

Voxel-Based Interactive Rendering of Translucent Materials under Area Lights using Sparse Samples

Ming Di Koa

School of Computer Science and Engineering
Nanyang Technological University
Singapore
mdkoal@e.ntu.edu.sg

Henry Johan

Fraunhofer Singapore
Singapore
hjohan@fraunhofer.sg

Alexei Sourin

School of Computer Science and Engineering
Nanyang Technological University
Singapore
assourin@ntu.edu.sg

Abstract—Interactive rendering of translucent materials in virtual worlds has always proved to be challenging. In our work, we develop a voxel illumination framework for translucent materials illuminated by area lights. Our voxel illumination framework consists of two voxel structures. They are the Enhanced Sub-surface Light Propagation Volumes (ESLPV), which handles the local translucent material appearance and the Light Propagation Volumes (LPV), which handles indirect illumination for the entire scene. Using a set of sparsely distributed Poisson disk samples in the ESLPV and LPV, illumination can be gathered from area lights. A uniform set of Poisson disk samples on the translucent objects is resampled and chosen as Translucent Planar Lights (TPLs) and is used to distribute lighting from translucent objects into the LPV by an additional gathering process. Our technique allows for direct and indirect illuminations from highly scattering translucent materials to be rendered interactively under area lighting at good quality. We can achieve similar effects, such as scattered light illumination from translucent materials, when compared to offline renderers without precomputations.

Keywords—Translucent Materials; Area Lights; Direct Illumination; Indirect Illumination; Interactive Rendering; Virtual Worlds

I. INTRODUCTION

Interactive rendering of scenes with highly scattering translucent objects is difficult to compute efficiently. Although approximation techniques can model the local illumination of translucent materials, techniques that model their indirect illumination interactively on surrounding surfaces are rarely researched. Translucent objects can act as area lights when they release scattered light to their surroundings to produce visual effects in the appearance as soft diffuse color bleeding, e.g., Figs. 1a, 1b. Although it is easier for rendering algorithms to work with translucent materials using point light sources, area lights often produce soft shadows which exhibit coloration by scattered lighting within the translucent materials. This effect contributes greatly to the visual realism in interactive rendering. This effect is not present on diffuse surfaces as they are usually optically thick (e.g., Fig. 1c). Additionally no indirect illumination will reach the external objects as it is physically impossible.

Monte Carlo techniques [1], [2] can render interreflections from translucent materials but they are too computationally expensive. Improvements in other methods, such as photon mappings, allow translucent materials to produce indirect illu-

mination effects [3], [4], but they are impractical for generating images at interactive rates. Current radiosity [5] approaches allow for real-time computations, but at the expense of long precomputation time and large storage data for precomputing form-factors.

In this paper, we present a voxel based illumination approach that renders translucent objects under direct area light illumination and renders indirect illumination from translucent and diffuse surfaces interactively without precomputations. We present three main contributions in this work.

- A Poisson disk sampling solution to allow lighting information from area lights to be injected into a voxel structure for rendering translucent objects
- A Poisson disk sampling solution to allow lighting information from area lights to be injected into a voxel structure for rendering indirect illumination for diffuse surfaces.
- An interreflection framework for distributing indirect illumination from translucent objects to their nearby diffuse surfaces. This allows translucent objects to be treated as area lights.

Our proposed illumination pipeline allows for interactive rendering without any precomputations and achieves the soft diffuse color bleeding effect of scattered light from translucent materials similar to offline renderers.

II. RELATED WORK

In recent years, several real-time techniques have been developed for translucent materials and its scattering phenomenon known as subsurface scattering. Lensch et al. [7] formulated the subsurface scattering effect as calculation of throughput between vertices using radiosity techniques. Mertens et al. [8] and Carr et al. [9] suggested using a hierarchy of clustered triangles and multi-resolution mesh respectively similar to hierarchical radiosity. Yu et al. [5] extended the radiosity method for subsurface scattering to support multiple bounce global illumination. However, radiosity approaches requires a large number of form factors which may in turn require a large amount of computation as well as storage. Wang et al. [10] and Adam et al. [11] devised tetrahedralization techniques for discretizing a mesh into a 4-connected mesh structure. Their techniques are able to handle




Fig. 1. Rendering of a water glass with a red scattering material illuminated with a single area light source inside the cup at 1024x768. We show our rendering in (a) with both the direct and indirect illumination components, and a reference image generated by Mitsuba Path Tracer (64 samples per pixel) in (b). The translucent material in the reference image is rendered with Photon Beam Diffusion method [6]. Image (a) is rendered in 146 ms. Image (b) is rendered in 54.96 seconds. Image (c) shows the same experiment setup but with a diffuse water glass. No light is scattered to the exterior of the water glass.

much more complex geometries, especially those with thin features. However, the time needed to process the mesh to a suitable structure, though automated, may still take minutes. Geist et al. [12] proposed using a voxel-based approach to simulate a light propagation process. Their process was much more storage efficient but computationally expensive. Bernabei et al. [13] proposed reducing the principal directions of propagation to six directions which enables rendering to be done at interactive rates. Borlum et al. [14] reformulated the LPV technique, and changed the LPV into a Subsurface Light Propagation Volumes (SSLPV). Unlike the LPV, the SSLPV is a voxelized structure of an object and not a scene. The flux is injected into the material and distributed by a propagation process. An extension using a hierarchical propagation approach was proposed by Koa et al. [15]. Their approach, the ESLPV, allows light to scatter across further distances simulating highly scattering materials. However, these works deal with local illumination models and do not extend to indirect illumination.

