


Real-time Global Illumination by Precomputed Local Reconstruction from Sparse Radiance Probes

ARI SILVENNOINEN, Aalto University, Remedy Entertainment

JAAKKO LEHTINEN, Aalto University, NVIDIA



a) Our result (fully dynamic lights), 3.90 ms

b) Path traced lightmaps

Fig. 1. a) *Real-time global illumination rendered using our method in 3.9ms with fully dynamic lights, cameras and diffuse surface materials.* b) *Ground truth (converged path tracing).*

We present a direct-to-indirect transport technique that enables accurate real-time rendering of indirect illumination in mostly static scenes of complexity on par with modern games while supporting fully dynamic lights, cameras and diffuse surface materials. Our key contribution is an algorithm for reconstructing the incident radiance field from a sparse set of local samples – radiance probes – by incorporating mutual visibility into the reconstruction filter. To compute global illumination, we factorize the direct-to-indirect transport operator into global and local parts, sample the global transport with sparse radiance probes at real-time, and use the sampled radiance field as input to our precomputed local reconstruction operator to obtain indirect radiance. In contrast to previous methods aiming to encode the global direct-to-indirect transport operator, our precomputed data is local in the sense that it needs no long-range interactions between probes and receivers, and every receiver depends only on a small, constant number of nearby radiance probes, aiding compression, storage, and iterative workflows. While not as accurate, we demonstrate that our method can also be used for rendering indirect illumination on glossy surfaces, and approximating global illumination in scenes with large-scale dynamic geometry.

CCS Concepts: • **Computing methodologies** → **Rendering**;

Additional Key Words and Phrases: Global illumination, real-time rendering

ACM Reference format:

Ari Silvennoinen and Jaakko Lehtinen. 2017. Real-time Global Illumination by Precomputed Local Reconstruction from Sparse Radiance Probes. *ACM Trans. Graph.* 36, 6, Article 230 (November 2017), 13 pages. <https://doi.org/10.1145/3130800.3130852>

1 INTRODUCTION

Indirect (global) illumination is widely acknowledged to be a key factor in perceived image realism. Over the past decade, the film industry has largely switched from more or less phenomenological models to physically-based rendering techniques [Christensen and Jarosz 2016] as they offer realistic, predictable, and controllable results. Yet, real-time rendering of accurate indirect lighting remains a challenge, as the necessary integration over the space of all light paths remains too costly to be performed for every pixel. The issues are only aggravated by high-DPI screens, stereo rendering, and the high frame rate requirements of VR headsets. Consequently, typical applications today do not support detailed indirect illumination with local, dynamic light sources, and instead mostly resort to static pre-computation or a limited form of dynamic illumination through precomputed radiance transfer [Sloan et al. 2002].

We present a technique that enables accurate real-time rendering of indirect diffuse illumination in mostly static scenes of complexity on par with modern games with fully dynamic lights, cameras and diffuse and emissive surface materials. Our key contribution is an algorithm for faithfully interpolating incident radiance captured at a sparse set of low-frequency “radiance probes” to nearby receiver points. While much previous work has built on the same basic ideas, our method features little inappropriate light leaking thanks due to our novel formulation of the interpolation at the *surfaces seen by both the probes and the receiver points*, not only spatially between the samples. Using our interpolator, we formulate a novel direct-to-indirect precomputed transport technique that yields results that compare favorably to path traced references with probe sets that are extremely sparse in comparison to previous methods at similar quality. We study the properties of our reconstruction using tools for analyzing light field sampling and reconstruction; this allows reasoning about the role of spatial sampling densities and angular bandwidths.

© 2017 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3130800.3130852>.

2 RELATED WORK

Real-time global illumination methods draw efficiency from two main sources: 1) geometric approximations that enable faster integration, e.g., using point-based models or voxels instead of accurate geometry [Christensen 2008; Crassin et al. 2011; Ritschel et al. 2009], and 2) interpolating illumination from sparse sample sets, or sampling it densely but poorly and attempting to remove noise after the fact using sophisticated non-linear filters (e.g., [Dammertz et al. 2010; Kontkanen et al. 2004]). Our method makes use of exact geometry and sparse samplings.

