# A Frequency Analysis and Dual Hierarchy for Efficient Rendering of Subsurface Scattering

David Milaenen
Weta Digital
Université de Montréal

Laurent Belcour
Unity Labs
Université de Montréal

Jean-Philippe Guertin
Université de Montréal

Toshiya Hachisuka
University of Tokyo

Derek Nowrouzezahrai
McGill University
Université de Montréal

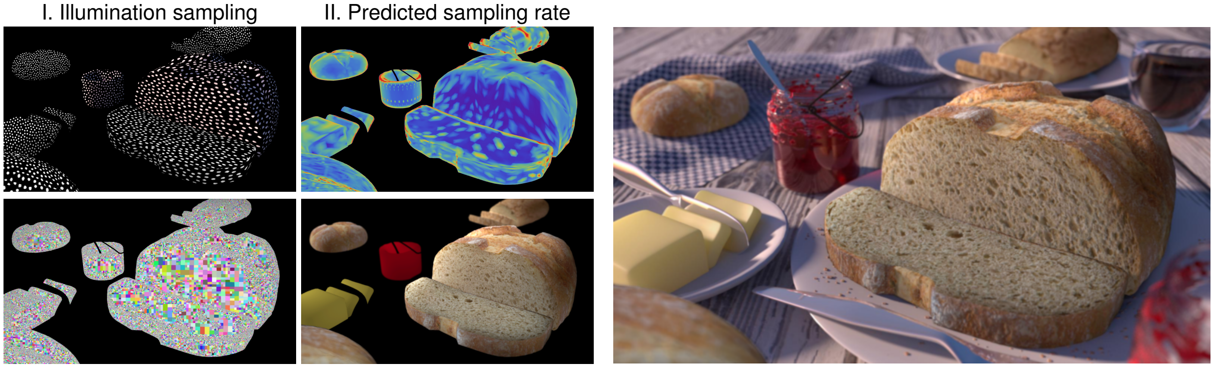


Figure 1: We introduce a hierarchical method to accelerate the rendering of multiple scattering with BSSRDFs (IV). We overview our approach in the PICNIK scene, above: our frequency analysis of BSSRDFs allows us to predict the screen-space sampling rates (II) which are used to devise bounds on the variation of outgoing radiance. These bounds allow us to efficiently integrate the BSSRDF using a dual hierarchy over clustered illumination samples (I) and shading points (i.e., pixels; III).

## Abstract

BSSRDFs are commonly used to model subsurface light transport in highly scattering media such as skin and marble. Rendering with BSSRDFs requires an additional spatial integration, which can be significantly more expensive than surface-only rendering with BRDFs. We introduce a novel hierarchical rendering method that can mitigate this additional spatial integration cost. Our method has two key components: a novel frequency analysis of subsurface light transport, and a dual hierarchy over shading and illumination samples. Our frequency analysis predicts the spatial and angular variation of outgoing radiance due to a BSSRDF. We use this analysis to drive adaptive spatial BSSRDF integration with sparse image and illumination samples. We propose the use of a dual-tree structure that allows us to simultaneously traverse a tree of shade points (i.e., pixels) and a tree of object-space illumination samples. Our dual-tree approach generalizes existing single-tree accelerations. Both our frequency analysis and the dual-tree structure are compatible with most existing BSSRDF models, and we show that our method improves rendering times compared to the state of the art method of Jensen and Buhler [18].

**Index Terms:** Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—Ray Tracing

## 1 Introduction

Including subsurface scattering effects in virtual scenes can significantly increase the realism of rendered images. Since many real-world materials exhibit subsurface scattering effects, modeling and simulating them remains an important problem in realistic image synthesis.

Accurate light transport in highly absorbing media can be modeled mathematically with the Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF). Many BSSRDF models exist with varying degrees of accuracy: classical dipole models [9, 19] and quantized diffusion [11] do not account for the angular variation of incident radiance, however more recent models do [10, 14, 16]. Unlike BRDFs, BSSRDFs describe light transport between two *different* locations on an object. As such, an additional spatial integration (over the surface) is required in order to render objects with BSSRDFs. Jensen and Buhler [18] introduced an adaptive hierarchical integration method to amortize the cost of this spatial integration using clusters of spatial illumination samples. While this approach has been successfully used in many applications, it does not take the smoothness of the resulting outgoing radiance (i.e., in screen-space) into account.

We propose a novel integration method that clusters **both** pixels and illumination points as illustrated in Figure 1. We conduct a frequency analysis of subsurface scattering that is agnostic to the underlying BSSRDF model. Specifically, we study the frequency content of the spatial and angular variation of radiance after its BSSRDF interaction. This leads us to a theoretically sound criterion for sparse sampling and adaptive integration. Using this criterion, we leverage a dual hierarchical data structure to accelerate the final evaluation of the multiple scattering term. Our hierarchical evaluation is motivated by the existing tree-based approach of Jensen and Buhler [18]; our dual-tree structure, however, amortizes computation cost across both pixels *and* illumination points. We are able to generate higher-quality results in less rendering time compared to the single tree method of Jensen and Buhler [18]. Concretely, we propose:

- a frequency analysis of shading with BSSRDFs,
- a numerical approach for estimating the BSSRDF spectra, which we use to determine the variation of outgoing radiance