Rendering dynamic indirect lighting requires many ray-triangle intersection tests. Many algorithms work on determining whether a pixel should be illuminated based on adaptive algorithms that simplify either the scene or the ray tracing process. In recent years, emphasis has been given to algorithms that work on voxel-based illumination to compute dynamic illumination on the fly. Crassin et al. [16] proposed an indirect illumination algorithm that uses a sparse voxel octree (SVO) to store illumination and geometric information in its voxels. The SVO structure allowed an efficient cone tracing solution to reduce the costs of ray tracing, as well as to generate filtered mipmaps which allows for blending of illumination information across different resolutions. However, it requires a large memory space for pre-allocating voxel information, which are usually larger than the acceleration structure used for ray tracing. Sugihara et al. [17] proposed using a layered shadow map to store illumination in the Reflective Shadow Maps (RSM) [18] while voxel structures continue to hold geometric information. This approach significantly reduces

memory consumption. Kaplanyan et al. [19] proposed the Light Propagation Volumes (LPV) as an extremely memory efficient solution for single and multi-bounce indirect diffuse illumination. The LPV uses 4-Spherical Harmonics coefficients for each color channel to represent the flux distribution in each voxel. The initial flux distribution is defined using a RSM. The position, normal and flux density of each texel obtained in the RSM are injected into the corresponding voxel in the LPV. The LPV then undergoes a propagation process, where each step of the propagation process allows the flux distribution in each voxel to be distributed to its neighbouring voxel. After the flux distribution in the LPV has reached equilibrium, the final voxel results can be used as a representation of indirect diffuse illumination. Unfortunately, using the RSM restricts the LPV to only point and spot lights. Hedman et al. [20] proposed using a dynamic distribution of point lights that contributes to the visible pixels for each frame. In each frame, visible point lights can be re-used and new ones can be created. A heuristic sample distribution is defined to obtain temporal stability. In our work, we use the LPV approach with a fixed sparse set of Poisson-disk samples which achieves temporal stability and speed at the expense of accuracy.

III. OUR PIPELINE

We combined the ESLPV, in Koa et al. [15] and LPV [19] for our illumination pipeline. The ESLPV renders local subsurface scattering effects while the LPV distributes the indirect illumination from diffuse surfaces as well as translucent objects. The 3D objects in our scene and the translucent objects are voxelized for the LPV. The translucent objects can be voxelized according to the solid voxelization algorithm described in Schwarz et al. [21]. Voxelization allows us to store geometric information and material properties into these voxels, which are used for propagation. In our work, we use the LPV structure from Doghramachi's work [22] because of the way geometric blockers are represented in the voxels. Their method of storing normals of blockers into a tetrahedron face allows us to extract the closest normals relevant to the surface

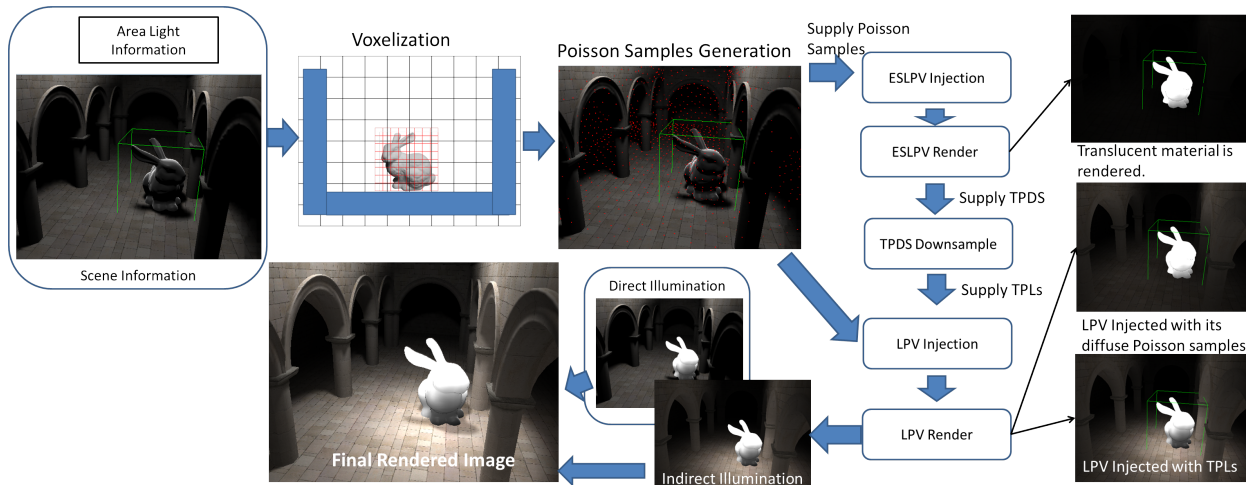


Fig. 2. The pipeline of our work with the indirect illumination combined with the direct illumination output.

we are rendering. This allows more accurate illumination computations.

