is used to specify the resolution for uniform caustic maps. However, this gives a view dependent solution and in general is not appropriate for a general pre-process, the results of which we wish to use subsequently in a walk through visualisation system.

Initially, illumination maps are not assigned to surfaces, and only assigned to surfaces that are actually hit by light pass rays. Thus surfaces which are not indirectly irradiated via specular transfers are not encumbered with expensive maps. An estimate of convergence/divergence of the wavefront emanating from the light source is used to determine the resolution of the map. This estimate is also used to determine the distribution of power deposited on the surface's illumination map and is outlined in section 3.2.1. The surface keeps track of the average behaviour of the wavefront as it interacts with the surface and increases the resolution of the map if the power distribution exhibits high frequencies. The illumination maps are increased in resolution in integer quantities to simplify the power redistribution. A map with an original resolution of (r_x, r_y) will have its resolution increased to $(2r_x, 2r_y)$ with the each original cell being redistributed over 4 new cells, thus the new power per cell is simply $\Phi_{\text{cell_old}}/4$.

3.1.2 Light Ray Distribution




Figure 1: Item Buffer Construction.

In order to determine the solid angle, about the light sources within the scene, which must be sampled, we use a technique similar to the item buffer of [Weghorst84] and the hemicube of [Cohen85]. We wish to trace light-rays only in directions that will potentially give rise to some form of specular to diffuse transfer, so we first determine the directions in which specular surfaces are visible from the light source. All the objects in the scene are projected onto a unit cube centred at the light source with Z-buffers attached to its faces. During the light pass, we sample only those 'pixels' on the cube's faces that can 'see' a specular surface. The number of rays sent to each cube pixel is determined by the solid angle subtended by the pixel and a user specified sample density (specified in rays per steradian). The delta solid angle per pixel is derived in a similar manner to the hemicube delta form-factor:

$$\begin{aligned}
 x &= \frac{h_x + 0.5}{R_x} & y &= \frac{h_y + 0.5}{R_y} \\
 A &= \frac{4}{R_x R_y} & r &= \sqrt{x^2 + y^2 + 1} \\
 \cos\theta &= \frac{1}{\sqrt{x^2 + y^2 + 1}} & & (4.1) \\
 \Delta\omega &= \frac{\cos\theta}{Ar^2} \\
 &= \frac{R_x R_y}{4 * \left[\sqrt{x^2 + y^2 + 1} \right]^3}
 \end{aligned}$$

Because of the cube's symmetry, we only need to compute $\Delta\omega$ for one quarter of one face. (R_x, R_y) is the resolution of the cube faces. (h_x, h_y) is the co-ordinate of the cube pixel in cube-space, θ , the angle between the face normal and the direction of $\Delta\omega$, r , the distance to the centre of the cube and A the area of the cube-pixel.

Each ray carries with it a fraction of the total power of the source. This fraction is proportional to the solid angle subtended by the ray at the source. The power shot towards each cube-pixel for a point light source is:

$$\Phi_p = \frac{\Phi_s \Delta\omega}{\omega_s} = \frac{\Phi_s \Delta\omega}{4\pi} \quad (4.2)$$

This power is evenly distributed among the rays shot through the pixel. Note that this distribution of rays will produce a bias towards the centre of the cube faces. This bias may be reduced by stratifying the sampling of the cube faces, whereby the area of the cube pixels are

proportional to the solid angle they subtend, rather than the uniform area used here. To reduce aliasing artefacts, the rays are jittered within their sub-pixels in order to approximate a Poisson distribution on the cube face.