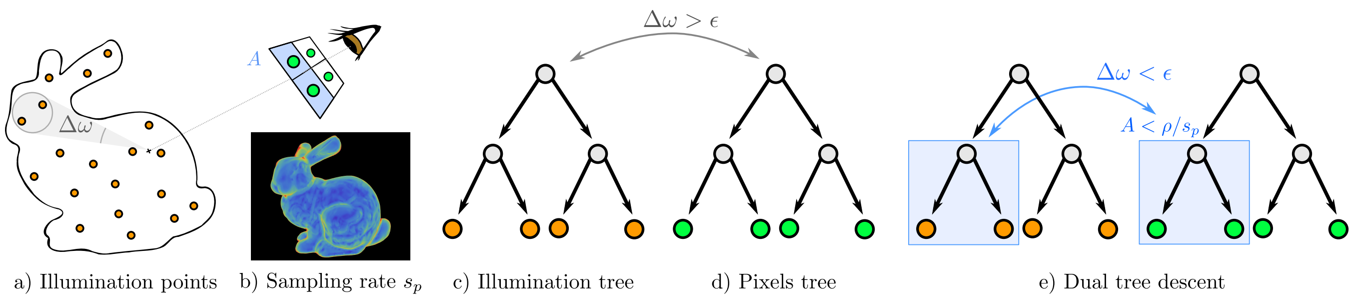


Figure 2: We sample incident illumination over the object (a) according to its subsurface scattering properties and construct two spatial acceleration structures: one over these samples (c) and one over pixels (d). To render, we simultaneously traverse the trees (e), using our outgoing radiance bandwidth estimate $s_p$ (b) to stop the tree traversal and shade super-pixels of area $A$.

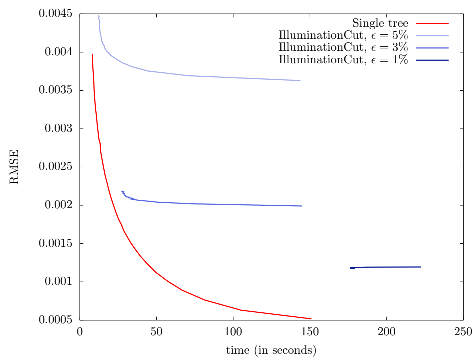


Figure 3: We compare IlluminationCut [8] to the method of Jensen & Bulher [18] on the BUNNY scene with various $\epsilon$ error bound settings. The cost of computing the upper-bound metric [23, Eq. 13], which requires multiple BSSRDF evaluations, precludes the direct applicability of IlluminationCut to adaptive BSSRDF shading.

over the surface of a translucent object, capable of supporting any underlying dipole model, and

- the application of a dual-tree structure to the problem of BSSRDF estimation in joint image- and object-space, directly leveraging our frequency analysis to adaptively traverse the structure and accelerate the final rendering.

## 2 Previous Work

We focus on work that most closely aligns with our approach: specifically, we review integration schemes for BSSRDF models, and frequency analyses of light transport.

BSSRDF Integration Techniques. In all cases, the bottleneck of dipole-like techniques remains the numerical evaluation of the spatial-angular integration in Equation 1. Jensen and Buhler [18] compute an approximate evaluation of this contribution from sparse irradiance samples distributed over a translucent object's surface. Here, the outgoing radiance at any shade point is computed by traversing a tree over the irradiance samples and terminating traversal according to a quality criterion. This two-pass approach introduces a controllable bias and remains compatible (often without modification) with many of the newer dipole models we discussed in Section . Notably, Frisvad et al. need only substitute the (diffuse) irradiance samples with a vector of differential irradiance samples, and d'Eon and Irving use a supplemental 1D radial directional radiance bin.

We are motivated by the lack of techniques that fully leverage image-space coherence to reduce the computation time of rendering translucent materials. Approaches based on LightCuts [24] fail to either efficiently treat BSSRDFs or amortize computation cost across similar pixels. Arbree et al. [1] propose a scalable approach to rendering large translucent scenes, based on multidimensional lightcuts [23], aggregating the computation of irradiance samples by simultaneously clustering lights and irradiance samples. Their clustering is designed to approximate the resulting contribution at a given shade point. While this methods also uses two trees, it treats each pixel **independently** and does not take the resulting image smoothness into account (see Figure 4 and Section 4.1 of [1]). We do not consider the evaluation cost of (ir)radiance samples, but we do cluster evaluation over pixels. In contrast, multidimensional-LightCut methods, such as IlluminationCut [8], could (in theory) be extended to BSSRDF shading but, in doing so, would require a prohibitive number of BSSRDF evaluations to evaluate their error threshold; indeed, we implemented such an extension of IlluminationCut to validate this claim (see Figure 3). On the other hand, our technique can also be used to provide a frequency-based cut threshold specifically designed for BSSRDFs, all while avoiding any explicit evaluation of the BSSRDF model.

The idea of applying a doubly-adaptive traversal originates from the particle simulation literature [15], and the implementation of d'Eon and Irving's quantized diffusion model [11] in Pixar's RenderMan implicitly leverages a similar principle (i.e., with REYES' adaptive micropolygon evaluation). One of our contributions is a well-founded oracle to terminate shading tree traversal based on our BSSRDF frequency analysis. Similarly, Jarabo et al. [17] leverage trees over virtual point lights and shading points without explicitly using dual trees, but their algorithm is specifically restricted to VPLs and diffuse materials, making the transition to translucent surfaces difficult.

Frequency Analyses of Light Transport. Durand et al. [13] presented the first comprehensive Fourier analysis of light transport in scenes with opaque surfaces, and a proof-of-concept adaptive image space sampling approach to reconstruct noise-free images at super-pixel sampling rates. Dubouchet et al. [12] use frequency analysis to construct a sampling cache which improves efficiency when rendering animations using distant direct lighting. Bagher et al. [2] derived atomic operators for bandwidth estimation in order to study environmental reflection with acquired BRDFs. Belcour et al. extend these frameworks to incorporate the study of defocus and motion blur [5], scattering in arbitrary participating media [3], and for global illumination [6], but do not directly tackle dense media or BSSRDFs. We bridge this gap with a frequency analysis of scattering in dense media, similarly leveraging matrix-vector formulations of frequency-space bandwidth operators.

