

Hero Wavelength Spectral Sampling

A. Wilkie^{2,1} and S. Nawaz¹ and M. Droske¹ and A. Weidlich¹ and J. Hanika^{3,1}

¹Weta Digital, New Zealand

²Charles University in Prague, Czech Republic

³Karlsruhe Institute of Technology, Germany




Figure 1: Wavelength-dependent subsurface scattering with single wavelength sampling (left) and our proposed method (right). The lower image halves are rendered with 4 paths per pixel, and the upper halves are rendered with 1024. For both sample counts, the hero wavelength images took only 3.5% longer than the single wavelength ones.

Abstract

We present a spectral rendering technique that offers a compelling set of advantages over existing approaches. The key idea is to propagate energy along paths for a small, constant number of changing wavelengths. The first of these, the hero wavelength, is randomly sampled for each path, and all directional sampling is solely based on it. The additional wavelengths are placed at equal distances from the hero wavelength, so that all path wavelengths together always evenly cover the visible range.

A related technique, spectral multiple importance sampling, was already introduced a few years ago. We propose a simplified and optimised version of this approach which is easier to implement, has good performance characteristics, and is actually more powerful than the original method. Our proposed method is also superior to techniques which use a static spectral representation, as it does not suffer from any inherent representation bias.

We demonstrate the performance of our method in several application areas that are of critical importance for production work, such as fidelity of colour reproduction, sub-surface scattering, dispersion and volumetric effects. We also discuss how to couple our proposed approach with several technologies that are important in current production systems, such as photon maps, bidirectional path tracing, environment maps, and participating media.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

Although a number of phenomena which are important for truly accurate prediction of object appearance, like e.g. metamerism, transmission colour, coloured scattering or fluorescence, can only be properly handled in a spectral renderer, a significant number of currently available produc-

tion renderers are still traditional RGB colour space systems. There are a few notable exceptions like Maxwell [Nex06] or Mitsuba [Jak10]. However, widespread adoption of such technology is hampered by the fact that industry workflows, and perhaps even more importantly, existing assets, are usually all exclusively in RGB. In spite of this, there is an increasing demand in the industry for realistic and predictive

rendering, e.g. for digital doubles or highly accurate colour matches between synthetic and real footage, which can only be provided by a spectrally based pipeline.

Most current spectral path tracers carry one spectral sample per path, which turns spectral sampling into just another dimension of the path space integral. This introduces considerable colour noise and is inefficient, as most paths are valid for more than just one wavelength. We show how to correctly use many wavelength samples per path, and how to efficiently incorporate this into a modern path tracer.

In section 2, we very briefly outline different approaches to spectral rendering, as well as the problem of wavelength dependent scattering functions. In section 3, we describe how we arrived at our solution, and in sections 4 and 5 we discuss technical implementation details and results.

2. State of the Art and Background

Monte Carlo Wavelength Sampling

The easiest way to incorporate spectral phenomena in a Monte Carlo path tracer is to use one wavelength per sampled path, which turns the spectral sampling into just another dimension of the path space integral formulation of the rendering equation. We can calculate the image by using an estimator

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

where X_i specifies a sample in path space and λ_i one sampled wavelength, $f(X_i, \lambda_i)$ is the measurement contribution and $p(X_i, \lambda_i)$ is the probability density of the sample, N is the number of sampled paths. Using one wavelength per sampled path introduces a lot of colour noise and in most cases it is quite wasteful, as each generated path is perfectly valid for all wavelengths if no spectral dependency is occurring during path propagation. Evans and McCool [EM99] make use of this fact by separating the spectral dimension and specifying a number of wavelengths to sample for one path.

$$\langle I \rangle = \frac{1}{N} \frac{1}{C} \sum_{i=1}^N \sum_{j=1}^C \frac{f(X_i, \lambda_i^j)}{p(X_i, \lambda_i^j)}, \quad (2)$$

where $p(X_i, \lambda_i^j)$ is the probability to sample the pair $[X_i; \lambda_i^j]$. If the path generation is independent of the wavelength, this probability reduces to

$$p(X_i, \lambda_i^j) = p(\lambda_i^j) \cdot p(X_i), \quad (3)$$

where $p(X_i)$ is the probability to sample the path specified by the sample X_i and $p(\lambda_i^j)$ is the probability to sample the wavelength λ_i^j . This formulation is used in e.g. in Luxrender [Ter07].