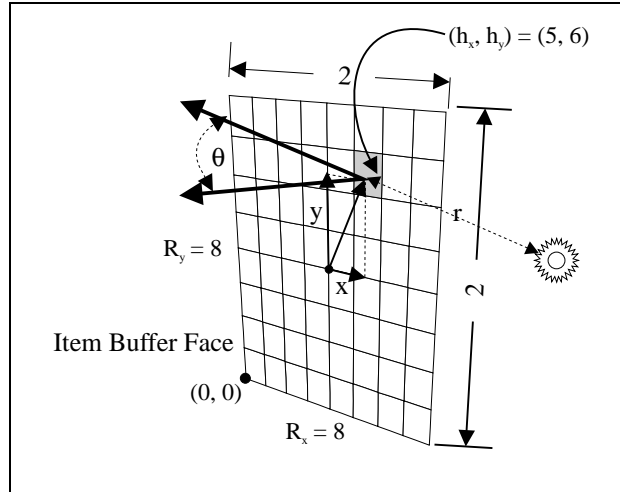


Figure 2: Delta solid angle per ray determination

3.2 Wavefront Behaviour Tracking

We model the behaviour of the wavefront emanating from the light source by examining the distribution of ray hits across surfaces. On a given surface we are interested in the convergence/divergence of the wavefront locally. Rather than attempt to reconstruct the shape of the wavefront by applying a post-filter to the intersection points we maintain connectivity information between rays. By tracking arrays of rays from the light source, we associate with each ray some fraction of the solid angle of the source. When these rays intersect surfaces we examine the distances between neighbouring rays to determine the shape of the wavefront. We use this information to determine how to distribute the power on the receiving surface.

Figure 3 demonstrates the problem of reconstructing the wavefront interaction with a surface. The distribution of the hit points is such that if a locality based filter is applied, the intensity is biased close to the sample points, and due to the coincidence of these sample points, we get a skewed estimate of the power distribution. Using connectivity information between rays, we can determine the area over which each ray must deposit its power and thus eliminate the bias introduced through the locality based filter schemes.

3.2.1 Deposit Area Determination

We maintain neighbourhood information between rays through the use of a dynamic caching system for each surface where the most recent set of rays and their collisions with the surfaces are recorded, thus for any new ray colliding with a surface, we can estimate local wavefront density by examining the last few ray hits. If enough information does not exist to determine the spread, deposition of power is postponed until new rays provide the required information.

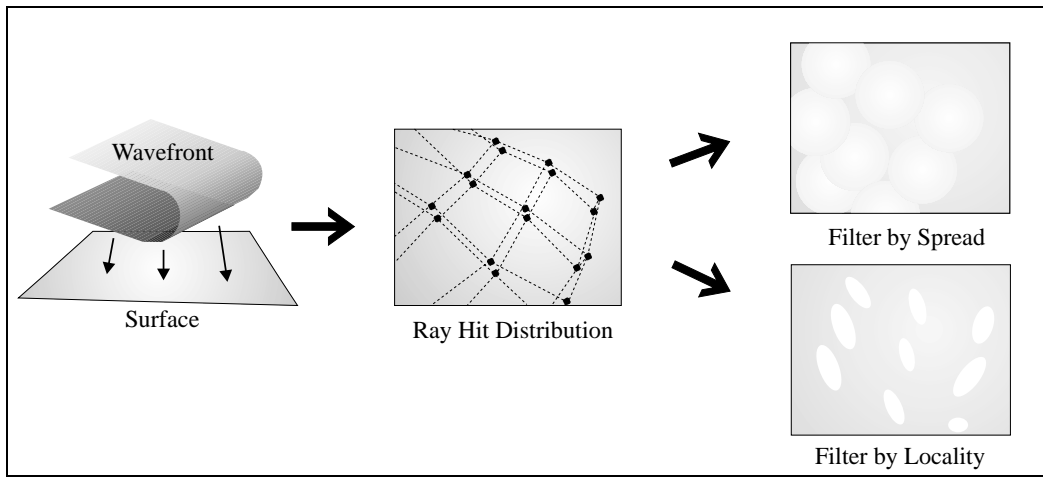


Figure 3: Effects of Different Filtering Approaches

Figure 4 depicts a subset of the rays from a light source incident on a surface, having passed through the scene, possibly interacting with specular surfaces. Only when ray number 4 hits the surface do we have enough information to estimate the spread of the wavefront. Rays 1 to 3 are cached until 4 arrives, and the estimated spread is used for all 4 deposits (as there will not be enough information to independently determine an estimate for rays 1 to 3). The spread estimate is simply the area of the polygon defined by the 4 ray hit points. This area is effectively the area over which we wish to deposit the power being carried by the ray.

Note that this area is proportional to the squared distance the ray has travelled and also the cosine of the angle between the ray direction and the normal to the surface. Thus when we are depositing the ray's power onto the surface we need not scale it by the usual cosine and distance squared factor.