Interpolation methods. Our work most naturally relates to sparse interpolation techniques. E.g. Ward [1992], Křivánek et al. [2005] interpolated irradiance and bandlimited radiance, respectively, from point sets whose local density is driven by distance to close-by geometry. Schwarzhaupt et al. [2012] introduced a second order, occlusion-aware error metric for irradiance caching. Lehtinen et al. [2008] perform similar interpolation, but address changing spatial frequency needs with a hierarchy of grids. We share these methods’ fundamental idea — computing illumination sparsely and interpolating — but address their greatest shortcoming: light leaks due to interpolation weights that do not respect visibility. Our algorithm is able to faithfully interpolate radiance from much sparser cache point sets.

Local-global separation. Arikan et al [2005] interpolate radiance represented as spherical harmonic expansions from sparse samples, and model the shadowing effect of local occluders by approximate visibility. In a sense, we share many of the same goals. Our technique accounts for visibility in a more precise manner, and consequently yields superior results as the number of probes decreases.

Precomputed light transport. Precomputed radiance transfer methods [Sloan et al. 2002], including “direct-to-indirect” techniques [Hašan et al. 2006], precompute linear operators that map emitted or direct illumination functions into indirect radiance or irradiance sampled over the scene surfaces. We use our interpolation method to formulate a novel direct-to-indirect technique, where direct illumination is captured at runtime at a sparse set of probes, and mapped to the incident radiance of densely sampled nearby receiver points through precomputed local interpolation operators. On a high level, our algorithm is precisely equivalent to earlier direct-to-indirect methods [Hašan et al. 2006; Kontkanen et al. 2006; Lehtinen et al. 2008; Martin 2010], but one with a transport operator explicitly factorized into global and local components. Both operators are *sparse*: there is a relatively small number of probes only, and each receiver point only needs to consult the radiance from a small subset of nearby probes. This is in contrast to earlier direct-to-indirect techniques that globally link senders and receivers together. The locality, which aids compression and streaming, furthermore makes it simple for us to support indirect illumination effects from limited forms of dynamic blockers. We make use of clustered principal component analysis for compression [Sloan et al. 2003].

Light Field Probes. Light field probes are spatial samples of the angular variation in the light field [Buehler et al. 2001]. They may be represented in various forms, e.g., cube maps [Hooker 2016; McGuire et al. 2017] or spherical harmonics expansions [Ramamoorthi and

Hanrahan 2001], which is also our choice. Hooker [2016] uses a sparse set of cube maps for accelerating final gathering in a pre-computation program. Like us, they make use of precise visibility between probes and query points; in contrast, they consult probes near the secondary hits of the final gather rays whereas we query the probes near the rays’ *origins*. In our direct-to-indirect setting, their approach would be prohibitively expensive, as it would link receiver points to potentially all probes in a scene. McGuire et al. [2017] performed ray tracing in cube maps with depth information; they also perform spatial interpolation, but with weighting that leads to inaccurate reconstruction with sparse samplings. In contrast to these methods, we sample the light field in an angle-bandlimited format using spherical harmonics. We reason about why and when such a representation is still sufficient for faithful reconstruction of irradiance below.

3 LOCAL TRANSPORT OPERATOR

Section 3.1 introduces our premise, interpolation of radiance from sparse samples of the directional radiance field. Section 3.2 presents a novel interpolation method that accounts for precise visibility and thus significantly reduces the signature “light leaking” of interpolation techniques. We also analyze the effect of angular bandwidth stored at the radiance samples. Section 3.3 builds a local precomputed radiance transfer method using the novel interpolation operator. Finally, Section 3.4 studies the properties of the interpolator by light field analysis techniques. Section 4 then describes a full working direct-to-indirect transport algorithm.

Assumptions. We assume throughout that the radiance field leaving the scene surfaces is diffuse; however, our framework supports glossy BRDFs for the final bounce towards the viewer. This is a common approximation in real-time global illumination techniques. We further assume that the scene is mostly static, but do offer support for approximate dynamic occluders. In addition, dynamic objects can receive indirect illumination by combining our method with (ir)radiance volumes [Greger et al. 1998].




