

# Stratified Wavelength Clusters for Efficient Spectral Monte Carlo Rendering

Glenn F. Evans

Michael D. McCool

Computer Graphics Laboratory  
Department of Computer Science  
University of Waterloo

## Abstract

Wavelength dependent Monte Carlo rendering can correctly and generally capture effects such as spectral caustics (rainbows) and chromatic aberration. It also improves the colour accuracy of reflectance models and of illumination effects such as colour bleeding and metamerism.

The stratified wavelength clustering (SWC) strategy carries several wavelength stratified radiance samples along each light transport path. The cluster is split into several paths or degraded into a single path only if a specular refraction at the surface of a dispersive material is encountered along the path. The overall efficiency of this strategy is high since the fraction of clusters that need to be split or degraded in a typical scene is low, and also because specular dispersion tends to decrease the source colour variance, offsetting the increased amortized cost of generating each path.

*Key words:* Monte Carlo methods, wavelength dependent (spectral) rendering, caustics, rainbows, refraction, reflectance, global illumination.

## 1 Introduction

The faults of using three component colour models for computing reflectance, absorption and dispersion are well known [13, 14].

Most physical processes that produce changes in radiant power spectra and colour (reflection, absorption, scattering, etc.) are poorly approximated by models that assume smooth spectra, linearity, or linear separability of colour computations into separate geometric and wavelength dependent factors. For instance, absorption computations involve the exponentiation of a spectrum, while effects such as dispersion involve almost a purely functional dependence between wavelength and geometry. A trichromatic model is suitable for perceptual representations of colour, but not for physical computations.

Any attempt to implement physically-based rendering should therefore consider the transport of *spectral* power distributions and should model reflectances and absorptions spectrally rather than approximate them with a trichromatic model. In this pa-

per we demonstrate not only that this is feasible, but also that with a strategy we call *stratified wavelength clustering* (SWC) the marginal cost is negligible for Monte Carlo ray tracing and bidirectional path tracing.

## 2 Outline

Prior work on wavelength dependent and spectral rendering is surveyed in Section 3. Section 4 presents the stratified wavelength cluster strategy. Three splitting strategies are also presented and compared. In Section 5 a bidirectional path tracer that uses stratified wavelength clusters is described. Finally, in Section 6 we present our results, comparing and contrasting crude Monte Carlo integration over wavelength, quasi Monte Carlo integration over wavelength using Halton sequences, and the stratified wavelength cluster strategy.

## 3 Background

Before presenting our approach for extending Monte Carlo global illumination algorithms to wavelength-dependent (spectral) rendering, we first review the relationship of spectral power densities to perceived colour and review wavelength-dependent phenomena that have a significant effect on the generated image. Previous algorithms and representations for spectral rendering are also discussed.

### 3.1 Colour and Spectra

Given a spectral radiant power distribution  $\Phi(\lambda)$ , its CIE XYZ colour coordinates can be computed with the following three integrations [2], using the CIE 1964 matching functions  $\bar{x}_{10}(\lambda)$ ,  $\bar{y}_{10}(\lambda)$ , and  $\bar{z}_{10}(\lambda)$ ,

$$X = k_{10} \int \Phi(\lambda) \bar{x}_{10}(\lambda) d\lambda,$$

$$Y = k_{10} \int \Phi(\lambda) \bar{y}_{10}(\lambda) d\lambda,$$

$$Z = k_{10} \int \Phi(\lambda) \bar{z}_{10}(\lambda) d\lambda,$$

where  $k_{10} = 683\text{lm/W}$ . These integrations should be interpreted as inner products that project the infinite-dimensional function  $\Phi(\lambda)$  down onto a finite-dimensional vector space. The functions  $\bar{x}_{10}$ ,  $\bar{y}_{10}(\lambda)$ , and  $\bar{z}_{10}(\lambda)$  only *analyze* the spectrum  $\Phi(\lambda)$ ; they do not constitute a basis to reconstruct the original spectrum. In fact, there are an infinite number of *reconstruction bases* that will generate different spectra with the same XYZ coordinates.

The significance of the XYZ coordinates is perceptual: all spectra with the same XYZ coordinates should be perceived as the same colour by normal human observers. Such perceptually equivalent spectra are called *metamers*.

A  $3 \times 3$  matrix multiplication can convert XYZ colour coordinates to any desired linear trichromatic colour coordinate system [7]. Given the XYZ coordinates of a set of phosphors on a monitor, a matrix  $M$  can easily be derived to convert from XYZ coordinates to the monitor's RGB phosphor space, at least for colours within the gamut of the monitor and within the limits of an assumption of linearity [3, 5].

The traditional approach to rendering immediately converts spectral distributions into linear trichromatic colour coordinates, such as RGB coordinates, at the time of definition of the light sources and reflectances. In fact, usually perceptual coordinates are all that is given, since colours are often chosen interactively by observing the colour produced by a monitor.

Unfortunately, while a trichromatic linear model is all that is needed for the perceptual *representation* of a colour, it is too coarse to accurately compute the physical *interactions* of spectra.