As previous work for the ESLPV [15] and LPV [19] did not deal with area light illumination, we design a complete framework for rendering translucent materials under area lights as shown in Fig. 2. Our system starts off with area light and scene information (e.g., mesh information). We generate Poisson disk samples [23] for both our translucent and diffuse objects. The Poisson disk sampling algorithm produces an ideal distribution with a minimum Euclidean distance between samples. This makes it ideal for generating sparse samples for rendering. We use the notation diffuse Poisson disk samples (DPDS) and translucent Poisson disk samples (TPDS) to describe the two sets of Poisson disk samples generated on diffuse and translucent surfaces respectively. Every TPDS will have its light intensity computed and transferred into the ESLPV for injection. The ESLPV renders the translucent object with the given set of TPDS. This set of TPDS are downsampled to a reduced set of translucent planar lights (TPLs). This set of TPLs is used to represent the translucent object as an area light source. The intensities of the TPLs can be obtained by sampling the ESLPV voxels with a rendering equation after the propagation stage has been completed.

Next, we compute the reflected flux at the location of each DPDS. The reflected flux is computed with both the area light information and the TPLs, indicating that each DPDS now contains the reflected flux of direct illumination from the area light and the TPLs. The DPDS are injected by accumulating their intensities with previous TPLs into the LPV and rendered. By propagating the light intensity in the LPV designed by Doghramachi [22], we are able to simulate the indirect illumination of translucent materials as well as diffuse materials in the scene. The final rendered image is combined with our direct illumination (for diffuse surfaces), which handles area light direct illumination. We implemented the previously proposed method of direct illumination (Koa et al. [24]) using multi-resolution rendering under area lighting.

This method is able to produce high quality soft shadows and is suitable for overlaying our indirect illumination component. The right side of Fig. 2 shows the various output of each injection stage.

A. Direct illumination for translucent objects from area lights

We first distribute a set of TPDS on the surface of a translucent object. For each Poisson disk sample, we perform a ‘gathering’ operation in which the transmitted flux from each refracted light ray from the area lights is computed. We create a ESLPV voxel structure of 32^3 resolution for the translucent objects. The light intensities are stored according to the refracted light directions. The transmitted flux entering the translucent medium from each refracted light ray is converted to its SH representation of a clamped cosine lobe and is accumulated into the voxels corresponding to their locations. We describe the transmitted flux at point x_i from a light ray emitted at a virtual point light (VPL) (from uniformly divided patch from the area light) location x_{light} from direction $\vec{\omega}$ as $L_t(x_i, \vec{\omega})$ in (1a). We use x_{light} as the representative point for each uniformly divided patch on the area light from uniform sampling. $T(\vec{\omega}, \vec{\omega}')$ refers to the Fresnel term for describing the fraction of light energy transmitted into the translucent material after entering from direction ω_k and refracted to direction $\omega_{k'}$. A represents the area of a uniform patch on the area light defined by the VPL from uniform point sampling. A is used as part of the form factor computation between area to point contribution. \vec{N} refers to the normal of the TPDS, and \vec{L} refers to the normalized ray direction from the TPDS to the light. $I_{intensity}$ refers to the light intensity from the VPL. Equation 1a can be converted to its Spherical Harmonics (SH) representation as in (1b) where the injected flux for each Poisson disk sample is represented in the SH coefficients, $c_{lm}^{(\vec{\omega}')}$, of a clamped cosine-lobe oriented at the refraction vector, $\vec{\omega}'$. y_{lm} refers to the SH basis functions. Fig. 3 describes how the refracted light energy is injected into the ESLPV voxels. The SH coefficients in each ESLPV

voxel are accumulated by summation and then inversely scaled by the total number of TPDS located in them. Once the light intensities are injected into the ESLPV voxels, we can perform the light propagation as usual, which creates the local information required for rendering the translucent appearance.

$$L_t(x_i, \vec{\omega}) = \frac{T(\vec{\omega}, \vec{\omega}')(\vec{N} \cdot \vec{L}) * (\vec{N}_{light} \cdot -\vec{L}) * A * I_{intensity}}{|x_i - x_{light}|^2} \quad (1a)$$

$$L_{t,lm}(x_i, \vec{\omega}) = \frac{[\sum_{l=0}^{\infty} \sum_{m=-l}^l c_{lm}^{(\vec{\omega}')} y_{lm}]}{\pi} * L_t(x_i, \vec{\omega}) \quad (1b)$$

B. Indirect illumination from area lights

1) Indirect illumination from direct area light sources:

As the usual LPV scheme does not handle area lights, our proposed solution generates a set of Poisson disk distributed samples [23] on the 3D scene we are rendering. The illumination for each Poisson disk sample is computed by performing a ‘gathering’ operation to the light source. We only gather contribution from rays that pass the visibility test. We compute the reflected flux, L_{VPL} , for each diffuse Poisson disk sample, x_i , which is a summation of the contribution of each individual light ray to the sample point (refer to (2b)):

$$F_s = (\vec{N} \cdot \vec{L}_k) * (\vec{N}_{light} \cdot -\vec{L}_k), \quad (2a)$$

$$L_{VPL}(x_i, \vec{\omega}) = \frac{1}{\pi} \sum_1^K \frac{F_s * A * I_{intensity} * \rho}{|x_i - x_{k_light}|^2} \quad (2b)$$

where \vec{N} refers to the normal of Poisson disk sample x_i . \vec{L}_k refers to the vector from x_i to a VPL, x_{k_light} on the light. \vec{N}_{light} refers to the normal direction of the light. F_s refers to the foreshortening factor. $I_{intensity}$ refers to the intensity at x_{k_light} with area A from uniform sampling. A is used for area to point form factor computation. ρ refers to the diffuse coefficients of the Poisson disk sample.