Fig. 2. Issues and solutions for interpolation from a sparse set of radiance probes. a) Using the same constant direction ω when interpolating radiance is incorrect. b) Reprojecting the point $\Gamma(\mathbf{x}, \omega)$ seen by the receiver \mathbf{x} in the direction ω to the probes improves interpolation significantly. Still, interpolating radiance in directions where Γ cannot be seen by some probes leads to incorrect results. c) Our reprojected interpolation with added visibility weighting enables much more faithful reconstruction.

3.1 Overview

Our goal is to compute indirect illumination at a dense set of *receiver points*, which are either surface points or points in a radiance volume. This requires a dense sampling of the incident radiance field $L(\mathbf{x}, \omega)$ over directions ω for all receiver points \mathbf{x} . As it continues to be too expensive to perform a real-time dense angular radiance sampling, e.g., by rendering cube maps or tracing hundreds of rays per receiver, we aim to instead *reconstruct* the angular incident radiance function at the receivers by interpolation from a much sparser set of radiance observations $L_i(\omega) := L(\mathbf{p}_i, \omega)$ (*radiance probes*), where each probe captures the angular variation as seen from the probe position \mathbf{p}_i (Figure 2).

Our goal is shared by much previous work, e.g., [Křivánek et al. 2005; Ward et al. 1988]. These reconstructions typically take spatial averages of the probes' radiance (or irradiance) by

$$L(\mathbf{x}, \omega) \approx \frac{\sum_i w_i(\mathbf{x})L(\mathbf{p}_i, \omega)}{\sum_k w_k(\mathbf{x})}, \quad (1)$$

where $w_i(\mathbf{x})$ are spatial interpolation weights that are functions of the positions of the probe and the receiver, and potentially differences in their surface normals (Figure 2a). Some methods also make use of the spatial gradient $\nabla_{\mathbf{x}}L$ [Ward and Heckbert 1992].

We observe two key issues in the reconstruction performed according to Equation (1). First, interpolating the radiance using the same fixed direction ω for all probes is incorrect, because these directions do not generally point towards the surface point $\Gamma(\mathbf{x}, \omega)$ seen by the receiver \mathbf{x} in the direction ω (Figure 2a). Reprojecting the directions appropriately, i.e., consulting the probes' radiance in the direction of the point $\Gamma(\mathbf{x}, \omega)$ seen by the receiver in the direction ω , remedies the issue; this solution ("sheared reconstruction") has been employed in several light field reconstruction techniques [Buehler et al. 2001; Chai et al. 2000; Egan et al. 2011, 2009; Gortler et al. 1996; Lehtinen et al. 2011]. Second, local visibility easily causes erroneous interpolation of radiance in directions that are not mutually visible by the probes and the receiver (Figure 2b); here, the blue radiance sample from \mathbf{p}_2 incorrectly contributes to the weighted estimate. (This is the cause of the well-known "light leaks" in irradiance volume-type methods.) Some techniques attempt to mitigate this issue by consulting the visibility between the receiver \mathbf{x} and the probe \mathbf{p}_2 [McGuire et al. 2017]; however, the unoccluded line of sight *does not* guarantee that the radiance sample would be valid. Note, however, that there also exist directions — marked with green arrows in Figure 2b — where \mathbf{p}_2 *could* appropriately contribute to the radiance estimate at \mathbf{x} ; hence leaving \mathbf{p}_2 unused is wasteful. While sheared reconstruction techniques alleviate some of these issues [Egan et al. 2011, 2009], they still fundamentally require denser spatial sampling in regions of complex visibility.

3.2 Visibility-Aware Interpolation

Our key idea is to perform an intermediate reparameterization from the angular domain onto the scene surfaces — where visibility can be accounted for precisely — and then back into the angular domain. That is, when querying the radiance from a receiver \mathbf{x} in a given direction ω , we *only use information from radiance probes that directly see the point $\Gamma(\mathbf{x}, \omega)$ seen by the receiver in that direction*. Furthermore,

finding $\Gamma(\mathbf{x}, \omega)$ allows us to consult the probes' radiance in the appropriately reprojected direction (Figure 2c).