### 3.2 Wavelength Dependent Phenomena

Due to metamerism, the apparent colour of an object can depend critically on the illumination spectrum. Two objects which have the same colour under one illuminant can look very different under another illuminant. Many illuminant spectra, such as those of fluorescent lights, can be far from smooth, and this can result in a significant dependence of the final colour on the detailed shape of a reflectance or absorption spectrum.

Many wavelength dependencies in physically-based reflectance models can be traced to the variation of the complex index of refraction with wavelength. The variation with wavelength is quite strong; Sellmeier's Law [4], which is a good fit to observed data, results in an  $O(\lambda^3)$  dependency in  $d\eta/d\lambda$ . The complex refractive index is used in the Fresnel formulas to compute the ratio of reflected to transmitted light at a specular surface. The Fresnel formulas are an important component of all physically-based models, since they result directly from the continuity conditions required by Maxwell's equations.

Unfortunately, the Fresnel formulas have a non-separable and highly nonlinear dependence on both the wavelength dependent refractive index and the incident angle. Other non-separable wavelength dependent effects, such as interference and diffraction, can also be important in reflectance models [6], and furthermore can lead to "spiky" reflectance and transmission spectra.

Consider also absorption in participating media. A homogeneous material absorbs some fraction  $\alpha(\lambda)$  of radiant energy

per unit length. Light travelling a length  $z$  through such a homogeneous material will have its spectral power modified by the multiplicative absorption factor given by Beer's Law:

$$\begin{aligned}\xi(\lambda, z) &= \exp(-\alpha(\lambda)z) \\ L_{\text{out}}(\lambda, 0) &= \xi(\lambda, z)L_{\text{in}}(\lambda, z).\end{aligned}$$

The cumulative absorption  $\xi(\lambda, z)$ , a function of both wavelength and distance, is a nonlinear function of another spectrum  $\alpha(\lambda)$ . The function  $\xi(\lambda, z)$  is also not separable into factors containing only  $z$  and  $\lambda$  alone. In fact, longer distances tend to increase the contrast (change the shape) of the absorption spectrum, resulting in a more saturated colour for thicker materials.

Participating media can also have wavelength-dependent scattering distributions [4, 10], which can result in similar distance-dependent spectral colour shifts.

Finally, the refractive index determines the angle of specular refraction, by Snell's Law. This effect is called *dispersion* because polychromatic light tends to be spread apart into monochromatic light, producing saturated colours from unsaturated ones. This is the strongest wavelength dependent effect and also the hardest to handle properly, since it results in a direct dependence of ray direction on wavelength [18, 22]

### 3.3 Linear Spectral Representations

One way to extend a renderer to handle spectral phenomena is to represent spectra explicitly and correctly implement the operations, primarily multiplication, that can take place on them. A simple dense vector with 80 samples can handle most physical interactions and can represent spiky spectra accurately, but such a representation is too expensive to be practical.

Other models with a smaller number of parameters have been used in an attempt to improve efficiency. In the linear representation, a vector of colour coordinates  $(w_1, w_2, \dots, w_K)$  weights a set of basis functions  $\{b_1(\lambda), b_2(\lambda), \dots, b_K(\lambda)\}$  which are summed to reconstruct the spectrum:

$$\Phi(\lambda) \approx \sum_{i=1}^K w_i b_i(\lambda).$$

The traditional trichromatic models can be interpreted as linear representations if we assign appropriate basis functions.

In the context of spectral nonlinearities (as in absorption) or when there is a strong geometric dependence on wavelength (as with specular dispersion) or when two spectra can interact by multiplication (reflection and absorption), using a trichromatic representation to compute the interaction of spectra can lead to major colour and illumination errors.

In particular, absorption cannot be correctly handled using only trichromatic colour computations [7] due to its nonlinear  $\exp(-\alpha(\lambda)z)$  factor.

In the case of reflection, elementwise multiplication of colour coordinates is equivalent to spectral multiplication if and only if the basis functions used to reconstruct the spectra are non-overlapping box functions. For only three colour coordinates such basis functions are subject to aliasing and/or significant approximation error.

Unfortunately, raising the dimensionality of linear representations does not help much unless the basis functions are very carefully chosen [7]. Certain illuminants, such as fluorescent lighting, can contain spikes at specific frequencies, and these spikes will be poorly captured by a non-adaptive representation scheme. As noted above, some kinds of reflective phenomena can also have very spiky spectra, for instance if thin-film interference or diffraction is involved.

For any given scene, especially in the context of indirect and global illumination, the large number of possible multiplicative interactions between illumination and reflectance spectra can lead to a large class of possible spectra.

Peercy [13] has studied generalized linear models, in which a set of orthogonal basis functions are chosen based on a characteristic vector analysis of the “important” spectra in a scene. In practice, it can be difficult to determine which of the possible spectra will actually be present in the scene and contribute significantly to the final image. If overlapping basis functions are chosen (the usual result of a characteristic vector analysis) the cost of computing reflectance (or any spectrum-to-spectrum multiplicative interaction) with this scheme also grows as  $O(K^2)$ , where  $K$  is the number of basis functions.