The reflected flux in each Poisson disk sample is then deposited into the respective voxel at one unit normal distance away from the Poisson disk sample in their Spherical Harmonics (SH) representation. Following Kaplanyan et al. [19], we can first use a clamped cosine lobe oriented in the z-axis, represented by zonal harmonics, rotated to the direction \vec{N} . The flux distribution when converted to SH Coefficients is represented in (3). As the accumulated reflected flux is defined as the integral of each individual light ray over a hemisphere, a normalization factor of $\frac{1}{\pi}$ is required to conserve the energy.

$$I(x_i, \vec{\omega}) = \frac{L_{VPL}(x_i, \vec{\omega})}{\pi} * \sum_{l=0}^{\infty} \sum_{m=-l}^l c_{lm}^{(\vec{N})} y_{lm} \quad (3)$$

In the LPV implementation, we are only required to store the four SH coefficients (two bands) multiplied by $\frac{L_{VPL}(x_i, \vec{\omega})}{\pi}$ into each voxel. Each color channel (R,G,B) is stored separately as four float values of SH coefficients each. Once we have obtained the SH coefficients by a cosine lobe orientated in

direction \vec{N} , we can scale these coefficients by the irradiance of the Poisson disk sample. For Poisson disk samples that lie in the same voxels, the intensities can be accumulated in the same voxel.

The remaining process follows the standard LPV [19] pipeline. Propagation is performed to distribute the light intensity throughout the LPV. The final result is an approximation to the indirect illumination distributed by diffuse surfaces. We refer to Fig. 3 for the light injection process from area lights.

2) *Indirect illumination from translucent materials:* Scattered light from the translucent material can be represented by a sparse set of planar lights, with a point location allocated to its center. We downsample the set of TPDS that were generated earlier in Section III-A to easily get a reduced set of samples known as TPLs. The main intention for downsampling is to ensure that we are able to reduce the total number of samples for representing the translucent object. This reduces computation time significantly. For our downsampling algorithm, we basically reduced the 32^3 ESLPV to a 8^3 representation and temporarily deposit each TPDS into their corresponding locations in the 8^3 voxel volume. We then iterate each voxel of the 8^3 voxel volume and find the TPDS that is nearest to the center of each of the voxels in the 8^3 volume. Since the original Poisson disk sampling approach by Bowers et al. [23] distributes samples according to some voxel resolution, our downsampling approach would still allow us to reasonably obtain samples that are evenly distributed. The nearest TPDS in each voxel is chosen as the TPL. This chosen sample maintains its location on the surface of the object, unlike interpolation methods which might provide a non-surface sample. TPLs are represented as virtual planar light sources with an area allocated to them. This representation corresponds to how voxels are represented as planar area lights in the SSLPV [14]. The output radiance of each TPL can be computed by (4). A_{voxel} is the area of a voxel face of the ESLPV. There is no outgoing Fresnel term as it is incorporated when DPDS ‘gather’ illumination in the irradiance term from the TPL (5b).

$$I_{TPL}(x_{TPL}, \vec{\omega}_o = \vec{n}) = \frac{2}{A_{voxel}} \left(\sqrt{\frac{1}{4\pi}} c_{00}^{(\vec{n})} + \frac{1}{2} \sqrt{\frac{3}{4\pi}} c_{10}^{(\vec{n})} \right) \quad (4)$$

Once a set of TPLs is computed, each of the DPDS will perform a ‘gathering’ operation on the TPLs (refer to Fig. 4). DPDS are represented in red circles while TPLs are represented in green circles. The reflected flux contribution from each TPLs is then added to the original illumination contribution $L_{VPL}(x_i, \vec{\omega})$ from area lights in Section III-B. This contribution is then stored into the respective voxel location corresponding to each of the DPDS’ location. The reflected flux contributed by K number of TPLs is computed in (5b), where K is the number of TPLs visible from x_i . $I_{TPL}(x_{TPL}, \vec{\omega}_o = \vec{n})$ is the intensity of the TPL computed by (4). $T(\vec{\omega}, \vec{\omega}')$ is the Fresnel transmission term describing the percentage of light energy leaving the translucent medium




Fig. 3. Injecting transmitted and reflected light flux into ESLPV and LPV, respectively, with area lights. For diffuse materials, light is injected into the LPV voxel above the surface. For translucent materials, light is injected into the ESLPV voxel on the surface.

from direction $\vec{\omega}'$ to $\vec{\omega}$. F_s is the foreshortening factor as described in (2a). ρ is the albedo of the material at x_i . A_{TPL} refers to the surface area represented by each TPL. This is computed by dividing the total surface area of the translucent object with the total number of TPLs on the object.

$$L_{TPL_Single}(x_i, \vec{\omega}) = \frac{F_s A_{TPL} T(\vec{\omega}, \vec{\omega}') I_{TPL}(x_{TPL}, \vec{\omega}' = \vec{n})}{|x_i - x_{TPL}|^2} \quad (5a)$$

$$L_{TPL}(x_i, \vec{\omega}) = \frac{\rho}{\pi} \sum_1^K L_{K_TPL_Single} \quad (5b)$$




Fig. 4. Injecting light intensities into LPV with TPLs. DPDS are represented in red circles while TPLs are represented in green circles.