Predefined Wavelengths and Basis Functions

Instead of incorporating the spectral sampling into the path space integral formulation, an extension of the historic tristimulus method is to use n spectral samples ($n > 1$) taken at fixed, pre-determined wavelengths. On the positive side, this allows one to perform the most common operation on spectral data in a renderer, the calculation of the product of two spectra, via simple pairwise multiplication of the matching values of the respective sampling points. However, sampling at a fixed number of fixed frequencies introduces an error if not enough samples are used – and crucially, the magnitude of this error can be hard to determine a priori.

How many samples are used, and how they are placed (equidistant vs. distributed according to $V(\lambda)$) varies between implementations, and no consensus on how it should be done exists. Common sample counts are between at least 8 for rendering, and at least 32 samples for colour verification. Result precision, memory usage and rendering time are directly related to the number of samples [Hal99] [JF99] – but also to the properties of the input data, and to some degree also on which physical effects one supports in a given system. So a main research area was how to find a good number and distribution of samples [CW05, ZCB97, WTP00] – albeit with inconclusive results so far.

A spectrum can also be expressed as a combination of a set of basis functions. Using predefined wavelengths can be seen as using a set of box basis functions to represent the spectrum. Other basis functions like e.g. the Fourier basis [Pee93] or Gaussian [Mey88] can be used, or a specialized basis out of the set of spectra of a given scene [DF03, BDM09, DMCP94, RP97] can be extracted. The Fourier basis can work well for arbitrary scenes if only smooth spectra occur. However, data with abrupt peaks, like fluorescent lamp emission spectra, would require a large number of basis vectors, and therefore cannot be represented efficiently with this technique. Also, using more involved basis functions is comparatively slow and more complicated to implement. A hybrid approach was proposed in the composite spectral model of [SFD00]. It takes advantage of the fact that reflection spectra are usually smooth, and that spikes only occur in the spectra of some light sources: each spectrum is decomposed into its smooth component and a collection of spikes.

Both the hybrid and the plain basis function approach suffer from a fundamental issue: choosing basis functions for a particular spectrum usually yields good results, but finding one that are capable of representing *all* spectra in a given scene with only a few coefficients is hard.

Wavelength-Dependent Scattering Functions

One advantage of adding spectral sampling as another dimension of the path space integral is that the problem of finding the right number of spectral samples, or the right basis functions is completely avoided. It also automatically deals

properly with BSDFs or volume scattering functions that are wavelength-dependent. When using more than one wavelength per path or any basis function approach one wavelength has to be chosen for path propagation and the probabilities to generate the path with the other wavelengths can diverge.

Evans and McCool [EM99] suggest to overcome this problem by either splitting the path or choosing one of the wavelengths and discarding all the others: the former option reintroduces colour noise for these paths, and the latter potentially results in exponential path generation. Unfortunately spectral dependency is not rare in real scenes as the absorption in volumes and especially in skin is wavelength dependent: to give a typical example, figures 14-13 and 14-14 in [Ngu08] show typical values for a realistic human skin model. Using just a single coefficient to determine importance sampling for such strongly wavelength-dependent media will inevitably yield poor results.

Spectral Multiple Importance Sampling

Radziszewski et al. [RBA09] propose to use multiple importance sampling (MIS) to overcome this problem. They randomly select one of the wavelengths to use for path propagation and weight the path contribution by

$$w_i(X_i, \lambda_i^s) = \frac{p(X_i|\lambda_i^s)}{p(\lambda_i^s) \sum_{k=1}^C p(X_i|\lambda_i^k)}, \quad (4)$$

where λ_i^s is the selected wavelength and $p(X_i|\lambda_i^k)$ is the probability of sampling the path specified by X_i if the wavelength λ_i^k would have been used for path propagation.