### 2.1 Overview

Figure 2 overviews our approach: after sparsely evaluating incident radiance on the surface of each translucent object (Figure 2a), we compute a per-pixel bandwidth estimate of the multiply-scattered outgoing radiance (Figure 2b). We build two spatial acceleration structures, one over illumination samples (Figure 2c) and another

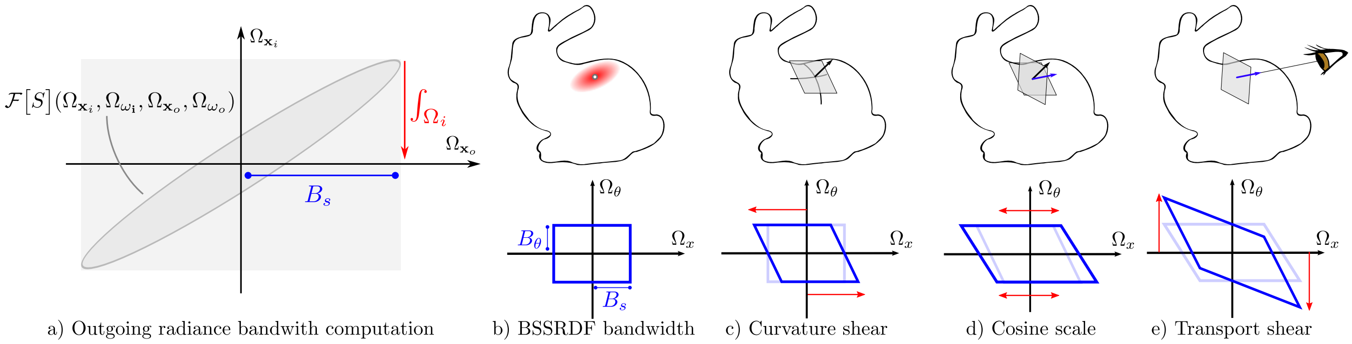


Figure 4: Assuming that the incoming light-field has infinite bandwidth, we estimate the bandwidth of the outgoing light-field $[B_s, B_\theta]$ as the bandwidth of the BSSRDF along the outgoing spatial positions and directions (a). The interaction with the material limits the spectrum of the local light-field by the BSSRDF spatial and angular bandwidth (b). To estimate the bandwidth at the camera position, we first shear spatially the spectrum to account for curvature (c). Then, we scale spatially to account for foreshortening (d) and finally shear angularly the spectrum to account for transport (e).

over pixels (Figure 2d). In order to compute the object's final shading, we simultaneously traverse both trees, hierarchically accumulating the contribution of groups of illumination samples to groups of pixels (Figure 2e). We use the frequency bandwidth of the outgoing radiance predicted by our theory (Section 3) to terminate traversal along each tree, significantly reducing the number of BSSRDF evaluations necessary to compute the final image without introducing visible artifacts.

We present our BSSRDF frequency analysis theory, as well as its numerical realization for computing image-space radiance bandwidths in Section 3. We introduce our variant of the dual tree construction and how the bandwidth predictions are used during hierarchical traversal in Section 4. Finally, we discuss our implementation details in Section 5 and compare our method to the state of the art in Section 6.

## 3 Fourier Analysis

We will derive conservative, numerical estimates of the frequency bandwidth of the outgoing radiance in image space, taking into account the effects of curvature, foreshortening, transport and multiple scattering on the incident light field's frequency content. We will show that the BSSRDF acts as a band-limiting filter on the incident radiance distribution, and we will derive an expression of the resulting spatio-angular bandwidth of the outgoing radiance spectrum (Section 3). We will use these bandwidth estimates, combined with the formulation of Bagher et al. [2], to predict the variation of outgoing radiance in image space (Section 3.2), which will in turn drive our hierarchical dual tree traversal and integration (Section 4).

### 3.1 Fourier Transform of a BSSRDF

Given a BSSRDF model $S(\mathbf{x}_i, \omega_i, \mathbf{x}_o, \omega_o)$, the outgoing radiance at the object surface $L_o$ in direction $\omega_o$ and at position $\mathbf{x}_o$ is expressed as:

$$L_o(\mathbf{x}_o, \omega_o) = \iint_{\mathcal{A}\times\mathcal{H}} S(\mathbf{x}_i, \omega_i, \mathbf{x}_o, \omega_o)\, L_i(\mathbf{x}_i, \omega_i)\, d\omega_i^{\perp}\, d\mathbf{x}_i, \qquad (1)$$

where $\mathcal{A}$ is the object's surface area, $\mathcal{H}$ is the set of (hemispherical) incident directions, $L_i$ is the incident radiance, and $d\omega_i^{\perp} = \cos\theta_i\, d\omega_i$ is the projected solid angle.

If we apply a Fourier transform to Equation 1, converting products in the primal domain to convolutions in the frequency domain and integration in the primal domain to DC evaluation in the frequency domain, we obtain:

$$\mathcal{F}\left[L_o\right](\Omega_{\mathbf{x}_o}, \Omega_{\omega_o}) = \left[\mathcal{F}\left[\widehat{S}\right] \circ \mathcal{F}\left[L_i\right]\right](0, 0, \Omega_{\mathbf{x}_o}, \Omega_{\omega_o}), \qquad (2)$$

where $\mathcal{F}\left[f\right]$ is the Fourier transform of $f$, $\circ$ the convolution operator, and $\Omega_x$ the frequency variation of $x$. Concretely, the outgoing radiance's spatial-angular frequency spectrum $\mathcal{F}\left[L_o\right](\Omega_{\mathbf{x}_o}, \Omega_{\omega_o})$ results from evaluating the convolution of the Fourier transform of the cosine-weighted BSSRDF $\mathcal{F}[\widehat{S}] = \mathcal{F}\left[S(\mathbf{x}_i, \omega_i, \mathbf{x}_o, \omega_o)\cos(\theta_i)\right]$ with the Fourier transform of the incident light $\mathcal{F}[L_i]$ at the incoming spatial and directional DC frequencies $(\Omega_{\mathbf{x}_i}, \Omega_{\omega_i}) = (0, 0)$.