Figure 4: Determine Wavefront Spread from Ray Intersections.

3.2.2 Power Deposition using Gaussian Kernels

Having determined the area over which we wish to deposit the power, we must now decide how to distribute this power across that area. Rather than distribute the power uniformly across the area which tends to result in aliasing at the edges of the regions (see [Watt90]) we deposit the power as a *splat*. This technique has been used to implement a progressive refinement algorithm for volume rendering [Laur91] where gaussian splats are used to represent the volume's octtree and the footprint of the splats scaled to match the projected area of the voxels making up the octtree. We deposit power, therefore, with the following distribution:

$$w(x,y) = e^{-\frac{r^2}{2\sigma}} \quad (4.4)$$

The footprint of the gaussian kernel is scaled to match the areas defined by the ray collisions. Intuitively we felt that we might further improve the estimate by adaptively scaling the gaussian spread factor, σ . This surprisingly gave less accurate results as detailed in section 5. Although the gaussian is generally accepted as a good choice of kernel for filtering what is essentially a warped image [Green86], we have also investigated box, disc and cone kernel topologies. Based on observation of the relative performance of the various kernels (see section 5), the gaussian was deemed the best choice.

In order to deposit the kernel of power on the surface, the kernel must be warped into the parameter space of the surface. The kernel is then sampled by the illumination map to extract the power from the kernel and store it in the map pixels. Having sampled the kernel, thus necessitating re-normalisation of the kernel, the application of the splat is a 2 pass procedure:

1. Sample the kernel into a temporary buffer.
2. Normalise the buffer and add it to the illumination map.

The current implementation assumes a rectilinear illumination map, thus the problem of warping is simplified to one of scaling. In order to speed up deposit times, an array of kernels is created before tracing begins. The problems presented by illumination maps on curved surfaces such as spheres has not been addressed. An improved kernel based on rotated elliptical gaussians ([Gotsman93]), designed to match the orientation and aspect ratio of the regions defined by ray-surface intersections is currently under investigation.

3.3 Eye pass

During the eye pass, the illumination map is queried in a similar fashion to a texture map (though, unlike texture maps, it is not subject to shadowing). The power in each pixel of the illumination map is converted to intensity making the assumption that the surface is perfectly diffuse:

$$L(x,y,\lambda) = \frac{\Phi_{(x,y)}}{\pi A_{(x,y)}} \quad (4.5)$$

Φ_p is the power accumulated in pixel (x,y) of the illumination map and A_p is the area of the pixel in world-space.

4. Spectral Considerations

Due to the potentially high intensity artefacts resulting from wavefront convergence through transmitting media, the wavelength dependency of the refractive index can be omitted in order to render realistic images.

Wavelength dependent reflection is taken into account in the Fresnel approximation presented by Cook et al. [Cook82], which approached the problem using a geometric optics model, where reflectance is a function of both wavelength and the angle of incidence of the incoming light. The model presented by He et al. [He91] is an extension of the Cook Torrance model, based on physical optics, whereby the reflectance of a surface is dependent on the wavelength, angle of incidence, surface roughness parameters and surface refractive index.

As a consequence of the *dispersion* of light as it passes through transmittive material we must sample the spectrum of light in order to correctly model light transmission. The sampling resolution is critical in order to resolve correct colour separation. The rendering system models rays carrying, not R, G and B values, but rather a spectral power curve, represented by n sample points.

4.1 Spectral Sampling

The choice of n is not clear. Musgrave, in his modelling of the rainbow, [Musgrave89], uses 13 samples, spread uniformly across the visible spectrum. Meyer [Meyer88], having examined a number of different sampling schemes, derived a new colour space, AC_1C_2 , and using gaussian quadrature techniques to integrate low order sampling polynomials applied to the spectral curve, samples into this new colour space. Having tried a number of sampling schemes, Meyer suggests a scheme where 4 samples are used.