Although their formulation results in the correct estimator it is not apparent which sampling techniques are weighted against each other to result in these MIS weights. In the next section, we explain the derivation of the MIS weights and show that their estimator relates to the correct probability densities for wavelength dependent paths. It can actually be explained by the interpretation of the balance heuristic described in [Vea97, Sec. 9.2.2.1].

We then discuss how a suitably simplified and streamlined version of this technique can actually be implemented with comparatively little effort, and can serve as an efficient solver for different flavours of path space sampling. In particular, we show how these MIS weights can be combined with other sampling techniques like next event estimation, bidirectional path tracing or photon maps.

3. Hero Wavelength Sampling

Most Monte Carlo path space samplers use single wavelength sampling. As outlined in section 2 adding the wavelength as another dimension of the path space integral is straightforward, unbiased, easy to implement and deals correctly with wavelength dependent phenomena. However, the

colour noise inherent to this approach is a highly objectionable artefact that is hard to remove via post-processing.

However, using the single wavelength path space sampler as the basis, one can implement a multiple wavelength estimator easily, by defining the first, randomly sampled wavelength as the *hero wavelength* that is used for path propagation decisions. As this wavelength is sampled in exactly the same way as in a single wavelength renderer, no additional random numbers have to be generated. The term is derived from movie industry jargon: shots in which a main character gets center stage are referred to as *hero shots* or *hero renders*. As the first wavelength gets center stage for all directional sampling decisions, it seems like an apt analogy.

The additional wavelengths are *not* sampled randomly: they are placed at equidistantly spaced locations around the hero wavelength instead, to evenly cover the visible spectral domain Λ . We define a rotation function $r_j : \Lambda \rightarrow \Lambda$,

$$r_j(\lambda_h) = (\lambda_h - \lambda_{\min} + \frac{j}{C}\bar{\lambda}) \bmod \bar{\lambda} + \lambda_{\min} \quad (5)$$

for $j = 1, \dots, C$, $\bar{\lambda} = \lambda_{\max} - \lambda_{\min}$, which we use to construct the C different wavelengths from one hero wavelength λ_h .

3.1. Combined Sample Density

When using multiple wavelengths per path, one reuses each spatial sample X_i on the path for more than one wavelength. If the sampling that generated X_i is wavelength dependent, this actually results in a different probability for the combined sample pair $[X_i; \lambda_i^j]$. As X_i is fixed for one set of wavelengths the probability to sample the pair $[X_i; \lambda_i^j]$ is

$$p(X_i, \lambda_i^j) = \frac{\sum_k p(\lambda_i^k) p(X_i|\lambda_i^k)}{C}, \quad (6)$$

where $p(\lambda_i^k)$ is the probability to sample λ_i^k and $p(X_i|\lambda_i^k)$ is the probability to sample X_i given the wavelength λ_i^k . The sum over the probabilities stems from the fact, that any of the C wavelengths could generate the path X_i and it would still result in the same pair $[X_i; \lambda_i^j]$ being evaluated. The sum is then multiplied by the probability of choosing this wavelength as the hero wavelength to generate the path $1/C$. Please note that this is true even if the additional wavelengths are sampled and not chosen deterministically. Most rendering systems that use more than one wavelength per path ignore this fact and use the probability

$$\bar{p}(X_i, \lambda_i^j) = p(\lambda_i^j) \cdot p(X_i|\lambda_i^j) \quad (7)$$

instead, which is only correct if the probability to sample the path specified by X_i is the same for all wavelengths of this path, i.e. no spectral dependency in the path generation.

3.2. MIS with shifted techniques

Even though the authors state a valid multiple wavelength sampling model in [RBA09], the connection to multiple im-

portance sampling is not derived directly, since the link of the weight and the sampling technique is missing. To verify the result we will define the sampling techniques, that are weighted against each other. This allows to apply different well-known weighting heuristics and to derive a simple extension for combining the proposed multiple wavelength technique with other sampling techniques, such as forward path tracing and next event estimation.