Assuming that $\mathcal{F}\left[L_i\right]$ contains all-frequency content, the resulting outgoing bandwidth (along $\Omega_{\mathbf{x}_o}$ and $\Omega_{\omega_o}$) after convolution against the spectrum of the cosine-weighted BSSRDF $\mathcal{F}[\widehat{S}]$ will match the bandlimit of $\mathcal{F}[\widehat{S}]$ (see Figure 4a). We will discuss how to compute the spatial and angular bandwidths $\{B_o, B_\theta\}$ of the cosine-weighted BSSRDF given its local orientation.

**Spatial Bandwidth.** We compute the cosine-weighted BSSRDF's spatial bandwidth numerically by sampling and projecting $S(\mathbf{x}_i, \omega_i, \mathbf{x}_o, \omega_o)\cos(\theta_i)$ into the frequency domain, across its different dimensions. Depending on the underlying BSSRDF model, the cosine-weighted BSSRDF may depend on the viewing direction, the incident lighting direction, and the distance and angle between $\mathbf{x}_o$ and $\mathbf{x}_i$.

For instance, the dipole model has a separable form:

$$\mathcal{F}[\widehat{S}] = \mathcal{F}\left[R_d(||\mathbf{x}_i - \mathbf{x}_o||)\, F_i(\theta_i)\, \cos(\theta_i)\, F_o(\theta_o)\right],$$

where $R_d$ is the diffuse reflectance, and $F_i$ and $F_o$ are the incident and outgoing Fresnel terms [19, Equation 5]. Here, we take advantage of the separability of the model (w.r.t. $\theta_i$ and $\theta_o$) to express its Fourier transform as

$$\mathcal{F}[\widehat{S}] = \underbrace{\mathcal{F}\left[R_d(||\mathbf{x}_i - \mathbf{x}_o||)\right] \mathcal{F}\left[F_i(\theta_i)\, \cos\theta_i\right]}_{\mathcal{F}[\widehat{S}_i](\Omega_{\mathbf{x}_i}, \Omega_{\mathbf{x}_o}, \Omega_{\omega_i})} \mathcal{F}\left[F_o(\theta_o)\right].$$

Since we are only concerned with the DC $[\Omega_{\mathbf{x}_i}, \Omega_{\omega_i}] = [0, 0]$ hyperplane, the spatial bandwidth is computed with the 1D diffuse reflectance spectrum $\mathcal{F}\left[R_d\right](\Omega_{x_o})$. We discuss the outgoing term $\mathcal{F}\left[F_o\right](\Omega_{\omega_o})$ below.

In contrast, the directional dipole [14] additionally takes $\omega_i$ and the direction between $\mathbf{x}_i$ and $\mathbf{x}_o$ into account:

$$\mathcal{F}[\widehat{S}_i] = \mathcal{F}\left[\frac{e^{-\sigma_{tr}||\mathbf{x}_i - \mathbf{x}_o||}}{4\pi^2||\mathbf{x}_i - \mathbf{x}_o||^3} M(\mathbf{x}_i - \mathbf{x}_o, \omega_{12})\, F_i(\theta_i) \cos\theta_i\right],$$

where $M(\mathbf{x}_i - \mathbf{x}_o, \omega_{12})$ models the spatial-directional scattering distribution and $\omega_{12}$ is the refraction of $\omega_i$ at $\mathbf{x}_i$ [14, Equation 17]. We extract the outgoing spatial bandwidth by taking the maximum 1D bandwidth for various angles between $\mathbf{x}_i - \mathbf{x}_o$, the normal at $\mathbf{x}_i$ and the refracted ray $\omega_{12}$.

In all instances, we compute a conservative — that is to say, such that our derived frequency domain bounds are strictly larger than the true underlying bounds — estimate of the outgoing spatial and directional frequency *bandwidths*, $B_s$ and $B_\theta$, as the values required to retain 95% of the energy of the discrete power Fourier spectrum.

**Angular Variation.** The angular variation of the BSSRDF is modulated by the outgoing Fresnel term above, and we use a windowed Fourier transform to compute the bandwidth of $\mathcal{F}[F_o](\Omega_{\omega_o})$, again as the 95[th] energy percentile spectrum value. We tabulate these bandwidths as a function of $\theta_o$, and use them to modulate $B_\theta$; this is particularly important at grazing angles, where the effects of the spectrum of the outgoing Fresnel term can significantly impact the angular bandwidth of the outgoing radiance.

### 3.2 Outgoing Radiance Bandwidth Computation

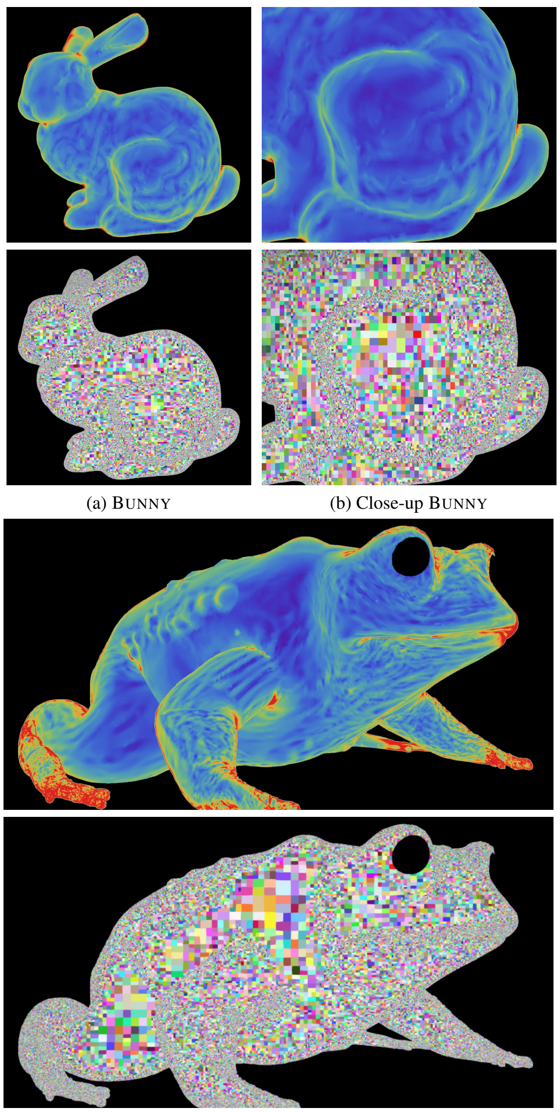

(a) BUNNY (b) Close-up BUNNY

(c) TOAD

Figure 5: First rows: The sampling rate $s_p$ computed from the screen-space bandwidth estimation. Second rows: Pixel areas from which the sampling rate predicts an adequate approximation of the outgoing radiance variation.