IV. IMPLEMENTATION

A. Poisson disk samples generation

We implemented Bowers et al. [23] Poisson disk sampling method as it properly distributes samples across a regular grid

structure. Although a GPU implementation of their algorithm has been created [25], we found that it is unnecessary to re-generate samples regularly. A new set of Poisson disk samples would only be required when a new diffuse surface or translucent object is added to the scene, or if any object in the scene is deformed.

B. Gathering operations on Poisson disk samples

As we need to generate multiple visibility rays from our Poisson disk samples to VPLs/TPLs, we utilize NVIDIA's CUDA and OptiX Prime [26] for our ray creation and ray tracing process. Firstly, a ray is created from every (diffuse and translucent) Poisson disk sample to a VPL on the area light source. The visibility tests are performed using OptiX Prime.

For every visible ray from a VPL to the TPDS, the transmitted flux is computed and compressed into its SH basis based on its refraction direction. $L_t(x_i, \vec{\omega})$ is projected into the SH coefficients of a clamped cosine lobe oriented at refracted direction $\vec{\omega}'$ in (1b). $c_{lm}^{(\vec{\omega}')}$ refers to the SH coefficients of the clamped cosine lobe based on the refracted ray direction $\vec{\omega}'$. Using CUDA's SHFL functions, the illumination can be gathered quickly if the samples are in multiples of warp sizes (16,32). Warps are units of threads which CUDA is able execute them concurrently. The gathered transmitted flux for each TPDS at position, x_i , is computed by accumulating the contribution of every visible ray in (6). We only need to inject the values of four SH coefficients for each color channel, as described by $L_{t,lm}^v(x_i, \vec{\omega})$ in (7), into the relevant voxel location of the ESLPV.

$$L(x_i, \vec{\omega}) = \sum_1^K L_{t,lm}^k(x_i, \vec{\omega}), \quad (6)$$

$$L_{t,lm}^v(x_i, \vec{\omega}) = \text{float4}\left(\frac{c_{0,0} * L_t(x_i, \vec{\omega})}{\pi}, \frac{c_{1,-1} v * L_t(x_i, \vec{\omega})}{\pi}, \frac{c_{1,0} * L_t(x_i, \vec{\omega})}{\pi}, \frac{c_{1,1} * L_t(x_i, \vec{\omega})}{\pi}\right), \quad (7)$$

Next, the TPDS is downsampled to TPLs, which are usually less than 256. All DPDS would then gather illumination (5b) from TPLs after the TPLs have been computed with the results from the ESLPV using (4). Each TPL now represents an area light source with radiance and normals.

C. Indirect illumination voxels

Reflected flux from all diffuse Poisson disk samples, after ‘gathering’ from VPLs and TPLs, are injected into the LPV. The SH coefficients of the reflected flux are computed by multiplying the reflected flux of the Poisson disk sample and the clamped cosine SH lobe of the normal vector of the Poisson disk sample. The SH coefficients of the reflected flux from the Poisson disk samples are injected into the LPV at an offset of one unit voxel away in the normal direction. We keep a counter using a 3D texture and GLSL’s `imageAtomicAdd` function to keep track of the number of samples inside each voxel. Ultimately, the intensities stored in the voxels should be normalized by the number of samples in them. This normalization is not required in the original ESLPV and LPV as the normalization has already been done by dividing the total light energy by the number of texels in the RSM. However, in our case we are unable to use a RSM to represent an area light.

V. RESULTS AND DISCUSSION

We show the rendering results of four scenes: a translucent water glass, a translucent Buddha model and a plane, a translucent bunny in a colored Cornell box, a translucent dragon in the Sponza scene in Figs. 1a, 5a, 6a, 7, respectively, in 1024x768 pixel resolution. We tabulate the computation timings in Table I. The rendering was performed on an Intel i5 3.40GHz CPU with an NVIDIA GeForce GTX 980 GPU. The direct illumination for diffuse surfaces in our results are rendered with a screenspace multi-resolution technique [24], with 64 samples per fragment. The direct illumination provides the soft shadows effects. The translucent materials in our reference images (generated by Mitsuba [27]) are rendered with Photon Beam Diffusion [6]. Due to the sparse distribution of the Poisson disk samples, light ‘gathering’ operations are considered to be cheap operations that can be computed in tens of milliseconds. Furthermore, these ‘gathering’ operations only need to be performed when there is a change in object positions, material properties or light positions.