We define sampling technique $s = 1, \dots, C$ as follows:

- Sample λ according to $p(r_s^{-1}(\cdot)) =: p_{\lambda,s}$, where $r_s^{-1}(\lambda)$ returns the original hero wavelength of technique s before rotation.
- Sample X according to the probability density $p(\cdot|r_s^{-1}(\lambda)) =: p_{X,s}$.
- Set $p_s(X, \lambda) := p_{X,s}(X|\lambda) \cdot p_{\lambda,s}(\lambda)$

This basically means every time one of the C different wavelengths is sampled as the hero wavelength we define it as a different technique to sample the pair $[X_i; \lambda_i^s]$.

The definition of the weights for the balance heuristic is now straightforward

$$w_s(X, \lambda) := \frac{p_s(X, \lambda)}{\sum_{k=1}^C p_k(X, \lambda)}, \quad (8)$$

which yields the estimator

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^N \sum_{s=1}^C w_s(X_i^s, \lambda_i^s) \frac{f(X_i^s, \lambda_i^s)}{p_s(X_i^s, \lambda_i^s)}. \quad (9)$$

An important ingredient of our method is to use the same (Q)MC sequence for each of these techniques, in order to generate equidistantly sampled wavelengths and to sample the same path $X_i^s = X_i$ for each of these techniques. Furthermore, note that $p_s(X_i, \lambda_i^s) = p(X_i|\lambda_i^s) \cdot p(\lambda_i^s)$, where λ_i^h is the hero wavelength of the sampled path.

The MIS weight for the shifted wavelength sampling can be rewritten as

$$w_s(X_i, \lambda_i^s) = \frac{p_s(X_i, \lambda_i^s)}{\sum_{k=1}^C p_k(X_i, \lambda_i^s)} = \frac{p(\lambda_i^h) p(X_i|\lambda_i^h)}{\sum_{k=1}^C p(\lambda_i^k) p(X_i|\lambda_i^k)}, \quad (10)$$

as the probability to sample $p_s(X_i, \lambda_i^s)$ for technique s is the probability of the hero wavelength of this sample. Note that the index k is different in the denominator, letting the sum rotate over all techniques that can generate λ_i^s as a sample.

Inserting the definition of the MIS weight results in the estimator

$$\langle I \rangle = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C \frac{f(X_i, \lambda_i^j)}{\sum_{k=1}^C p(\lambda_i^k) p(X_i|\lambda_i^k)}. \quad (11)$$

We may recognize the connection to eq. (6) in [RBA09]. As described in [Vea97, Sec. 9.2.2.1] this corresponds to viewing the balance heuristic as the combined sample model described in section 3.1. However, with the MIS formulation we can easily switch to different weighting heuristics and

crucially, we can now combine these weights with other MIS weights from different path tracing techniques.

It is important to note that most applications of multiple importance sampling in rendering concern themselves with finding the best weighted combination of different sampling techniques, to favour those that have the best efficiency for any given path. In contrast, weighting is used here to allow additional path evaluations at different wavelengths without causing variance problems. If the sampling is highly wavelength dependent, such as in the extreme case of dispersion via a Dirac interface, our method degenerates to using only the hero wavelength. If, on the other hand, the shape of the integrand only varies slightly with wavelength, as is the case for most materials, the shifted techniques are still efficient techniques, so that we obtain an almost equal weighting for all bands and fully benefit from sampling multiple wavelengths.

It is also noteworthy that our spectral weight formulation could be applied to conventional spectral renderers that use n fixed sampling wavelengths. We did not investigate this in detail, as this would retain the spectral aliasing issues that random selection of the hero wavelength entirely removes, without offering any conceivable benefits. But this insight can still be useful to improve legacy spectral rendering codebases to handle arbitrary scattering media correctly: in such systems, the hero wavelength actually does not necessarily have to be one of the n sample wavelengths.