Given the spatial-angular bandwidth of the outgoing radiance at a shade point, estimated as the BSSRDF bandwidth, we need to compute the associated pixel frequency bandwidth. To do so, we are motivated by Bagher et al.'s [2] bandwidth tracking approach, applying bandwidth evolution operators defined by Durand et al. [13] to the bandwidth vector $[B_s, B_\theta]^T$. Figure 4 (c – e) illustrates the transport operators in the following order:

1. we transform from local shade point coordinates to global coordinates by projecting the outgoing spectrum onto the shade point's tangent plane, which amounts to a shear in the spatial frequency according to the local curvature $k$,
2. we take the foreshortening towards the viewpoint due to $\cos\theta_o$ into account, stretching the spectrum spatially, and
3. we evaluate the spectrum at the sensor, after transport through free-space with a distance $d$, by applying an angular shear to the spectrum.

These operations can be compactly expressed as matrix operators, if we act directly on frequency bandwidths instead of the full spectra [2], as:

$$\mathbf{T}_d = \begin{bmatrix} 1 & 0 \\ d & 1 \end{bmatrix}, \; \mathbf{P}_x = \begin{bmatrix} 1/\cos\theta_x & 0 \\ 0 & 1 \end{bmatrix}, \text{ and } \mathbf{C}_k = \begin{bmatrix} 1 & k \\ 0 & 1 \end{bmatrix}.$$

We apply these operators, in order, to the outgoing radiance bandwidth (i.e., the BSSRDF bandwidth $[B_s, B_\theta]$), to predict the final screen space bandwidth vector for a pixel as:

$$\begin{bmatrix} B_p & B_a \end{bmatrix}^{\mathrm{T}} = \mathbf{T}_d \, \mathbf{P}_x \, \mathbf{C}_k \, \begin{bmatrix} B_s & B_\theta \end{bmatrix}^{\mathrm{T}}. \tag{3}$$

Isolating the screen space angular bandwidth $B_a$ above,

$$B_a = B_\theta + d \, (B_s + kB_\theta) / \cos\theta \,, \tag{4}$$

and applying the Nyquist criterion, we arrive at the pixel sampling rate $s_p$ (in units of pixel$^{-1}$) as twice the angular screen space bandwidth,

$$s_p = 2 \, B_a \max\left(f_x/W, f_y/H\right), \tag{5}$$

for a $W \times H$ image resolution and a horizontal and vertical field of view of $f_x$ and $f_y$. Figure 5 visualizes the screen space sampling rate for the scenes we render.

## 4 HIERARCHICAL APPROACH

We now explain how to utilize our bandwidth estimation in order to accelerate rendering with BSSRDFs. First, we review the single hierarchy approach of Jensen and Buhler [18], then explain how we can use a dual hierarchy to adaptively cluster both illumination samples and pixels simultaneously.

### 4.1 Hierarchical Surface Integration

Jensen and Buhler [18] pointed out that we can cluster illumination samples over the surface in order to reduce the cost of BSSRDF evaluations. The underlying observation is that we can aggregate contributions from illumination samples that are distant from a given shading point. We can thus evaluate the BSSRDF only once for a cluster of such illumination samples, resulting in fewer BSSRDF evaluations.

This approach has two passes. In the first pass, pre-integrated illumination samples are inserted into a tree data structure where each inner node $i$ represents the aggregated information of its children. For example, each node stores the average illumination, the total surface area $A_i$, and the irradiance-weighted average location $\mathbf{p}_i$ of its children. In the second pass, we traverse this tree until the current node accurately represents all the contributions of its children to a given shading point. If the shading point is in the bounding volume of the current node, we keep traversing the tree and consider contributions from the children nodes. Otherwise, we traverse to the child nodes only if the estimate of the solid angle subtended by the illumination samples , $\Delta\omega = A_i \Big/ ||\mathbf{x}_o - \mathbf{p}_i||^2$, is larger than the user-defined quality threshold $\varepsilon$ (Algorithm 1). While this approach significantly reduces the cost of integration over the surface, it is repeated for each shading point without considering the smoothness of resulting pixels values in screen-space.

**Algorithm 1** Single-hierarchy tree traversal: $\mathbf{x}_o$ is the shading point/pixel, with $I_L$ and $I_R$ as children of the active node.

```
procedure SINGLE(x_o, I)
    if I is leaf or (Δω < ε and x_o ∉ BBOX(I)) then
        c ← contribution of I to x_o
        add c to x_o
    else
        SINGLE(x_o, I_L), SINGLE(x_o, I_R)
    end if
end procedure
```

### 4.2 Dual Hierarchy for Pixel-Surface Integration.

We leverage a dual hierarchy to avoid traversing the illumination tree at *every* pixel. Similar to the spatial hierarchy of illumination samples in the previous approach, we also cluster pixels in the screen space and traverse two trees simultaneously. Each node in our pixel-tree stores the average world-space position $\mathbf{p}_o$ corresponding to the pixel group, its bounding box, the average normal direction, the average view direction, and the list of pixels covered by the node. This dual-tree approach allows us to evaluate the contribution from a cluster of illumination samples to a cluster of pixels. Algorithm 2 is a pseudocode of our dual-tree approach.