For the sake of computational efficiency, we have opted for 7 sample points. The distribution of these sample points is driven by our desire to accurately capture dispersive effects that are commonly visible in real scenes. Our sample points are located around the wavelengths corresponding to the 7 ‘colours of the rainbow’. This clearly does not represent a definitive sampling scheme, however it produces satisfactory results. Colour Plates 3 and 4 show the effects of different spectral sampling resolutions. The test scene shows a prism and sphere of silicate flint glass. The prism is illuminated by a thin beam of white light which is dispersed to produce the spectrum on the back wall. Colour Plate 3 details the results of 2 sampling schemes. The top 6 spectra result from initial rays hitting the prism being refracted, generating a number of new rays distributed (jittered) through the visible spectrum. The bottom 6 spectra detail a ‘fairer’ scheme where each initial ray spawns only 1 refracted ray, the wavelength of which is determined stochastically.

All computations are performed with respect to these 7 sample points. For final image display however, we must convert to RGB space. This is achieved [Hall89][Peercy93] by multiplying the reconstructed spectrum by the CIEXYZ tristimulus matching curves, using Riemann summation to evaluate the integral, and finally converting the XYZ colour to RGB space using chromaticity data for the monitor phosphors.

4.2 Light Source Data

Clearly to model the interaction of light with surfaces at a spectral level we need information regarding the spectral power curve of the light source and the frequency dependent refractive index curve for the surfaces. The light source data is usually obtainable from manufacturers and is available in the form of spectral energy distribution $\Phi(\lambda)$ curves and the total luminous light power and must first be converted to spectral power (from luminous power) using the relationship: [Languenou92]

$$\Phi_L = 683 \int_{480}^{700} \Phi(\lambda) V(\lambda) d\lambda \quad (5.1)$$

Φ_L is the luminous power, $\Phi(\lambda)$ the spectral power, $V(\lambda)$ the *spectral luminous relative efficiency curve* describing the sensitivity of the eye and 683 is a conversion factor from Watts to Lumens (the Lumen having been defined in 1979 at a meeting in Paris of the General

Conference on Weights and Measures as “the amount of light of monochromatic radiation whose frequency is $540 \cdot 10^{12}$ Hz. and whose power is $\frac{1}{683}$ Watt.”).

4.3 Material Data

The *mean dispersion* ($\eta_F - \eta_C$) of a material defines the change of index of refraction of that material with respect to frequency and the following ratio is used to define the dispersive behaviour of a material (usually glass):

$$v = \frac{\eta_d - 1}{\eta_F - \eta_C} \quad (5.2)$$

η_F, η_C are the refractive indices of the material at the F and C Fraunhofer lines (the emission lines of hydrogen at 486.1nm(blue) and 656.3nm(red) respectively) and η_d the refractive index at the yellow d line of helium (587.6nm) rather than the sodium D doublet (589.3nm) used in older literature. This ratio is *Abbe's number*, also called the *v value* or *V-number* of the material. A high V-number indicates a material of low mean dispersion. Manufacturers of glass classify the glass by its glass number, a six digit code, the first three digits of which is the refractive index $\eta_d - 1$ and the last three digits Abbe's number scaled by 10. For example a glass with number 523588 has a refractive index of $\eta_d = 1.523$ and mean dispersion of 58.8 [Meyer-Arendt89]. Glasses with low mean dispersion (V-number above 55) are called crowns and glasses of high dispersion (V-number below 50) are called flints.

Using this information we can reconstruct a curve relating frequency and refractive index, and sample this curve at our spectral sample points to determine the index to use for our light-rays. Another approach is to adopt one of the many attempts to formulate a quantitative relationship between refractive index and frequency such as that of Herzberger [Herzberger59]:

$$\eta = A + B\lambda^2 + \frac{C}{\lambda^2 - \lambda_0^2} + \frac{D}{(\lambda^2 - \lambda_0^2)^2} \quad (5.3)$$

A, B, C and D are empirically determined constants, and λ_0 is set at 168nm. Values for A, B, C, and D are difficult to come by and therefore the glass number is a more convenient representation. The glass used in the example scenes included in this paper has been modelled on Silicate Flint Glass with glass number 635577.

5. Results

Table 1 demonstrates the effect of light ray density on image quality. The corresponding images are shown in Colour Plate 5.

	Image A	Image B	Image C
Ray Density (rays/sr)	7957.7	79577.5	795774.7
No. of Light Rays	3467	34680	346813
Light-Pass Time	11.57 secs.	53.27 secs.	379.22 secs.
Eye-Pass Time	148.33 secs.	144.72 secs.	135.96 secs.