3.3. Combination with Path Tracing Techniques

Using the simple definition of these shifted techniques makes it straightforward to construct a combined multiple importance sampling estimator that incorporates also other path tracing techniques, such as bidirectional path tracing and next event estimation.

Let us denote by \mathcal{T} a set of such techniques with probabilities p_t , $t \in \mathcal{T}$, for each of which a shifted variant $p_{t,s}$ is defined accordingly. Analogously to the previous section, this results in the following MIS weight for technique t' :

$$\frac{p_{t',s}(X_i, \lambda_i^s)}{\sum_{t \in \mathcal{T}} \sum_{k=1}^C p_{t,k}(X_i, \lambda_i^s)} = \frac{p_{t'}(X_i, \lambda_i^h)}{\sum_{t \in \mathcal{T}} \sum_{k=1}^C p(\lambda_i^k) p_t(X_i|\lambda_i^k)} \quad (12)$$

Combining all techniques into a single estimator in contrast to weighting the techniques in \mathcal{T} becomes crucial if the techniques vary with wavelength. For example, consider the setup shown in figure 2, which uses forward path tracing and next event estimation, where the latter is importance sampled based on wavelength. If the weighting would only take into account the hero wavelength, the secondary wavelengths which are carried with the path that hit discs with their wavelength resulting in strong contributions with forward path tracing, would not be weighted down properly since the probability of sampling the direction in the hero




Figure 2: Joint MIS with wavelength dependent light sampling. From left to right: Environment used to light the scene, using MIS based on the hero wavelength only, single wavelength sampling, hero wavelength path tracing with joint MIS.

wavelength is low. Naturally, in such an extreme case, we cannot expect any advantage from hero wavelength sampling. However, by using the weighting strategy shown in equation (12), hero wavelength sampling extends robustly to multiple techniques.

3.4. Number of Additional Wavelengths

For performance reasons, we made the engineering choice to always go with exactly four wavelengths per path – the hero wavelength, and three additional ones. This allows all BSDF and scattering computations that are performed at each path vertex to be natively vectorised. As SSE floating point instructions use only a few clock cycles more than normal float operations, this keeps the computational overhead introduced by the additional wavelengths to a minimum. Also, this is a simplification compared to [RBA09], where computational effort is invested into figuring out how many samples one should take. The system can be trivially re-compiled to work with multiples of 4, but as we will show in section 5, the original choice of 4 proved to be sufficient in practice.

4. Implementation Details

We now describe implementation details needed to get our approach to work in a nontrivial rendering system.

4.1. Sampling and Weight Computation

Since we base all sampling decisions during path propagation on hero wavelength PDFs, it is *not* sufficient to just replace all single wavelength references and computations in the rendering codebase with SSE-compatible 4-vectors.

To compute the weight (10), the sampling probabilities have to be computed for all four wavelengths. All BSDF and scattering code is vectorised via SSE, so these are obtained for similar cost as a single probability. But the actual sampling weights require more care, since typical BSDF implementations compute the ratio of the function value divided by the probability for the same wavelength. But we have to make sure that we divide by the probability which generated the path, i.e. the hero wavelength probability.

Sampling Dirac distributions requires only some minor additional logic, to ensure that paths that are only possible for the hero wavelength only use the first spectral channel, and other contributions are blanked out. On the other hand contributions through Diracs which are aligned in other wavelengths are still distributed over all channels (cf. Fig. 6).

4.2. Photon Map Queries

An issue with all systems that work with random wavelength sampling is that photon map reconstruction operations are considerably less efficient than in systems which use a fixed representation for light. If all photons carry RGB information, all photons within a given radius can be used for a particular reconstruction that is taking place anywhere. The same applies if a fixed representation with the same n wavelengths is used throughout. However, if all photons carry different wavelengths (i.e. if they are the result of a single wavelength photon tracing pass), and if the path tracing pass also operates in a mode that it only works with one single wavelength, only a small subset of all photons are available for each reconstruction operation. Which greatly increases noise and decreases coherency. This situation is not made much better if one adds more random wavelengths to each path, as discussed in [RBA09]. In particular, if a photon is found, usually only a single one of the wavelengths stored in it can be used for any given reconstruction in such a system.