The key difference from the single tree approach is that, at each traversal step, we have a choice of refining the pixel and/or illumination point clusters. For refining clusters of illumination samples, we use a criterion similar to the single tree approach. We always traverse down the tree if bounding volumes of pixels and illumination samples intersect. Otherwise, we decide if we want to keep traversing the tree based on the extended solid angle measure, $\Delta\omega = A_i \Big/ ||\mathbf{p}_o - \mathbf{p}_i||^2$, which uses the average position $\mathbf{p}_o$ of clustered pixels.

Criterion to Refine Pixel Clusters. To refine pixel clusters, we use our frequency analysis to predict the potential variation in pixels. Given a pixel sampling rate $s_p[i]$ for the $i^{\text{th}}$ pixel in a pixel tree node, an estimate of a screen-space filter extent, centered about the node, is

$$P = \rho \Big/ \max_i \left(s_p[i]\right), \tag{6}$$

where $\rho$ is a user-defined parameter that intuitively corresponds to the fraction of captured outgoing radiance required to avoid discontinuity artifacts. The $\rho$ setting influences pixel cluster refinement during traversal.

We refine the cluster only if our criterion predicts a high variation of outgoing radiance in the parent node's pixels (the SHADE routine in Algorithm 2). During shading (SHADE procedure) we do not adaptively refine the illumination cluster and conservatively assume that $\Delta\omega < \varepsilon$ is satisfied for all the children nodes. We could alternatively continue refining along the illumination tree for sub-nodes of the pixel tree. However, not refining results in higher performance without any noticeable visual artifacts.

## 5 Implementation

We implemented our approach in the G3D Innovation Engine [21] and our results were measured on a 3.9 GHz Intel Core i3-7100 with 12 GB of RAM. Both our illumination and pixels hierarchies are kd-trees, split along the largest bounding volume dimension. Our single- and dual-tree implementations use the same underlying kd-tree structure.

**Algorithm 2** Dual-hierarchy traversal: $S$ and $I$ are the root nodes of the shading point and illumination trees, with $S_{\{L|R\}}$ and $I_{\{L|R\}}$ their respective left and right children.

```
procedure DUAL(S, I)
    if Δω < ε and BBOX(S) ∩ BBOX(I) = ∅ then
        SHADE(S, I)
    else if S is leaf and I is leaf then
        SHADE(S, L)
    else if S is leaf then
        DUAL(S, I_L), DUAL(S, I_R)
    else if I is leaf then
        DUAL(S_L, I), DUAL(S_R, I)
    else
        DUAL(S_L, I_L), DUAL(S_R, I_L), DUAL(S_L, I_R), DUAL(S_R, I_R)
    end if
end procedure
procedure SHADE(S, I)
    if LENGTH(S) < ρ / BANDWDITH(S) then
        SHADE(S_L, I), SHADE(S_R, I)
    else
        c ← contribution of I to S
        add c to all pixels x in S
    end if
end procedure
```

We uniformly sample points on translucent objects with Bowers' et al. [7] blue noise approach, and image-space curvature values are interpolated from object-space values precomputed with the robust curvature estimator of Kalogerakis et al. [20]. In Sections 3.1 we compute BSSRDF bandwidths as the 95[th] percentile of the discrete spectrum, since we find this setting balances numerical stability and accuracy. We use $\rho = 0.75$ (Equation 6) in all our scenes and plots, as we found this value avoids discontinuity artifacts while providing good performance. We discuss the performance vs. accuracy trade-offs of $\rho$ and $\varepsilon$ in Section 6.

## 6 Results and Discussions

We have tested our approach on objects with a range of scattering parameters, as well as adapting our frequency analysis to support *several* BSSRDF models: the standard dipole [19], the "better dipole" [9], and the directional dipole [14]. We use three scenes of increasing radiometric complexity: BUNNY, TOAD, and PICNIK (Figures 8, 7, and 1). TOAD uses the directional dipole, and the remaining scenes use the better dipole.

We compare root mean square error (RMSE) of our technique to the single hierarchy of Jensen and Buhler [18], for *total* render time, on the BUNNY and TOAD scenes (Figure 6). We sampled $\varepsilon$ to generate the plots, and our approach consistently reaches equal quality in less time.

Comparisons in the BUNNY scene (Figure 8) illustrate our scalability with pixel coverage: the performance discrepancy between the full-view (Figure 6a) and zoom-in (Figure 6b) renderings is due to the total number of pixels present in the pixel hierarchy. As expected, the benefit of our approach increases with the number of translucent pixels: one can expect our approach to scale sub-linearly here, which is particularly favorable given recent trends towards higher resolution renderings and higher pixel supersampling rates.

We provide computational timing breakdowns when rendering an image with our technique in Table 1: specifically, we measured the Fourier precomputation, construction of the illumination tree as well as the shading tree, and final rendering times (all on the BUNNY scene and for different values of $\varepsilon$). All timings are reported on a

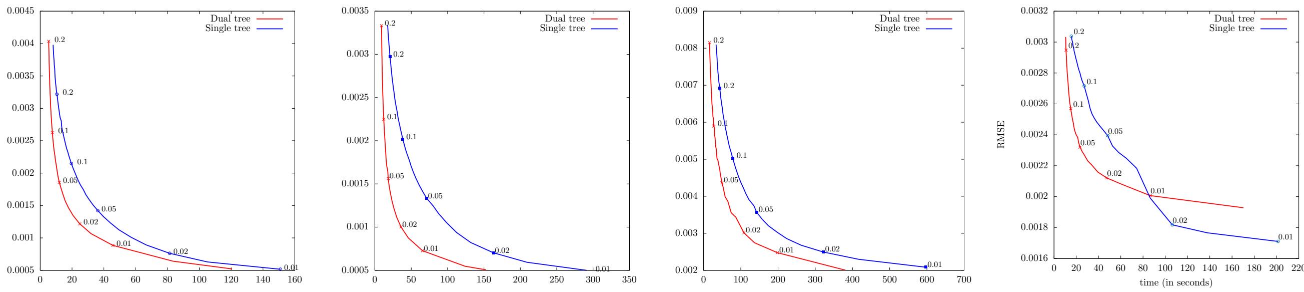


(a) BUNNY (b) BUNNY (CLOSE-UP) (c) TOAD (d) PICNIK

Figure 6: We compare our approach (red) to Jensen and Bulher [18] (blue) for different settings of $\varepsilon$. We highlight the $\varepsilon \in [0.01, 0.2]$ values and consistently reach equal-quality (measured in RMSE; $y$-axis) in less render time (in seconds; $x$-axis). The PICNIK scene challenges the assumptions of our work, and we only obtain equal-quality benefit at lower rendering times (albeit enough for visual convergence).