The generalized linear approach is only feasible if spectral effects are linearly separable from geometric influences (a poor physical assumption). Otherwise, a matrix with  $O(K^2)$  entries needs to be recomputed for every surface reflection, with the computation of each entry in the matrix involving an integration of a spectrum against the product of two basis functions. Even if nonoverlapping basis functions are chosen (which reduces approximation power) the cost of the  $O(K)$  integrations required to project the spectra against the basis functions can be prohibitive.

Raso [14] uses piecewise polynomials to represent spectra. When spectra must be multiplied, this raises the degree of the polynomials, and so they must be projected down onto lower degree polynomials. This approach can be considered a generalization of the linear model where approximation error is handled more explicitly using polynomial approximation algorithms. Unfortunately, as with Peercy’s scheme, if the reflectance models and absorption models are not separable, the coefficients for the polynomial representations of surface colour cannot be precomputed, and the approach’s marginal cost becomes too high to be practical.

### 3.4 Point Sampling Approaches

In general, for computing nonlinearities and multiplication of spectra, point sampling strategies have  $O(n)$  complexity for  $n$  wavelengths. Unfortunately, any deterministic point sampling strategy can be subject to aliasing. Choosing a high enough sampling rate can overcome aliasing, if the spectra are “smooth enough”. It has been shown that in the absence of aliasing, a small number of spectral samples, on the order of eight or nine, can capture most perceivable spectrally-generated phenomena [7].

Practically speaking, the marginal cost of using nine point wavelength samples (rather than three RGB samples) is small, since the reflectance computation usually grows slowly with the

number of point samples: overhead costs and common terms can be shared, and the floating point code typical of reflectance models can be pipelined. Typically, rendering costs are usually dominated by geometric intersection computations, not reflectance computations.

Meyer [11] has studied point-sampling approaches to spectral rendering. A Gaussian integration scheme can be used, in conjunction with a specific class of basis functions, to select the best set of wavelengths to sample to minimize perceptual spectral error during rendering. However, Meyer’s approach to optimizing sampling considers only the case of smooth spectra, and so cannot handle the disrupting influence of nonsmooth illuminants.

Gondek, Schramm, Meyer, and Newman [6, 16] have studied the simulation and representation of wavelength dependent BRDFs based on explicitly modelled subsurface scattering and interference. However, their microgeometry simulator [16] uses only a single wavelength per ray, which we will show is fairly inefficient.

Dispersion can also be handled with adaptive splitting at dispersive interfaces, as in Thomas’ work [18]. However, splitting can result in a large number (100+) of secondary rays per primary ray spread over many wavelengths, far more than are needed to simply integrate over the spectrum. Individual rays are also more complex; a spread angle and spread vector must be carried along after the first dispersive interaction, and splitting rules must take this spread angle into account.

Thomas’ approach has been optimized by Yuan *et al* [22] in the case of polyhedral objects, using a pencil of three rays to estimate and simplify the splitting rules. Yuan *et al* also use quadratic interpolation of spectral effects to approximate dispersion directions, improving performance considerably.

None of the point-sampling approaches considered to date have been integrated with Monte Carlo global illumination algorithms, such as bidirectional path tracing.

## 4 Stratified Wavelength Clusters

An ideal and physically correct approach to spectral rendering should defer projection of spectra onto the perceptual matching functions until after all light transport has been calculated.

In this case, all transport quantities are spectral distributions and the colour coordinate integrals should be incorporated into the sensor responses of the measurement integral [4, 19]:

$$Y_{ij} = \iiint W_{ij}^Y(\lambda, \underline{\mathbf{x}}, \hat{\omega}) L(\lambda, \underline{\mathbf{x}}, \hat{\omega}) d\underline{\mathbf{x}} d\hat{\omega} d\lambda$$

with  $W_{ij}^Y(\lambda, \underline{\mathbf{x}}, \hat{\omega}) = k_{10} \bar{y}_{10}(\lambda) W_{ij}(\underline{\mathbf{x}}, \hat{\omega})$  being the spectral sensor sensitivity<sup>1</sup> for pixel  $ij$ , for colour coordinate  $Y$ . Similar definitions need to be used for  $X_{ij}$  and  $Z_{ij}$ .

As a simple alternative approach to rendering colour, these integrals can be estimated using a Monte Carlo approach. All Monte Carlo rendering techniques generate paths from light

<sup>1</sup>Adaptive renderers typically use  $W_{ij}^Y(\lambda, \underline{\mathbf{x}}, \hat{\omega})$  or an appropriate nonlinear transformation of  $Y$  as the “importance”, since the  $Y$  colour component is proportional to the human visual system’s luminance sensitivity under photopic viewing conditions.

sources to the pixel sensors. In a spectral renderer, the generation of each path should depend on a wavelength  $\lambda_i$  generated at random using an importance probability distribution  $p(\lambda)$ , just as it depends on other random choices to select a subpixel location, light source, point on the lens aperture, each reflection direction at every glossy surface encountered, etc.