Technically, we combine the spatial interpolation weights $w_i(\mathbf{x})$ with directional, per-probe binary visibility weights $V_i(\omega)$ that pick out, for each direction separately, only probes that see the same surface point the receiver \mathbf{x} sees. Concretely, our radiance reconstruction is defined by

$$L(\mathbf{x}, \omega) \approx \frac{\sum_i w_i(\mathbf{x})V_i(\omega)L(\mathbf{p}_i, \psi_i(\omega))}{\sum_k w_k(\mathbf{x})V_k(\omega)}, \quad (2)$$

where the sum is over all probes, $w_i(\mathbf{x}) = w(\mathbf{x}, \mathbf{p}_i; r_i)$ is a spatial weight kernel with a finite support radius r_i , $V_i(\omega) = V(\Gamma(\mathbf{x}, \omega), \mathbf{p}_i)$ is the visibility between the hit point given by the ray-cast operator Γ and the probe at \mathbf{p}_i , and $\psi_i(\omega)$ is the direction from \mathbf{p}_i to the hit point $\Gamma(\mathbf{x}, \omega)$ (see Figure 2c). The probes' contributions are glued together by a partition of unity forced by the division by the sum of the combined weights, while the spatial weights guarantee smooth reconstruction across space. If no probe sees $\Gamma(\mathbf{x}, \omega)$, we define its radiance to be zero.

3.3 Bandlimited Probes and PRT

Notably, given infinite angular resolution at the probes, and remembering the diffuse radiance field assumption, the reconstruction of Equation (2) is *exact* as long as the point $\Gamma(\mathbf{x}, \omega)$ is seen by at least one probe. No set of spatial interpolation weights only can guarantee this. However, a high-resolution angular sampling of the probes' radiance is still overly expensive in terms of both computation storage, as we desire to perform this operation at real-time rates at a probe set that covers a complex scene. This motivates us to replace the probes' radiance fields by finite-dimensional basis expansions given in terms of spherical harmonics. We study the effect of this approximation in Section 3.4.

Concretely, we approximate the probes' radiance fields by

$$L(\mathbf{p}_i, \omega) = \sum_j \lambda_{ij} Y_j(\omega), \quad (3)$$

where λ , the *probe radiance vector*, contains the basis expansion coefficients for *all* probes flattened into a long vector, and Y_j are the spherical harmonic basis functions. For convenience, we index the probe radiance λ using probe index i and basis function index j . Each color band has its own vector $\lambda_r, \lambda_g, \lambda_b$, but we omit these for simplicity of notation except when otherwise noted.

For a fixed receiver \mathbf{x} , we define a *local transport operator* $\mathcal{P}_{\mathbf{x}}\{\lambda\}(\omega)$ that transforms the probe radiance vector λ to interpolated incident radiance at \mathbf{x} as a function of the continuous direction ω by evaluating Equation (2) using the probes' basis expansions:

$$\begin{aligned} \mathcal{P}_{\mathbf{x}}\{\lambda\}(\omega) &= \frac{\sum_i w_i(\mathbf{x})V_i(\omega) \sum_j \lambda_{ij} Y_j(\psi_i(\omega))}{\sum_k w_k(\mathbf{x})V_k(\omega)} \\ &= \sum_i \sum_j \lambda_{ij} \frac{w_i(\mathbf{x})V_i(\omega)Y_j(\psi_i(\omega))}{\sum_k w_k(\mathbf{x})V_k(\omega)} \\ &= \sum_i \sum_j \lambda_{ij} K_{ij}(\mathbf{x}, \omega), \\ \text{with } K_{ij}(\mathbf{x}, \omega) &= \begin{cases} \frac{w_i(\mathbf{x})V_i(\omega)Y_j(\psi_i(\omega))}{\sum_k w_k(\mathbf{x})V_k(\omega)} & \text{if } \sum_k w_k(\mathbf{x})V_k(\omega) > 0 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (4)$$