Table 1: Effect of light pass ray density on image quality and rendering times¹.

In order to demonstrate the effects of illumination map resolution on the image quality, we have fixed the illumination map to 3 separate resolutions in Colour Plate 6. The images were rendered with illumination maps of 30x30, 100x100 and 300x300 respectively. The resolution differences have little effect on the rendering times.

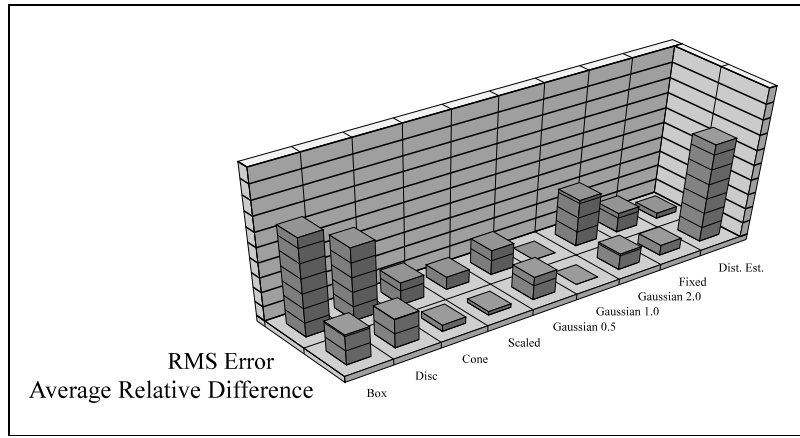


Figure 5: Relative performance of various filter topologies and size.

Figure 5 details the relative performance in terms of RMS error and average relative difference of some different choices of filter shape (box, disc, cone, gaussians with $\sigma = 0.5, 1.0, 2.0$, adaptive) and size (scaled according to distance travelled by the ray, and a fixed size kernel). RMS error and Average relative error are defined as:

$$RMSError = \sqrt{\sum (I_s - I_r)^2} \quad (5.1)$$

$$ARDerror = \frac{1}{n} \sum \frac{|I_s - I_r|}{I_r} \quad (5.2)$$

¹ All the images were rendered and the timings taken on a 486 66Mhz PC running BSD Unix. (©Berkeley Software Design Inc.)

where I_r is the intensity of the reference map (centre of Colour Plate 1), I_s the intensity of the source map, and n the number of pixels in the map.

Colour Plate 1 depicts visually the performance of the filters with Colour Plate 2 indicating the test scene used to generate the data. As can be seen the gaussian kernel with $\sigma = 1.0$ performs most satisfactorily. Note how the fixed kernel size results in an image that is indistinct, failing to capture the high frequency aspects of the epicycloidic curve. The other filter shapes also fail to capture the clarity of the gaussian kernel. The distance estimate performs surprisingly well in terms of RMS error, but results in a noisy image captured by the ARD error metric indicating that the error was present despite a low error variance.




Figure 6: (a) Image quality with varying numbers of light-pass rays and (b) different gaussian kernel shapes.

Figure 6(a) shows the performance of the system by increasing the numbers of rays. Figure 6(b) depicts the relationship between image quality and gaussian spread factor σ . The error is at a minimum when $\sigma = 0.9..1.0$.

To demonstrate the application of the system on complex objects, a classic example of caustic formation, the swimming pool floor, was rendered following the example of [Watt90]. Colour Plate 7 is a view from above, and Colour Plate 8 a view from below the surface of the water. As remarked by Watt, the triangular patch resolution of the water's surface is amplified by the caustic image, requiring very high patch subdivision resolutions. To avoid this we modelled the surface as a single polygon having a bump map composed of overlapping cosine terms [Max81]. Plate 9 is a view of the caustics through a complex object made up of 3958 polygons. Note that even with this large number of polygons, the caustics exhibit regular aliasing artefacts due to the resolution of the dataset.