After a set of paths have been generated, a weighted average of the power density transported over  $N$  paths can be computed to estimate the measurement integral:

$$Y_{ij} \approx \frac{1}{N} \sum_i \frac{W_{ij}^Y(\lambda_i, \mathbf{x}_i, \hat{\omega}_i) L(\lambda_i, \mathbf{x}_i, \hat{\omega}_i)}{p(\lambda_i)}.$$

This approach can tolerate spikes in the illuminants’ spectral power distribution if  $p(\lambda)$  is proportional to the chosen illuminant’s power distribution. If other importance sampling distributions are used, for example for glossy reflection, correction factors should likewise be included in the above calculation. As is usual in Monte Carlo methods, aliasing energy is converted into noise rather than systematic error.

Unfortunately, we have found that this direct Monte Carlo approach, even with quasi Monte Carlo sampling of wavelength as well as the other important dimensions, leads to an increase in source variance of up to an order of magnitude relative to the standard trichromatic approach, with a very objectionable loss of interpixel colour coherence. This more than triples the error and renders the naive direct scheme unusable. The problem is that each path takes almost as much time to calculate as before, but each carries only roughly a third of the information.

We therefore modify the direct scheme as follows: rather than choosing a *single* random wavelength per path, a *cluster* of  $K$  importance sampled, stratified random wavelengths is generated. This is done by first picking a light source at random, then picking a set of  $K$  random numbers stratified into  $K$  cells  $[i/K, (i+1)/K)$  for  $i = \{0, 1, 2, \dots, K-1\}$ , and then warping this set of random numbers through the inverse of the cumulative normalized spectral power distribution of the light source. The normalized spectral power distribution itself is used as  $p(\lambda)$  in the above estimator.

The rest of the path is then constructed as usual, but the power transfer is computed for all wavelengths simultaneously at every interaction with a surface or with a participating medium. As noted previously, the incremental cost for the computation of power transfer for a cluster of wavelengths *vs.* a single wavelength is small.

Certain effects, such as specular dispersion (wavelength dependent refraction), can decrease the effectiveness of this strategy, since the direction of the refracted ray becomes a deterministic function of the wavelength. Energy at different wavelengths is forced to travel along different paths. In this case, one of the following strategies must be used:

**Strategy 1: Degradation.** Transport light energy using a cluster, but if a specular refraction is encountered, discard all but one of the wavelengths (chosen randomly) and continue as with the naive single-wavelength approach.

**Strategy 2: Splitting.** If dispersion forces a change in direction directly dependent on wavelength, then continue trac-

ing a set of single wavelength paths from the point of dispersion [1].

**Strategy 3: Deferral.** Compute only power transfer for a single “primary” wavelength while generating the path. If no specular dispersive interaction is found, generate a cluster and compute power transfers at other wavelengths.

Note that all these approaches degrade performance to that of the direct Monte Carlo approach, but only in the difficult case of specular dispersion, and only for paths that encounter a specular interface on an object made of a material with a wavelength dependent refractive index.

In the case of specular dispersion, the visual effects (increased colour saturation and monochromaticity) tend to counter the increased amortized cost of only using one wavelength per path.

However, the strategies differ in their practical implications. Degradation throws away some information, but does not have to go back and do any extra computation afterwards if there is no specular dispersion. It will be the most successful if specular dispersion is rare. Splitting will result in many correlated samples with low power transfer, but does not waste any computation. It will be most successful when computing surface intersections is relatively expensive, and we are not supersampling heavily anyways. Deferral will unfortunately incur significant overhead, invoking reflectance computations twice for what we assume will be the most common case, non-dispersive reflection. This will reduce the potential for reuse of computation and will increase the marginal cost.

Splitting is unbiased since the cluster can just be interpreted as an optimization for carrying along many paths at once in the absence of dispersion. Degradation is unbiased because it is mathematically equivalent to a discrete approximation of splitting. Instead of continuing all  $K$  wavelengths through the scene, the summation of the split wavelengths is approximated by choosing one of the wavelengths at random with a probability of  $1/K$  (thereby increasing its weight by a factor of  $K$ ). Effectively, the accuracy of the spectral estimator is decreased during dispersive interactions.

Finally, deferral is unbiased if degradation is, we’ve just reordered the operations and avoided rather than discarded computations.

Effectively, degradation and deferral change the accuracy of the spectral estimator by increasing or decreasing the number of strata. Different estimators of the same value can be combined in an affine combination to produce an unbiased estimator [19].

Note as well if wavelength-dependent importance functions are used, for example to choose sampling directions for glossy reflections, that one sample in the cluster, the “primary wavelength”, should be chosen at random to be representative of the cluster. This representative sample will have the same probability distribution with respect to wavelength as the rest of the cluster.

## 5 Bidirectional Path Tracing

We have modified a bidirectional path tracer [19] to use stratified wavelength clusters. One of our goals was to incorporate

spectral rendering into a global illumination framework, since indirect illumination can cause a wide variety of spectra to be generated in the scene. When paths are shot from the light source and from the eye, naturally both paths must use the same wavelength(s) in order for a “join” to be made.