In Fig. 1a, we placed an area light inside a water glass so that it would be easier to only see the indirect illumination produced by the translucent materials as no direct illumination would be able to penetrate the water glass. Our result produced a reddish color bleed as observed on the floor outside the water glass. The reddish color is also observed in the reference image

in Fig. 1b. In Fig. 5, we illuminate a Buddha model floating on top of a planar surface. Fig. 5a shows a soft shadow produced by the Buddha model. The shadow is not completely dark and has a yellowish faint partially illuminated by light from the translucent material. These effects are more evident in the soft shadow region of the shadows. A reference image is provided in Fig. 5b. In Fig. 6a, we can see that there is faint coloration of reddish and greenish in the shadows of the bunny. This coloration is created by the indirect illumination from both the colored walls of the Cornell box and the scattered light from the bunny. We show that the rendered reference image with the Mitsuba renderer’s path tracer [27] in Fig. 6b and the coloration in their shadows are similar to ours. In Fig. 7a, we specifically render the indirect illumination only. It can be seen that scattered green light is present around the diffuse surfaces near the dragon. The scattered green light is also present in the soft shadows areas formed by the pillars or the dragon in Figs. 7b, 7c.

With reference to Table I, ‘Direct Illum.’ refers to the time used for rendering direct illumination (excluding translucent material rendering) using a multi-resolution screenspace approach [24]. ‘Poisson Gathering’ refers to the time used in ‘gathering’ illumination from Poisson disk samples generated on diffuse and translucent surfaces after being lit by VPLs. ‘TPL Gathering’ refers to the time needed for all diffuse Poisson disk samples to gather light from TPLs. ‘ESLPV’ represents the time needed to render the illumination on the translucent object using our ESLPV method. ‘LPV’ represents the time required for the LPV to render indirect illumination on the entire scene. We can observe that direct illumination of translucent materials using area lights can be quickly computed and rendered in less than 10 ms. Indirect illumination for the entire scene can be rendered in less than 20 ms for most of our results. ‘Total time’ refers to the time required for rendering a single frame with both direct and indirect illuminations. The total time for rendering appears high due to the direct illumination component. However, our indirect illumination technique is independent of the direct illumination technique used. The columns ‘DPDS’, ‘TPDS’ and ‘TPL’ describe the number of samples required for diffuse Poisson disk samples, translucent Poisson disk samples and Translucent Planar Lights, respectively. One advantage of voxel-based techniques is that the timings for ESLPV and LPV do not scale up with the scene complexity, which allows our indirect illumination for translucent materials to be rendered at fast speed.

We do note that there are some differences compared to the reference images in some areas of the images. In Fig. 6, we note that our result (Fig. 6a) has a different ceiling color compared to the reference image. This is because the LPV is not a physically-based algorithm and it uses a very abstract representation of geometry and reflectance property to distribute indirect illumination. Hence, we cannot expect the same quality of results as that in radiosity, photon mapping or path-tracing methods when rendering multi-bounce illumination. However, we can still expect color-bleeding effects to

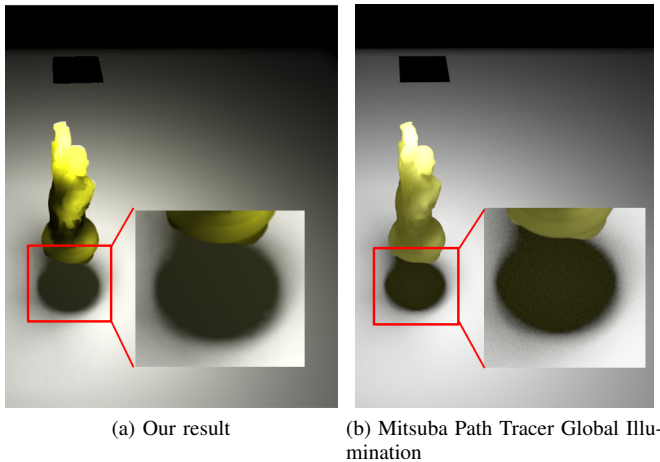


Fig. 5. Rendering of a Buddha object with a material property simulating an apple in (a) and (b). We show our rendering in (a) with both the direct and indirect illumination components, and the reference generated by Mitsuba Path Tracer (64 samples per pixel) in (b) with the Photon Beam Diffusion method [6]. Image (b) is rendered in 28.97 seconds.

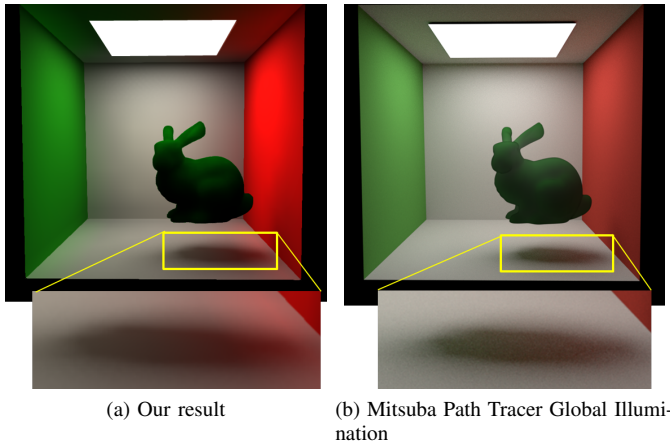


Fig. 6. Rendering of a Stanford bunny with material properties simulating jade in a colored Cornell box. We show our rendering in (a) with both the direct and indirect illumination components, and the reference generated by Mitsuba Path Tracer (64 samples per pixel) and the Photon Beam Diffusion method [6] in (b). Image (b) is rendered in 54.21 seconds.

be rendered. We are unable to simulate high frequency effects for translucent materials with weak scattering properties due to the highly compressed nature of our spherical harmonics representation. However, this is compensated through its low-memory consumption and fast rendering speed.