6. Conclusions

We have presented an algorithm extending existing specular to diffuse transport models with a view to capturing both high and low frequency artefacts through the use of adaptive gaussian

kernels or splats. We also extend the method to allow for accurate modelling of frequency dependent refraction of light with a view to capturing the beauty of dispersion. The algorithm has been designed to produce illumination maps to be used as texture maps in a real-time walk-through simulation system but can be used as part of a general purpose global illumination algorithm as an extension to existing ray-tracing/radiosity systems.

Through the use of the item buffer we can cut down on the number of first generation light rays and by adaptively scaling the illumination maps according to the average splat size we concentrate expensive map memory where it is most needed.

Many problems remain. In order to successfully handle surfaces with non linear parameter space we must warp the gaussian kernels to the parameter space of the surface before applying the kernel. Our current implementation assumes rectilinear texture maps.

More intelligently constructed item buffers should be used in order not to introduce bias into the distribution of energy from the light source.

The illumination maps are typically very expensive in terms of memory usage, despite the adaptive strategy adopted here. We feel that a hybrid approach might be used analogous to that adopted for ray tracing space subdivision, where the best speeds have been achieved by a shallow octree of uniform grids. This would entail the illumination maps being added to surfaces as surface details (similar to the approach of [Watt90]), with potentially many separate illumination maps spread across the surface.

The current spectral sampling scheme is quite primitive and could be improved through more attention to perceptual issues [Meyer88]. By importance sampling the spectrum of visible light with sample distribution proportional to the spectral luminous relative efficiency curve of the eye, more work will be done where the eye will perceive the difference.

7. Acknowledgements

I wish to thank my supervisor, Dr. Dan McCarthy for his continuous advice and support, Prof. John Byrne for continuing to support and fund this work, Pete Shirley for helpful information regarding global illumination and finally my colleagues in the Image Synthesis Group, Trinity College Dublin.

8. References

- [Arvo86] Arvo J., "Backward Ray Tracing", *SIGGRAPH '86 Developments in Ray Tracing seminar notes*, Vol. 12, Aug. 1986.
- [Chen91] Chen S.E., Rushmeier H.E., Miller G., Turner D., "A Progressive Multi-Pass Method for Global Illumination", *Computer Graphics*, Vol. 25, No. 4, pp. 165-174, July 1991.

- [Cohen85] Cohen M.F., Greenberg D.P., "The Hemi-Cube; A Radiosity Solution for Complex Environments", *Computer Graphics*, Vol. 19, No. 3, pp. 31-40, July 1985.
- [Cook82] Cook R.L., Torrance K.E., "A Reflectance Model for Computer Graphics", *ACM Transactions on Graphics*, Vol. 1, No. 1, pp. 7-24, 1982.
- [Cook84] Cook R.L., Porter T., Carpenter L., "Distributed Ray Tracing", *Computer Graphics*, Vol. 18, No. 4, pp. 165-174, July 1984.
- [Goral84] Goral C., Torrance K.E., Greenberg D.P. and Battaile B., "Modelling the Interaction of Light between Diffuse Surfaces", *Computer Graphics*, Vol. 18, No. 3, pp. 311-317, August 1984.
- [Gotsman93] Gotsman C., "Constant Time Filtering by Singular Value Decomposition", *Proceedings of the 4th Eurographics Workshop on Rendering*, pp. 145-155, June 1993.
- [Green86] Green N., Heckbert P.S., "Creating Raster Omnimax Images from Multiple Perspective Views using the Elliptical Weighted Average Filter", *IEEE Computer Graphics and Applications*, Vol. 6, No. 11, pp. 21-29, 1986.
- [Hall89] Hall R., *Illumination and Color in Computer Generated Imagery*, Monographs In Visual Communication, Springer Verlag, pp. 45-62, 1989.
- [He91] He X.D., Torrance K.E., Sillion F.X., Greenberg D.P., "A Comprehensive Physical Model for Light Reflection", *Computer Graphics*, Vol. 25, No. 4, pp. 175-186, July 1991.
- [Heckbert84] Heckbert P., Hanrahan P., "Beam Tracing Polygonal Objects", *Computer Graphics*, Vol. 18, No. 3, pp. 119-127, July 1984.
- [Heckbert90] Heckbert P., "Adaptive Radiosity Textures for Bidirectional Ray Tracing", *Computer Graphics*, Vol. 25, No. 4, pp. 145-154, August 1990.
- [Heckbert92] Heckbert P., "Introduction to Global Illumination", *SIGGRAPH '92 Global Illumination Course*, Vol. 18, July 1992.
- [Herzberger59] Herzberger M., "Colour Correction in Optical Systems and a New Dispersion Formula", *Opt. Acta (London)*, Vol. 6, pp. 197-215, 1959.
- [Immel86] Immel D.S., Cohen M.F., Greenberg D. P., "A radiosity Method for Non-Diffuse Environments", *Computer Graphics*, Vol. 20, No. 4, pp. 133-142, August 1986.
- [Languenou92] Languenou E., Tellier P., "Including Physical Light Sources and Daylight in a Global Illumination Model", *Third Eurographics Workshop on Rendering*, pp. 217-226, May 1992.
- [Laur91] Laur D., Hanrahan P., "Hierarchical Splatting: A Progressive Refinement Algorithm for Volume Rendering", *Computer Graphics*, Vol. 25, No. 4, pp. 285-288, July 1991.
- [Max81] Max N.L., "Vectorised Procedural Models for Natural Terrain: Waves and Islands in the Sunset", *Computer Graphics*, Vol 15, No. 3, pp 317-324, August 1981.
- [Meyer88] Meyer G.W., "Wavelength Selection for Synthetic Image Generation", *Computer Vision, Graphics & Image Processing*, Vol. 41, pp 57-79, 1988.
- [Meyer-Arendt89] Meyer-Arendt J.R., *Introduction to Classical and Modern Optics*, Third Edition, Prentice-Hall International, pp. 13-26, 1989.