We assume specular dispersion will be rare in real scenes and so we used degradation, rather than splitting or deferral.

Importance sampling was used for wavelength selection from the light sources (using their power spectra) and reflection direction selection (using a Phong approximation to the BRDF as an importance function). One “primary wavelength” was chosen at random from each cluster and used for BRDF importance calculations as well as for Russian roulette path termination.

Antialiasing was implemented by integration over the support of a filter function centered at each pixel. The balance heuristic was used to combine samples from multiple estimators [20]. Depth of field and area sources were not used in our test images. Adaptive sampling was not used, nor were the images shown here postprocessed to reduce noise.

We implemented the following four variations for comparison purposes:

**CMC: Crude Monte Carlo.** A pseudorandom number generator was used (`drand48`) to generate Monte Carlo parameters. Jittering was used to stratify subpixel positioning.

**QMC: Quasi Monte Carlo.** A Halton sequence generator was used rather than a pseudorandom number generator. The primes used to generate the sequence were permuted from pixel to pixel, to avoid structured aliasing artifacts.

The first few dimensions of a Halton sequence are more effective at lower sampling rates. As with stratification, above 10 dimensions or so Halton sequences have effectively the same convergence rate as crude Monte Carlo [12], and so we only use Halton sequences for the most important dimensions.

After the selection of a wavelength and subpixel position, the remaining odd dimensions were used for reflectance directions on the way from the light source, and the even dimensions were used for reflections on the way from the eye.

**CMC with SWC:** As above, but with stratified wavelength clusters. Jittered wavelengths were chosen to represent a spectrum, using the cumulative spectral power distribution of the illuminant of the light source to perform spectral importance sampling.

If a wavelength dependent specular interaction with a surface occurs, only power transfer for the primary wavelength is carried forward (*i.e.* the algorithm degrades to CMC).

**QMC with SWC:** As above, but with a Halton sequence generator rather than a pseudorandom number generator.

However, we found that use of Halton sequences for generating wavelengths did not work very well, as almost all the useful dimensions of the Halton sequence would be used up in the generation of the wavelength cluster, since each sample in the cluster would require a new dimension.

Instead, a pseudorandom number generator was used to generate the cluster, which was stratified in the usual way.

Unfortunately, the bidirectional path tracing algorithm has a flaw: some of the interesting direct and indirect spectral effects we would like to observe cannot be generated. Specifically, since the path segments joining the eye and light subpaths must be deterministically generated, power transfer along these segments can only be computed between two non-specular surfaces. Even light tracing requires that a deterministic (next-event) ray be shot towards the eye.

Certain effects we would like to observe cannot be generated with only point sources. For example, the “fire” observed in diamonds is the result of specular dispersion, but the path through the diamond to a point source is entirely specular [22, 18]. An area source or indirect illumination would permit the generation of fire, but the specular paths that hit the light source would still be generated with low probability. This is not a problem unique to our algorithm; Thomas’ algorithm suffers from a similar problem and will have similarly low efficiency if the solid angle subtended by the light source is small.

On the other hand, we were able to generate spectral caustics, indirect illumination due to these caustics, and chromatic aberration of nonspecular objects when seen through dispersive materials.

## 6 Results

Figures 3, 4, 5, and 6 were four of the scenes we used for testing.

Figure 3 (Scene 1) compares direct lighting *vs.* global illumination. Most lighting in this scene is indirect. Colour bleeding can be seen off the red wall to the right; the yellow colour of the ceiling is due to the green reflectance of the lamp shades combined with the spectral power distribution of the light source. The sphere is rendered with an HTSG model of gold, using a spectrally dependent complex refractive index.

Figure 4 (Scene 2) has only diffuse surfaces. The figure compares the results with and without stratified wavelength clusters, using CMC sampling. On the top row, we see that at low sampling rates wavelength clusters greatly improve colour coherence. On the bottom row we see that clusters improve efficiency, obtaining much lower RMS error for the same rendering time.

Figure 5 (Scene 3) has mostly diffuse surfaces, although the glass ball is specular, slightly dispersive, and has a wavelength dependent absorption function.

In Figure 6 (Scene 4) we demonstrate some of the caustic effects possible with a spectral renderer. Prisms built with wavelength dependent glass can be used to cast rainbows on walls, given a suitably broadband and collimated light source. Likewise, although we have not implemented this, rainbows in participating media could be produced by simply using the correct wavelength dependent phase (scattering) function.

This particular scene is somewhat unnatural in that almost all the light transport paths pass through objects that are specular and dispersive, since the light source is narrowly collimated. This scene is a worst case; almost all paths were degraded to a single wavelength, and some of the effects were sampled by only a small number of paths. The result is unfortunately rather noisy.