VI. CONCLUSION AND FUTURE WORK

We have presented an efficient and low cost solution for illumination of translucent materials from area lights. This is first done by generating a set of Poisson disk samples over the translucent materials and performing an illumination ‘gathering’ operation over them. The Poisson disk samples are then injected into the ESLPV [15], which distributes the light intensities through the translucent materials. Similarly, Poisson

disk samples are distributed around diffuse surfaces in the LPV for gathering direct illumination from the area lights.

In order to further illuminate the scene using the translucent objects, each of the Poisson disk samples in the LPV will gather illumination from a downsampled set of translucent planar lights (TPLs) on the translucent objects. These samples are then injected into the LPV. The LPV distributes the light, simulating indirect illumination with contributions from the diffuse and translucent objects. Our method efficiently distributes indirect illumination with very little computation overheads.

Using a Poisson disk sampling approach, we ensure that translucent materials and their indirect illuminations can be rendered with good quality under area lighting. This method achieves interactive rendering while remaining precomputationless and maintaining low storage costs.

The diffuse Poisson disk samples (DPDS) in the scene increase with scene complexity, hence leading to an overall increase in time for ‘gathering’ illumination from TPLs. A multi-resolution approach should be designed to select an appropriate set of samples. Other optimizations such as Debevec’s light probe sampling [28] could be more efficient in choosing our TPLs from our TPDS. However, it was designed for sampling a 2D environment map but it can well be extended to simulate importance sampling for clustering of TPLs.

ACKNOWLEDGMENT

This research is partially supported by the National Research Foundation, Prime Minister’s Office, Singapore under its International Research Centres in Singapore Funding Initiative. The 3D models used in our examples are obtained from: <http://graphics.cs.williams.edu/data>.

REFERENCES

- [1] M. Pharr and P. Hanrahan, “Monte Carlo evaluation of non-linear scattering equations for subsurface reflection,” in *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’00. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 75–84.
- [2] H. W. Jensen, J. Legakis, and J. Dorsey, “Rendering of wet materials,” in *Proceedings of the 10th Eurographics Conference on Rendering*, ser. EGWR’99. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 1999, pp. 273–282.
- [3] C. Donner and H. W. Jensen, “Rendering translucent materials using photon diffusion,” in *ACM SIGGRAPH 2008 Classes*, ser. SIGGRAPH ’08. New York, NY, USA: ACM, 2008, pp. 4:1–4:9.
- [4] W. Jarosz, H. W. Jensen, and C. Donner, “Advanced global illumination using photon mapping,” in *ACM SIGGRAPH 2008 Classes*, ser. SIGGRAPH ’08. New York, NY, USA: ACM, 2008, pp. 2:1–2:112.
- [5] Y. Sheng, Y. Shi, L. Wang, and S. G. Narasimhan, “A practical analytic model for the radiosity of translucent scenes,” in *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, ser. I3D ’13. ACM, 2013, pp. 63–70.
- [6] R. Habel, P. H. Christensen, and W. Jarosz, “Photon beam diffusion: A hybrid monte carlo method for subsurface scattering,” *Computer Graphics Forum (Proceedings of EGSR 2013)*, vol. 32, no. 4, Jun. 2013.
- [7] H. P. A. Lensch, M. Goesele, P. Bekaert, J. Kautz, M. A. Magnor, J. Lang, and H.-P. Seidel, “Interactive rendering of translucent objects,” in *Proceedings of the 10th Pacific Conference on Computer Graphics and Applications*, ser. PG ’02. IEEE Computer Society, 2002, pp. 214–

TABLE I

RENDERING STATISTICS FOR 1024x768 IMAGES. DIRECT ILLUM REFERS TO THE TIME REQUIRED TO RENDER DIRECT ILLUMINATION [24] ONLY FOR DIFFUSE SURFACES. THE REMAINING COLUMNS (POISSON GATHERING, TPL GATHERING, ESLPV, LPV) ARE TIMINGS FOR OUR WORK.

Figure	Triangles	Direct Illum. [24](ms)	Poisson Gathering (ms)	TPL Gathering (ms)	ESLPV (ms)	LPV (ms)	Total time (ms)/(FPS)	DPDS	TPDS	TPL
1a - Water glass	9,280	35	3	9	5	9	61 /16.4	3,711	17,077	179
5a - Buddha	110,346	77	3	5	2	10	97/10.1	3,680	4,913	236
6a - Bunny	80,020	90	4	9	5	9	117/8.55	2,603	8,438	177
7b - Dragon	238,076	80	6	24	5	9	124/8.06	21,366	8,495	84



(a) Indirect Illumination only

(b) Dragon View 1

(c) Dragon View 2

Fig. 7. Using TPLs for injection produces a soft diffuse look from the green translucent Stanford dragon model. The scene is illuminated by a rectangular area light source from the top. Image (a) shows the indirect illumination. The walls and pillars of the lower level which are closer to the dragon appear much greener compared to the top level walls and pillars. The light source used in the scene is a small white rectangular area light. Images (b) and (c) show other rendered views with the direct and indirect illumination components. The soft shadow regions exhibit indirect scattered illumination from the translucent material.