One of the the key benefits of our scheme of stratified wavelength selection is that it increases the chances of finding useful photons for a given reconstruction three-fold: given a hero wavelength for an eye path that requests a photon map evaluation, we can test all four wavelengths contained in each photon if they are close enough to the hero wavelength: usually a few nanometers distance are considered acceptable. If any of the photon wavelengths matches, the channels of the photon are cyclically rotated, so that the matching wavelength becomes the photon hero wavelength. This automatically ensures that the additional channels match the additional channels of the path that requested the reconstruction. In this way, one can use the *entire* information contained in the photon for reconstruction purposes,

which considerably reduces variance. As figure 4 demonstrates, the number of photons that are shot can be cut to a quarter compared to single wavelength rendering, with no loss in quality.

4.3. Light Source Sampling

A more subtle issue is that direct sampling of light sources has to be handled properly as well. If area light sources with varying spectral distributions (e.g. environment maps) are only importance sampled according to their luminance pattern, this is sub-optimal, but does not cause immediate issues. However, if the more optimal approach of importance sampling a given area light is performed separately for each wavelength, and if the environment exhibits strongly different patterns according to wavelength, this has to be taken into account. Figure 2 shows a particularly extreme case, in which separate regions in the environment map emit in the short, medium and long wavelength domain. Naive sampling causes excessive variance, and it takes the combined multiple importance sampling, which is made possible by the formulation we derive in section 3.2, to bring convergence of the hero wavelength sampling technique to the levels of single wavelength approach.




Figure 3: Glossy reflections of coloured area lights. Top: single wavelength, bottom: hero wavelength, both 16 spp; the insets were rendered with 1024 spp. Even at low sample numbers far from convergence colour noise is eliminated by the hero wavelength approach due to the stratified sampling of the spectral domain. Also note the weak but still noticeable residual colour noise in the single wavelength 1024 sample insets.




Figure 4: Single (left) vs. hero wavelength (right) renders of a photon map scene with a non-dispersive glass sphere. 4 samples per pixel, inset: 256 samples. To render the hero wavelength version, only a quarter of the photons used for the single wavelength scene was shot: due to the increased ability of the renderer to find photons that can be used, no decrease in quality can be seen.

4.4. Volumes

Two main sampling techniques of path tracing in volumes are scattering and sampling a distance for a given ray direction. An efficient technique for sampling distances along a ray is to importance sample with a probability as close as possible to the attenuation

$$\tau(x, x', \lambda) = e^{-\int_0^s \sigma_t(x(s), \lambda) ds}. \quad (13)$$

In volumes in which the extinction only depends on λ but not on x , the transmittance can be perfectly importance sampled and the probability directly computed. Hence, it is straightforward to compute sampling weights and probabilities for all relevant wavelengths.

This is more tricky for volumes with heterogeneous, i.e. spatially changing extinction coefficients. Even though the distance can be efficiently importance sampled for a given wavelength using Woodcock tracking [RSK08], the probability and the evaluation of the integrand now include the transmittance term, which is an unknown integral.

To compute the sampling weights and probabilities for all channels we may resort to numerical integration of the transmittance parallel to the Woodcock tracking computation, but this introduces some bias due to integration error. If the method can be accepted to be unbiased, we found that with a sufficient number of integration steps to achieve convincing results (see Fig. 5) even though rendering of heterogeneous volumes with multiple wavelengths per path has so far been discouraged in the literature (cf. [RSK08], Sec. 3.2).