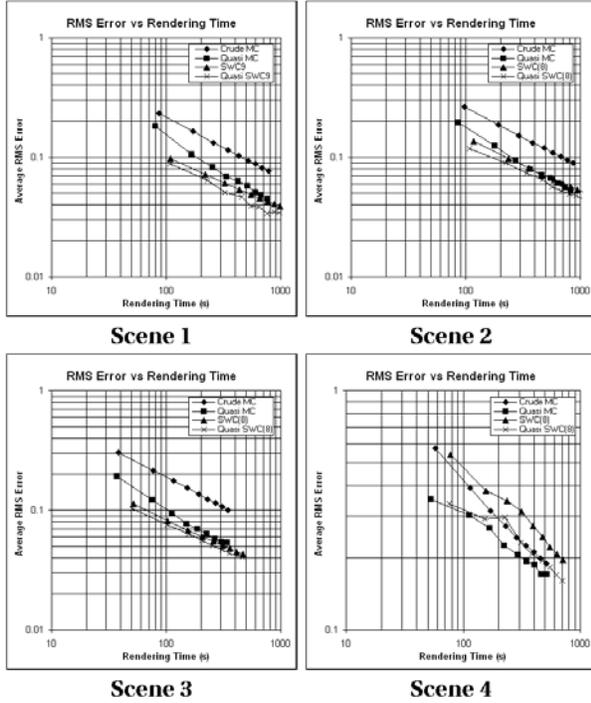


Figure 1: *Relative efficiency. RMS error versus rendering time for the scenes presented here, on log-log scales. The slope of the line on a log-log graph is an estimate of the exponent  $e$  of the asymptotic convergence in the form  $O(N^e)$ . Note the change of scale for Scene 4 (which is a worst case).*

Graphs of the average RMS error vs. CPU time for each of these scenes is shown in Figure 1. For Scenes 1 and 2 the best algorithm by far was quasi Monte Carlo sampling combined with stratified wavelength clusters. Eventually plain (single-wavelength per path) QMC catches up, but only at unreasonably high sampling rates. At low sampling rates, the majority of the reduction in the RMS error can be attributed to the introduction of wavelength clustering.

Scenes 3 and 4 test the degradation of the clustering algorithm in the presence of specular and dispersive materials. In Scene 3, the degradation is slight because most paths are not degraded to a single wavelength.

In Scene 4, however, simple QMC is the winner, and Crude Monte Carlo actually beats the algorithm with clusters, since the overhead of degradation does not occur.

It would be possible to automatically switch to one of the alternative strategies as needed. For example, if a pilot sample shows that most paths are specularly dispersive, then simple single-wavelength QMC can be used, or deferral. Likewise, if more complex scenes are used and it is found that computing surface intersections is much more time-consuming than computing reflectances, then the splitting strategy can be used.

We have also experimented with different numbers of wavelengths in a cluster. As shown in Figure 2, for Scenes 1 and 2 most of the improvement in performance was for the first three

wavelengths. There was very little benefit beyond nine wavelengths, and beyond fifteen wavelengths efficiency started decreasing. Most other scenes we tested had almost identical behaviour, except for Scene 4, in which a cluster size of one was optimal. An optimal cluster size of eight or nine is in agreement with previous observations that eight or nine spectral point samples can capture most spectral effects. However, our technique avoids aliasing, since all wavelength samples are also randomized.

## 7 Conclusions

A strategy for incorporating spectral phenomena into Monte Carlo renderers has been presented. This technique has little impact on performance when geometric and spectral phenomena are not deterministically related (for example, in the absence of specular dispersion), but improves colour accuracy.

A Monte Carlo global illumination algorithm such as bidirectional path tracing can split or degrade wavelength clusters only when needed to render spectral caustics, chromatic aberration, and other dispersive phenomena. The performance degradation is therefore proportional to the fraction of paths that encounter the dispersive material. However, the decrease in path generation efficiency is offset by the colour coherency (and lower source variance) typically generated by dispersion.

The clustering technique is applicable not only to bidirectional path tracers, but can also be applied easily to Monte Carlo ray tracers and potentially extended to other Monte Carlo global illumination algorithms. A renderer using Metropolis' light transport algorithm [21] could mutate one wavelength in a cluster at a time, or could mutate a cluster by warping it, for example. This would be particularly interesting in the case of "spiky" reflectance functions; by importance sampling using only illumination spectra, we implicitly assume that reflectance spectra are smooth.

In a naive extension of the photon map approach proposed by Jensen and Christensen [9], the bundle of wavelengths carried by a cluster would have to be stored for every hit. A  $k$ -nearest neighbour density estimation could be performed in the spectral domain as well as the spatial domain during the gather pass. Unfortunately, storing this amount of data might be infeasible.

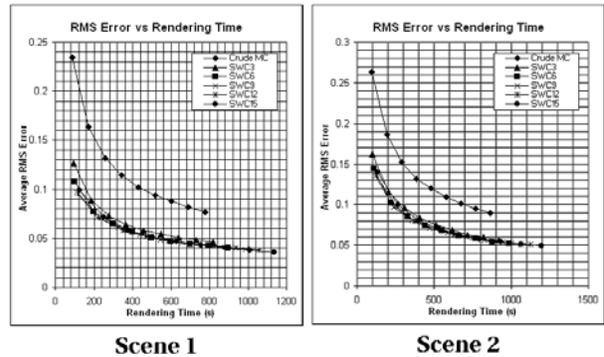


Figure 2: *RMS error vs. cluster size for Scenes 1 and 2, on a linear scale.*

Projection onto a fixed set of basis functions might be a more suitable approach in this circumstance.