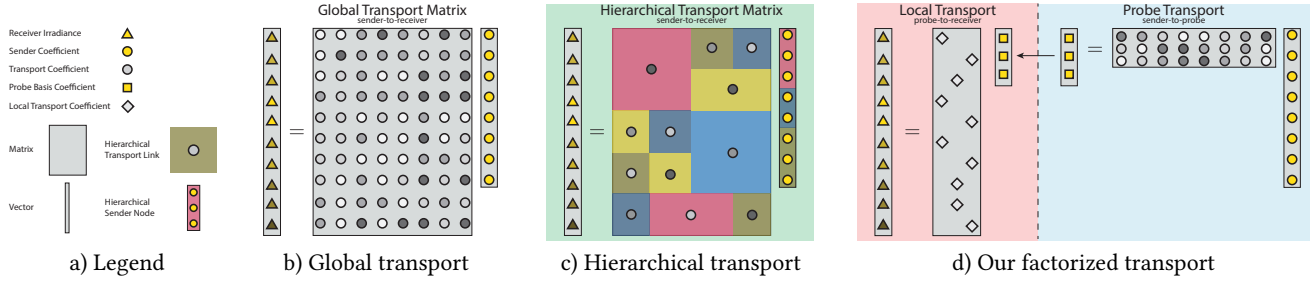


Fig. 3. *a) Legend.* Typical direct-to-indirect transport methods aim to compress the dense global transport operator (b) using a global, fixed hierarchy over the senders and receivers (c). We instead factorize the global transport operator into two parts: global probe transport and sparse “last-leg” local interpolation (d).

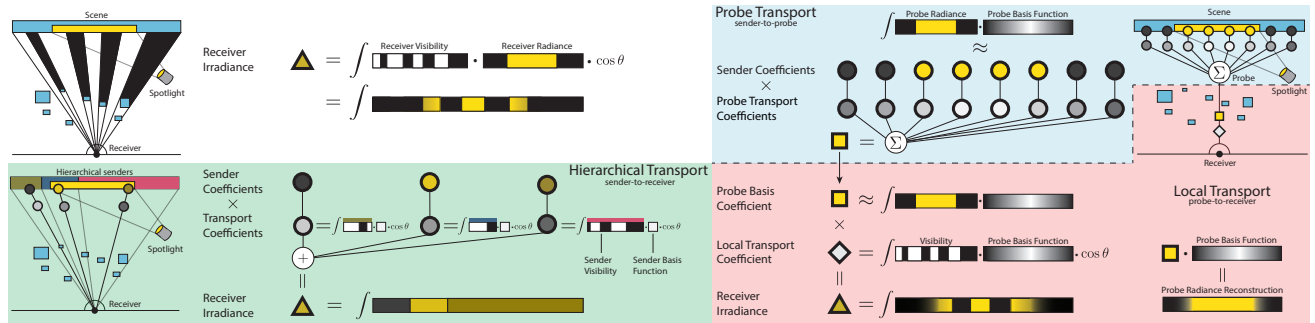


Fig. 4. *Scene with a spotlight and a receiver (top left).* We are seeking to reconstruct the irradiance at the receiver (top left). The ceiling geometry is subdivided into three piecewise constant hierarchical sender patches (bottom left). Hierarchical transport with piecewise constant basis functions amounts to multiplying the direct light coefficients in the sender basis with the precomputed transport coefficients, but using only three coefficients results in a poor approximation to the receiver irradiance (bottom left). Our factorized approach allows efficient decoupling of high-frequency near-field visibility from low-frequency illumination (left). The direct lighting is projected to the probe basis (top right). Our local transport results in more accurate reconstruction of the receiver irradiance with fewer coefficients due to decoupled near-field visibility (bottom right).

The *transport kernel* $K_{ij}(\mathbf{x}, \omega)$ is the angular function that results from projecting the value of the j th spherical harmonic from the i th probe onto the scene surface seen in the direction of $\Gamma(\mathbf{x}, \omega)$, modulated by visibility, and divided by the weights of other probes. The intuitive meaning of the linear combination is that the probes project their approximate radiance functions onto surfaces visible to them, and the receiver blends together results from all probes that see the point $\Gamma(\mathbf{x}, \omega)$.