single core. Note that we use the same precomputed illumination tree across trials. Unsurprisingly, since the Fourier precomputation and the shading tree construction are independent of $\varepsilon$, we obtain similar timings across trials.

| $\varepsilon$ | Illumination Tree | Fourier Precomputation | Shading Tree | Final Shading |
|---|---|---|---|---|
| 0.2 | | 0.10s | 5.59s | 6.77s |
| 0.1 | 1.74s | 0.11s | 5.47s | 11.73s |
| 0.05 | | 0.09s | 5.12s | 21.21s |

Table 1: Computation times for various parts of the algorithm.

Dual-tree (Ours) Single-tree [18]

Figure 7: The TOAD scene has a bumpy geometry with detailed textures. We compare the difference images of the multiple scattering term against the ground truth for an equal rendering time (196s). The difference images are scaled by 50 for visualization. Our approach achieves more accurate estimation than the single-tree in the same rendering time.

Our screen space adaptive sampling rate accounts both distance, local curvature, foreshortening and BSSRDF properties from first principles. Moreover, it properly explains (and it subsumes) most of the previously used heuristics in the literature, e.g., depth and normal min/max methods [22]. Our sampling rate formulation (Equation 4) is simple and only requires the precomputation of two values $(B_\theta, B_s)$ *per material*. We do not require an additional pass to aggregate min/max statistics over the G-buffer.

We introduce a new error metric for aggregating pixel rendering cost and reducing shading cost in scenes with BSSRDFs. In doing so, we opted to follow the the solid angle metric methodology of Jensen and Buhler [18] in order to avoid the cost of evaluating upper-bound metrics that rely on BSSRDF evaluation, such as in LightCuts methods [8, 23]. We observed that such an upper-bound metric cannot scale to more complex BSSRDF shading models (i.e., the complexity of the material evaluation). We do, however, note that our frequency metric could be used as a well-found replacement of the upper-bound metrics for the specific case of BSSRDFs. Indeed, Belcour and Soler have shown [4] that frequency criterion can be used to provide an approximate relative error measure.

**Limitations.** The PICNIK scene (Figure 6d) is a "failure" case: specifically, our current implementation creates a separate dual tree per object in order to prevent illumination from bleeding between neighboring objects, and since the PICNIK scene includes several (smaller) translucent objects, we only obtain a benefit for a sub-region of the quality/performance range. Moreover, the solid angles $\Delta\omega$ spanned by pixel-tree nodes are more sensitive to errors for small objects and small BSSRDF scales. Since our technique approximates $\Delta\omega$ for a group of pixels, it is sensitive to these scenarios and we plan to address this issue in the future by devising more appropriate $\Delta\omega$ estimates. Overall, the fact that the additional tree construction time is amortized over fewer pixels, and the nature of our non-conservative $\Delta\omega$ estimate in the presence of smaller objects (in image-space), contribute to the suboptimal performance profile in this scene. This also explains the reduced error reduction rate for small $\varepsilon$.

In some difficult scenarios, high frequencies may be missed due to pixel discretization: for instance, a worst-case scenario would involve a camera facing an object with staggered depth discontinuities, which may miss small depth changes due to pixel aliasing. Here, we would group pixels that should not have been grouped.

## 7 CONCLUSION

We presented a new frequency analysis of BSSRDFs in order to predict the variation of outgoing radiance for multiple subsurface scattered light. We build and traverse a dual hierarchy over illumination samples and pixels using a well-founded refinement strategy that leverages our frequency bandwidth estimates. This yields an adaptive rendering strategy that almost consistently outperforms the state-of-the-art. Moreover, our frequency analysis and bandwidth estimates apply to a variety of existing BSSRDF models with negligible precomputation, our rendering technique scales positively with shading resolution, all without introducing any additional approximation error.

Our approach leads to several interesting open questions:

1. An interesting avenue would be to combine our work and the one of Arbree et al. [1]. They cluster both light source

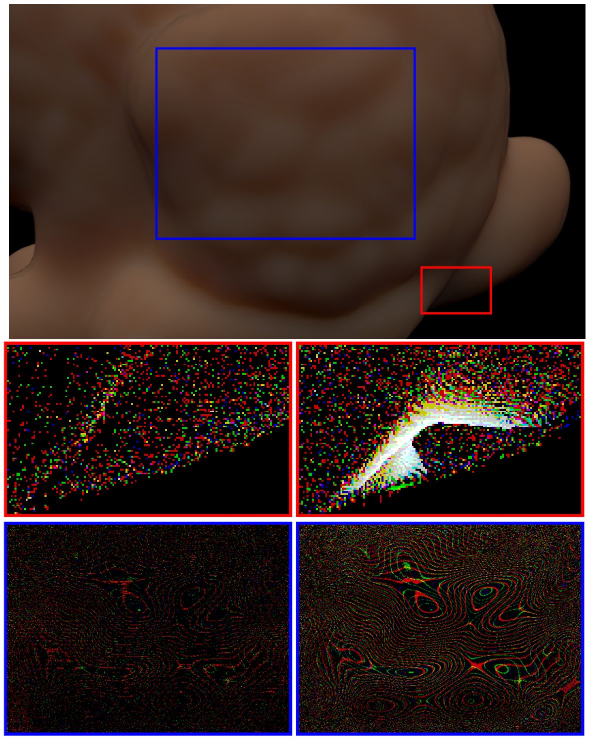

Dual-tree (Ours) Single-tree [18]