More research needs to be done to determine how to implement spectral rendering effectively in these classes of rendering algorithms.

### Acknowledgements

We would like to acknowledge the help and advice of members of the Computer Graphics Lab at the University of Waterloo, especially Dan Milgram, Pete Harwood (who wrote part of the renderer), Ian Stewart (who implemented another part of the renderer) and Jan Kautz.

This research was sponsored by a grant from the National Research Council of Canada.

### 8 References

- [1] J. Arvo and D. Kirk. Particle Transport and Image Synthesis. In *Proc. SIGGRAPH*, pages 63–66, August 1990.
- [2] Commission Internationale de l’Eclairage. CIE Colorimetry Standard. Technical report, Central Bureau of the CIE, 1986.
- [3] W. B. Cowan. An inexpensive scheme for calibration of a colour monitor in terms of CIE standard coordinates. *Proc. SIGGRAPH (Computer Graphics)*, 17(3):315–321, July 1983.
- [4] A. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufman, 1995.
- [5] Andrew S. Glassner, Kenneth P. Fishkin, David H. Marimont, and Maureen C. Stone. Device-directed rendering. *ACM Transactions on Graphics*, 14(1):58–76, January 1995.
- [6] J. Gondek, G. Meyer, and J. Newman. Wavelength Dependent Reflectance Functions. In *Proc. SIGGRAPH*, pages 213–220, July 1994.
- [7] Roy Hall. *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, 1989.
- [8] X. He, K. Torrance, F. Sillion, and D. Greenberg. A Comprehensive Physical Model for Light Reflection. In *Proc. SIGGRAPH*, pages 175–186, July 1991.
- [9] H. Jensen and P. Christensen. Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps. In *Proc. SIGGRAPH*, pages 311–320, July 1998.
- [10] R. Victor Klassen. Modeling the effect of the atmosphere on light. *ACM Transactions on Graphics*, 6(3):215–237, 1987.
- [11] G. W. Meyer. Wavelength Selection for Synthetic Image Generation. *CVGIP*, 41:57–79, 1988.
- [12] D. Mitchell. Consequences of Stratified Sampling in Graphics. In *Proc. SIGGRAPH*, pages 277–280, August 1996.
- [13] M. Peercy. Linear Color Representations for Full Spectral Rendering. In *Proc. SIGGRAPH*, pages 191–198, August 1993.
- [14] Maria Raso and Alain Fournier. A Piecewise Polynomial Approach to Shading Using Spectral Distributions. In *Proc. Graphics Interface*, pages 40–46, June 1991.
- [15] R. Redner, M. Lee, and S. Uselton. Smooth B-Spline Illumination Maps for Bidirectional Ray Tracing. *ACM Transactions on Graphics*, 14(4), October 1995.
- [16] M. Schramm, J. Gondek, and G. Meyer. Light Scattering Simulations using Complex Subsurface Models. In *Proc. Graphics Interface*, pages 56–67, May 1997.
- [17] P. Shirley, C. Wang, and K. Zimmerman. Monte Carlo Techniques for Direct Lighting Calculations. *ACM Transactions on Graphics*, 15(1), January 1995.
- [18] S. W. Thomas. Dispersive refraction in ray tracing. *The Visual Computer*, 2(1):3–8, January 1986.
- [19] E. Veach and L. Guibas. Bidirectional Estimators for Light Transport. In *Fifth Eurographics Workshop on Rendering*, pages 147–162, June 1994.
- [20] E. Veach and L. Guibas. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *Proc. SIGGRAPH*, pages 419–428, August 1995.
- [21] E. Veach and L. Guibas. Metropolis Light Transport. In *Proc. SIGGRAPH*, pages 65–76, August 1997.
- [22] Ying Yuan, Toshiyasu L. Kunii, Naota Inamoto, and Lining Sun. Gemstone fire: Adaptive dispersive ray tracing of polyhedrons. *The Visual Computer*, 4(5):259–270, November 1988.