- [8] T. Mertens, J. Kautz, P. Bekaert, F. Van Reeth, and H.-P. Seidel, "Efficient rendering of local subsurface scattering," in *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, ser. PG '03. IEEE Computer Society, 2003, pp. 51–58.
- [9] N. A. Carr, J. D. Hall, and J. C. Hart, "GPU algorithms for radiosity and subsurface scattering," in *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, ser. HWW'S '03. Eurographics Association, 2003, pp. 51–59.
- [10] Y. Wang, J. Wang, N. Holzschuch, K. Subr, J.-H. Yong, and B. Guo, "Real-time Rendering of Heterogeneous Translucent Objects with Arbitrary Shapes," *Computer Graphics Forum*, vol. 29, no. 2, pp. 497–506, May 2010.
- [11] A. Arbree, B. Walter, and K. Bala, "Heterogeneous subsurface scattering using the finite element method," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 7, pp. 956–969, Jul. 2011.
- [12] R. Geist, K. Rasche, J. Westall, and R. Schalkoff, "Lattice-Boltzmann lighting," in *Proceedings of the Fifteenth Eurographics conference on Rendering Techniques*, ser. EGSR '04. Eurographics Association, 2004, pp. 355–362.
- [13] D. Bernabei, A. Hakke-Patil, F. Banterle, M. Di Benedetto, F. Ganovelli, S. Pattanaik, and R. Scopigno, "A parallel architecture for interactively rendering scattering and refraction effects," *Computer Graphics and Applications, IEEE*, vol. 32, no. 2, pp. 34–43, 2012.
- [14] J. Børllum, B. B. Christensen, T. K. Kjeldsen, P. T. Mikkelsen, K. O. Noe, J. Rimestad, and J. Mosegaard, "SSLPV: Subsurface light propagation volumes," in *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics*, ser. HPG '11. ACM, 2011, pp. 7–14.
- [15] M. D. Koa and H. Johan, "ESLPV: Enhanced subsurface light propagation volumes," *The Visual Computer*, vol. 30, no. 6-8, pp. 821–831, Jun. 2014.
- [16] C. Crassin, F. Neyret, M. Sainz, S. Green, and E. Eisemann, "Interactive indirect illumination using voxel cone tracing," *Computer Graphics Forum (Proceedings of Pacific Graphics 2011)*, vol. 30, no. 7, Sep 2011.
- [17] M. Sugihara, R. Rauwendaal, and M. Salvi, "Layered Reflective Shadow Maps for Voxel-based Indirect Illumination," in *Eurographics/ ACM SIGGRAPH Symposium on High Performance Graphics*, I. Wald and J. Ragan-Kelley, Eds. The Eurographics Association, 2014.
- [18] C. Dachsbacher and M. Stamminger, "Reflective shadow maps," in *Proceedings of the 2005 Symposium on Interactive 3D graphics and games*, ser. I3D '05. ACM, 2005, pp. 203–231.
- [19] A. Kaplanyan and C. Dachsbacher, "Cascaded light propagation volumes for real-time indirect illumination," in *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, ser. I3D '10. ACM, 2010, pp. 99–107.
- [20] P. Hedman, T. Karras, and J. Lehtinen, "Sequential Monte Carlo instant radiosity," in *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, ser. I3D '16. New York, NY, USA: ACM, 2016, pp. 121–128.
- [21] M. Schwarz and H.-P. Seidel, "Fast parallel surface and solid voxelization on gpus," *ACM Trans. Graph.*, vol. 29, no. 6, pp. 179:1–179:10, Dec. 2010.
- [22] H. Doghramachi, *GPU Pro 4: Advanced Rendering Techniques - Rasterized Voxel-Based Dynamic Global Illumination*. Natick, MA, USA: A. K. Peters/CRC Press, 2013.
- [23] J. Bowers, R. Wang, L.-Y. Wei, and D. Maletz, "Parallel Poisson disk sampling with spectrum analysis on surfaces," *ACM Trans. Graph.*, vol. 29, no. 6, pp. 166:1–166:10, Dec. 2010.
- [24] M. D. Koa, H. Johan, and A. Sourin, "Interactive screenspace fragment rendering for direct illumination from area lights using gradient aware subdivision and radial basis function interpolation," *Computers And Graphics*, vol. 64, pp. 37 – 50, 2017.
- [25] X. Ying, S.-Q. Xin, Q. Sun, and Y. He, "An intrinsic algorithm for parallel Poisson disk sampling on arbitrary surfaces," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 9, pp. 1425–1437, Sep. 2013.
- [26] S. G. Parker, J. Bigler, A. Dietrich, H. Friedrich, J. Hoberock, D. Luebke, D. McAllister, M. McGuire, K. Morley, A. Robison, and M. Stich, "Optix: a general purpose ray tracing engine," *ACM Trans. Graph.*, vol. 29, no. 4, pp. 66:1–66:13, Jul. 2010.
- [27] W. Jakob, "Mitsuba renderer," 2010, <http://www.mitsuba-renderer.org>.
- [28] P. Debevec, "A median cut algorithm for light probe sampling," in *ACM SIGGRAPH 2006 Courses*, ser. SIGGRAPH '06. New York, NY, USA: ACM, 2006.