Figure 8: The BUNNY scene. We compare the difference images of the multiple scattering term against the ground truth for an equal rendering time (60s). The difference images are scaled by 200 for visualization. In this example, our approach removes artifacts under the tail and reduces Moiré patterns present in the single-tree approach.

positions and illumination points at the surface of object while we cluster both illumination points and shading points. Based on the same multiple cluster idea, it should be possible to build a trial-tree that accounts for those three components during rendering;

2. Our frequency analysis does not account for surface global illumination transport: we ignore visibility, and the use of spatial illumination samples ignores incident radiance variation. Modeling this behavior more accurately could lead to less conservative bandwidth estimates and traversal criteria;
3. There are no reasons why our theory and implementation could not support other existing diffusion models (e.g., *quantized diffusion* [11]), and so implementing these models under our framework is interesting even if only for the sake of completeness;
4. Investigating how increases in $\varepsilon$ should affect our choice of $\rho$, and vice versa, leads to the interesting question of whether an "optimal" parameter setting for both these values could be computed automatically;
5. The effects on performance and accuracy of replacing our position-based solid angle approximation with the actual projected solid angle of the underlying surface elements would also be worth investigating.
6. Lastly, there is much potential in analyzing our algorithm's temporal properties, notably in terms of information reuse across neighboring frames as well as in ensuring that no temporal artifacts occur due to image-space filtering.

## Acknowledgments

We also graciously acknowledge funding from the Fonds de recherche sur la nature et les technologies, through their Établissement de nouveaux chercheurs universitaires program (2014-NC-173734), as well as a Master's scholarship.

## References

[1] A. Arbree, B. Walter, and K. Bala. Single-pass scalable subsurface rendering with lightcuts. In *Computer Graphics Forum*, vol. 27, pp. 507–516. Wiley Online Library, 2008.

[2] M. M. Bagher, C. Soler, K. Subr, L. Belcour, and N. Holzschuch. Interactive rendering of acquired materials on dynamic geometry using bandwidth prediction. In *ACM I3D*, pp. 127–134, 2012.

[3] L. Belcour, K. Bala, and C. Soler. A local frequency analysis of light scattering and absorption. *ACM Trans. on Graph.*, 33(5):163:1–163:17, Sept. 2014.

[4] L. Belcour and C. Soler. Frequency based kernel estimation for progressive photon mapping. In *SIGGRAPH Asia 2011 Posters*, p. 47. ACM, 2011.

[5] L. Belcour, C. Soler, K. Subr, N. Holzschuch, and F. Durand. 5d covariance tracing for efficient defocus and motion blur. *ACM Trans. Graph.*, 32(3):31:1–31:18, July 2013.

[6] L. Belcour, L.-Q. Yan, R. Ramamoorthi, and D. Nowrouzezahrai. Antialiasing complex global illumination effects in path-space. *ACM Trans. Graph.*, 36(1), Jan. 2017.

[7] J. Bowers, R. Wang, L.-Y. Wei, and D. Maletz. Parallel poisson disk sampling with spectrum analysis on surfaces. *ACM Trans. on Graph.*, 29(6):166:1–166:10, Dec. 2010.

[8] N. Bus, N. H. Mustafa, and V. Biri. IlluminationCut. *Computer Graphics Forum (Proceedings of Eurographics 2015)*, 34(2):561 – 573, 2015.

[9] E. d'Eon. A better dipole. `http://www.eugenedeon.com/project/a-better-dipole/`, Nov 2012. [accessed 7-April-2015].

[10] E. d'Eon. A dual-beam 3d searchlight bssrdf. In *ACM SIGGRAPH 2014 Talks*, pp. 65:1–65:1, 2014.

[11] E. D'Eon and G. Irving. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph.*, 30(4):56:1–56:14, July 2011.

[12] R. A. Dubouchet, L. Belcour, and D. Nowrouzezahrai. Frequency Based Radiance Cache for Rendering Animations. In M. Zwicker and P. Sander, eds., *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*. The Eurographics Association, 2017.

[13] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. X. Sillion. A frequency analysis of light transport. *ACM Trans. Graph.*, 24(3):1115–1126, July 2005.

[14] J. R. Frisvad, T. Hachisuka, and T. K. Kjeldsen. Directional dipole model for subsurface scattering. *ACM Trans. Graph.*, 34(1):5:1–5:12, Dec. 2014.

[15] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comput. Phys.*, (2):325–348, Dec. 1987.

[16] R. Habel, P. H. Christensen, and W. Jarosz. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. In *Eurographics Symposium on Rendering*, pp. 27–37, 2013.

[17] A. Jarabo, R. Buisan, and D. Gutierrez. Bidirectional clustering for scalable vpl-based global illumination. In *CEIG*, 2015.

[18] H. W. Jensen and J. Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graph.*, 21(3):576–581, July 2002.

[19] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *ACM SIGGRAPH*, pp. 511–518, 2001.

[20] E. Kalogerakis, P. Simari, D. Nowrouzezahrai, and K. Singh. Robust statistical estimation of curvature on discretized surfaces. In *Eurographics Symposium on Geometry Processing*, pp. 13–22.

[21] M. McGuire, M. Mara, and Others. G3D Innovation Engine, 2014. http://g3d.sourceforge.net/.

[22] G. Nichols and C. Wyman. Multiresolution splatting for indirect illumination. In *ACM I3D*, pp. 83–90, 2009.

[23] B. Walter, A. Arbree, K. Bala, and D. P. Greenberg. Multidimensional lightcuts. *ACM Transactions on Graphics (TOG)*, 25(3):1081–1088, 2006.

[24] B. Walter, S. Fernandez, A. Arbree, K. Bala, M. Donikian, and D. P. Greenberg. Lightcuts: A scalable approach to illumination. *ACM Trans. Graph.*, 24(3):1098–1107, July 2005.