5. Performance / Results

In this section, we present results from several scenarios that can cause problems with spectral rendering, and demonstrate how the proposed technique improves the quality of the results. We do not give absolute performance figures for any

of these images, as too many factors are involved for absolute rendering times to be meaningful. We also do not give individual relative timings of single vs. hero wavelength renderings. The reason for this is that the performance penalty for the proposed hero wavelength extension was always in the range of 3% – 10%, if wall clock times of renderings with similar sample counts were compared.

A small performance decrease with respect to how many paths can be traced for a given wall clock time interval is expected: SSE instructions are slightly slower than normal floating point instructions, more data has to be handled which increases memory bus pressure, and the spectral MIS described in section 4.1 adds a small overhead. However, convergence speed for a given number of paths is considerably faster, as colour noise is eliminated by the hero wavelength approach. In all scenes we encountered, this gain in efficiency more than compensated the slight slowdown with respect to the numbers of paths one can trace per unit time.




Figure 5: Single (top) vs. hero wavelength (bottom) for heterogeneous volumes, 16 samples per pixel. Note how the colour of the volumes is already stable even for low sample counts with hero wavelength rendering. The converged insets are rendered with 1024 samples. The seeming brightness difference between the insets and their surround arises from the large amount of high energy fireflies in the 16 sample images, which are cut off by the tone reproduction operator.

When choosing our test cases, we attempted to systematically cover scene types that can be problematic in practice:

- Figures 1 and 7: sub-surface scattering in a medium

with a strongly wavelength-dependent scattering function (in this case, skin). This demonstrates a key point: namely, that the hero wavelength technique provides tangible gains for real materials with strongly divergent PDFs at different wavelengths.

- Figure 3: no colour noise even for low sample counts. Getting stable colours for non-converged preview images is an important feature during scene development.
- Figure 4: photon tracing only requires a quarter of the photons needed in a normal single wavelength renderer, with no loss in quality (section 4.2).
- Figure 5: ability to handle heterogeneous scattering and emissive media (section 4.4).
- Figure 6: dispersion. Demonstrates that our technique behaves as well as one can expect for an unbiased technique: for paths that traverse dispersive Dirac interfaces, it regresses to single wavelength behaviour, while retaining colour noise elimination for all others.
- Figure 2: efficient sampling of area lights even if they have a strongly divergent spectral distribution (section 4.3).

For all tests we conducted, our results were consistent: the hero wavelength extension yielded significantly better results for the same elapsed wall clock time, mainly through the complete elimination of colour noise, and did not cause any degradation in scene quality. In particular, and somewhat surprisingly, it even yielded significant gains in scenes where paths have to traverse media such as skin, with scattering functions that strongly vary with wavelength. The only case in which the technique gracefully regressed to single wavelength behaviour was dispersion through Dirac interfaces: and in this case, it does not perform worse than the original technique, and still eliminates colour noise in all parts of the scene that are not affected by dispersion.

6. Conclusion

In this paper, we have presented a spectral rendering technique that we believe to be a stable, workable solution with manifest advantages over existing techniques. In particular, it does not suffer from colour noise, while still avoiding the sampling bias found in techniques that sample input spectra at a number of fixed wavelengths. Crucially, just like simple random wavelength sampling, our technique still allows efficient sampling of arbitrary, wavelength-dependent distributions. In contrast to existing spectral MIS approaches, it is directly usable with photon mapping techniques, which means that it integrates well into multi-pass global illumination techniques.

The only small drawback of the technique, the fact that the spectral data present in a scene has to be explicitly evaluated for all randomly chosen wavelengths, adds little overhead, and is intrinsic to all techniques that avoid spectral sampling bias by accurately evaluating input spectra.




Figure 6: Two glass spheres, rendered with 16 samples per pixel: single wavelength (top) and hero wavelength (bottom). Hero wavelength transport helps where it can (e.g. the reflections seen through the left sphere, and the caustics), and as expected falls back to single wavelength behaviour in case of wavelength dependent refraction through a Dirac interface. The converged insets on the right sphere are rendered with 1024 samples.

7. Acknowledgements

We wish to thank Jonathan Swartz for providing figure 5, and Martin Šik for his valuable comments on the mathematical formalisms used in the paper.