- [Mitchell92] Mitchell D., Hanrahan P., "Illumination from Curved Reflectors", *Computer Graphics*, Vol. 26, No. 2, pp. 283-291, July 1992.
- [Musgrave89] Musgrave F.K., "Prisms and Rainbows: a Dispersion Model for Computer Graphics", *Graphics Interface '89*, pp. 227-234, 1989.
- [Peercy93] Peercy M.S., "Linear Color Representations for Full Spectral Rendering", *Computer Graphics Proceedings, Annual Conference Series 1993*, pp. 191-198, August 1993.
- [Poston78] Poston T., Stewart I., "Optics and Scattering Theory", *Catastrophe Theory and its Applications*, pp. 246-252, 1978.
- [Shinya87] Shinya M., Takahashi T., Naito S., "Principles and Applications of Pencil Tracing", *Computer Graphics*, Vol. 21, No. 4, pp. 45-54, July 1987.
- [Shirley90] Shirley P., "A Ray Tracing Algorithm for Global Illumination", *Graphics Interface '90*, pp. 205-212, May 1990.
- [Sillion91] Sillion, F.X., Arvo J.R., Westin S.H., Greenberg D.P., "A Global Illumination Solution for General Reflectance Distributions", *Computer Graphics*, Vol. 24, No. 4, July 1991.
- [Wallace87] Wallace J.R., Cohen M.F., Greenberg D.P., "A 2-Pass Solution to the Rendering Equation: A Synthesis of Ray Tracing and Radiosity Methods", *Computer Graphics*, Vol. 21, No. 4, pp. 311-320, July 1987.
- [Ward88] Ward G.J., Rubinstein F.M., Clear R.D., "A Ray Tracing Solution for Diffuse Interreflection", *Computer Graphics*, Vol. 22, No. 4, pp. 315-324, August 1988.
- [Ward92] Ward G.J., Heckbert P.S., "Irradiance Gradients", *Proceedings of the 3rd Eurographics Workshop on Rendering*, pp. 85-98, May 1992.
- [Watt90] Watt M., "Light Water Interaction using Backward Beam Tracing", *Computer Graphics*, Vol. 24, No. 4, pp. 377-385, August 1990.
- [Weghorst84] Weghorst H., Hooper G., Greenberg D.P., "Improved Computational Methods for Ray Tracing", *ACM Transactions on Graphics*, Vol. 3, No. 1, pp. 52-69, January 1984.
- [Whitted80] Whitted T., "An Improved Illumination Model for Shaded Display", *Communications of the ACM*, Vol. 23, No.6, pp. 343-349, June 1980.