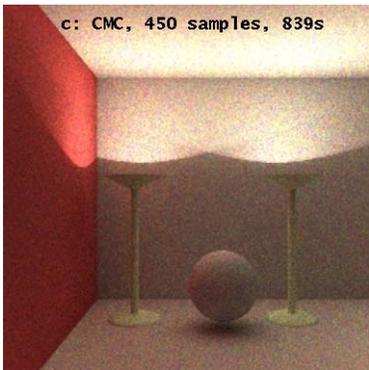
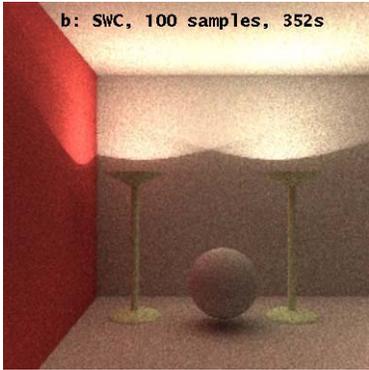
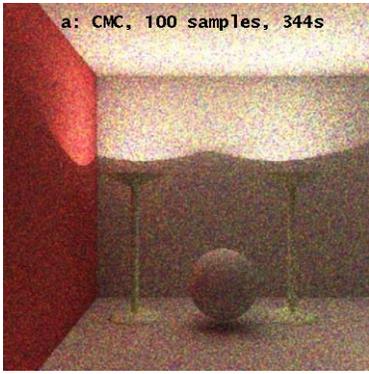


Figure 4: Scene 2: Diffuse environment, indirect illumination.

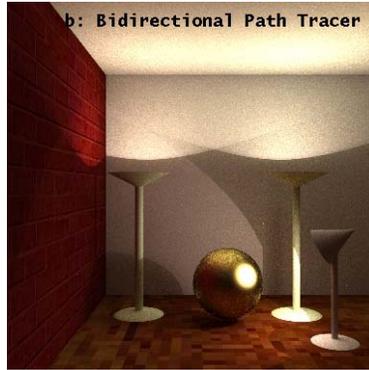
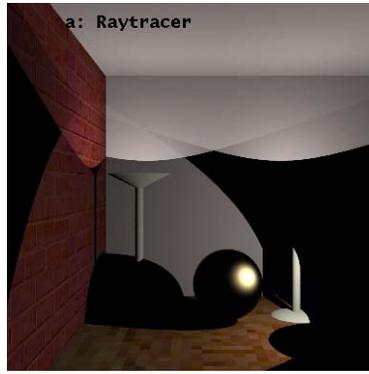


Figure 3: Scene 1: Raytracing vs. bidirectional path tracing.

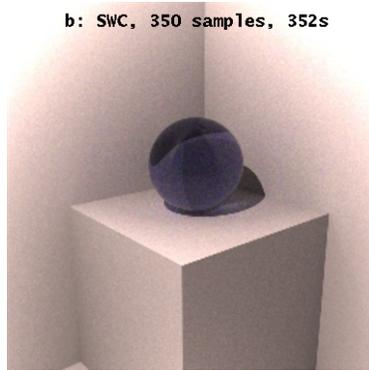
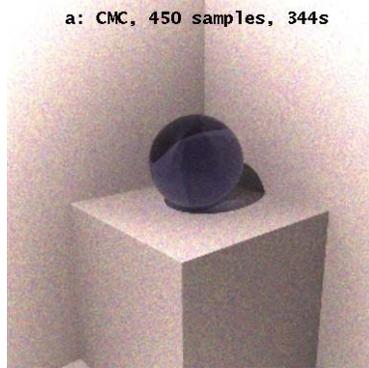


Figure 5: Scene 3: Dispersion and absorption.

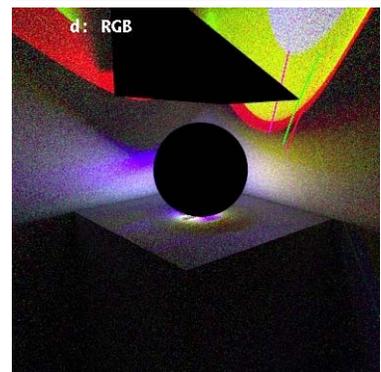
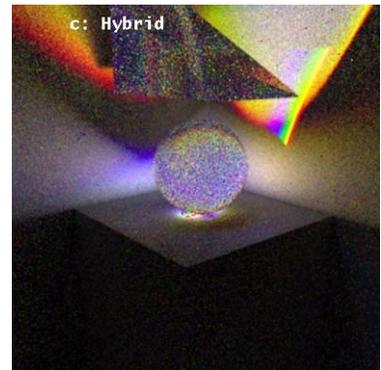
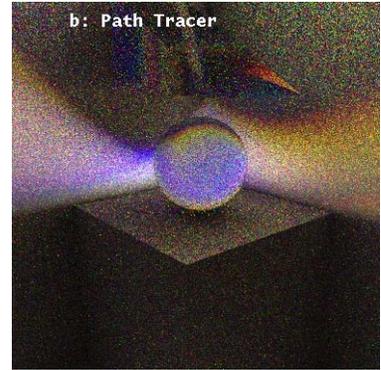
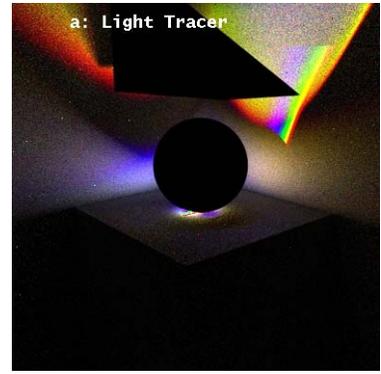


Figure 6: Scene 4: Spectral caustics cast by a prism and a sphere from a collimated broadband light source.