This work was partially supported by the Czech Science Foundation grant P202-13-26189S.

References

- [BDM09] BERGNER S., DREW M. S., MÖLLER T.: A tool to create illuminant and reflectance spectra for light-driven graphics and visualization. *ACM Trans. Graph.* 28, 1 (2009), 1–11. 2
- [CW05] CHERN J.-R., WANG C.-M.: A novel progressive refinement algorithm for full spectral rendering. *Real-Time Imaging* 11, 2 (2005), 117–127. 2
- [DF03] DREW M. S., FINLAYSON G. D.: Multispectral processing without spectra. *J. Opt. Soc. Am. A* 20, 7 (2003), 1181–1193. 2
- [DMCP94] DEVILLE P. M., MERZOUK S., CAZIER D., PAUL J.-C.: Spectral data modeling for a lighting application. *Comput. Graph. Forum* 13, 3 (1994), 97–106. 2
- [EM99] EVANS G. F., MCCOOL M. D.: Stratified wavelength




Figure 7: Closeups of the thumb of the hand scene in Figure 1. From left to right and top to bottom: single wavelength transport, hero wavelength no MIS, balance heuristic, power heuristic, and finally bottom row: single wavelength and hero wavelength at 1024 spp. Note the residual colour noise even at 1024 spp in the single wavelength rendering.

clusters for efficient spectral Monte Carlo rendering. In *Graphics Interface* (1999). 2, 3

- [Hal99] HALL R.: Comparing spectral color computation methods. *IEEE Comput. Graph. Appl.* 19, 4 (1999), 36–45. 2
- [Jak10] JAKOB W.: Mitsuba renderer. <http://www.mitsuba-renderer.org>, 2010. 1
- [JF99] JOHNSON G. M., FAIRCHILD M. D.: Full-spectral color calculations in realistic image synthesis. *IEEE Comput. Graph. Appl.* 19, 4 (1999), 47–53. 2
- [Mey88] MEYER G. W.: Wavelength selection for synthetic image generation. *Comput. Vision Graph. Image Process.* 41, 1 (1988), 57–79. 2
- [Nex06] NEXT LIMIT: Maxwell render. <http://www.maxwellrender.com/>, 2006. 1
- [Ngu08] NGUYEN H. (Ed.): *GPU gems*, vol. 3. Addison-Wesley, 2008. 3
- [Pee93] PEERCY M. S.: Linear color representations for full speed spectral rendering. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques* (1993), pp. 191–198. 2
- [RBA09] RADZISZEWSKI M., BORYCZKO K., ALDA W.: An improved technique for full spectral rendering. *Journal of WSCG* 17, 1-3 (2009), 9–16. 3, 4, 5
- [RP97] ROUGERON G., PEROCHE B.: An adaptive representation of spectral data for reflectance computations. In *Proceedings of the Eurographics Workshop on Rendering Techniques* (1997), pp. 127–138. 2



Figure 8: A hairball rendered with single wavelength (top) and hero wavelength (bottom) at 4, 16, and 256 samples per pixel in each slice (left to right).

- [RSK08] RAAB M., SEIBERT D., KELLER A.: Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer Berlin Heidelberg, 2008, pp. 591–605. 6
- [SFD00] SUN Y., FRACCHIA F. D., DREW M. S.: Rendering light dispersion with a composite spectral model. In *1st International Conference on Color in Graphics and Image Processing* (Nov. 2000), pp. 51–56. 2
- [Ter07] TERRENCE VERGAUWEN: Luxrender. <http://www.luxrender.net/>, 2007. 2
- [Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997. 3, 4
- [WTP00] WILKIE A., TOBLER R. F., PURGATHOFER W.: Ray-tracing of dispersion effects in transparent materials. In *WSCG* (2000). 2
- [ZCB97] ZEGHERS E., CARRE S., BOUATOUCH K.: Error-bound wavelength selection for spectral rendering. *The Visual Computer* 13, 9+10 (1997), 424–434. 2