Precomputed Transport. The local transport operator \mathcal{P}_x can be used to measure the interpolated incident radiance at \mathbf{x} in various ways. Owing to linearity, and, in line with the voluminous literature on precomputed radiance transfer, the result of the measurement with a general spherical function Φ is a linear function of the probe radiance vector λ :

$$\langle \mathcal{P}_x\{\lambda\}, \Phi \rangle = \int \mathcal{P}_x\{\lambda\}(\omega) \Phi(\omega) d\omega \quad (5)$$

$$= \sum_i \sum_j \lambda_{ij} \phi_{ij}, \quad (6)$$

where the *transport coefficients* ϕ_{ij} are given by $\langle K_{ij}(\mathbf{x}, \cdot), \Phi \rangle$.

Irradiance Transport. In particular, the irradiance $I(\mathbf{x})$ at a receiver \mathbf{x} is computed by measuring the interpolated approximate radiance field by the cosine lobe:

$$I(\mathbf{x}) \approx \int \mathcal{P}_x\{\lambda\}(\omega) \cos \theta d\omega = \sum_{ij} \lambda_{ij} \alpha_{ij}, \quad (7)$$

where the double indices i and j have been flattened using a single, linear index ij for notational convenience, and the entries of the irradiance transport vector α_x are given by $\alpha_{ij} = \langle K_{ij}(\mathbf{x}, \cdot), \cos \theta \rangle$. This corresponds to the SH transport vectors of, e.g., Sloan et al. [2002]. Note that most entries in α_x are zero, as only few probes contribute to each receiver.

Radiance Transport. We approximate the interpolated incident radiance field at \mathbf{x} using a directional basis expansion with orthogonal basis functions B_k :

$$L(\mathbf{x}, \omega) \approx \mathcal{P}_x\{\lambda\}(\omega) \approx \sum_k b_k B_k(\omega). \quad (8)$$

Finding the expansion coefficients b_k reduces to measuring $\mathcal{P}_x\{\lambda\}(\omega)$ by the basis functions $B_k(\omega)$, which yields the radiance transport matrix \mathbf{B}_x with entries $\beta_{(ij)k} = \langle K_{ij}(\mathbf{x}, \cdot), B_k \rangle$, again in line with the transport matrices of Sloan et al. and others. The coefficients of

transported radiance are given by $\mathbf{B}_x \lambda$. \mathbf{B}_x is a $n_k \times n_i n_j$ radiance transport matrix, where n_k is the number of receiver basis functions, n_i is the number of probes, and n_j is the number of probe basis functions. The rows contain the transport vectors for the corresponding basis functions:

$$\mathbf{B}_x = \begin{bmatrix} \beta_{(1,1),1} & \beta_{(1,2),1} & \cdots & \beta_{(n_i,n_j),1} \\ \vdots & \ddots & & \vdots \\ \beta_{(1,1),n_k} & \beta_{(1,2),n_k} & \cdots & \beta_{(n_i,n_j),n_k} \end{bmatrix}. \quad (9)$$

Note, again, that due to the local support of the weights $w_i(\mathbf{x})$, this matrix is sparse: only few probes contribute to each receiver.

A Note on Color. Importantly, as the color of surfaces seen by the probes is encoded in the probe radiance vectors, the transport coefficients ϕ_{ij} are decoupled from surface color, and encode only geometric information. In particular, this means that we are free to change surface properties, such as diffuse albedo and emissivity, without changing the transport coefficients.

Discussion. Previous direct-to-indirect transport methods aim to efficiently compute and store the global transport operator that maps basis expansions of (ir)radiance defined at the scene surfaces into basis expansions that describe indirect illumination, again on the scene surfaces, linking senders and receivers globally across the scene [Hašan et al. 2006; Jendersie et al. 2016; Lehtinen et al. 2008; Loos et al. 2011]. In this framework, our approach is to instead factorize the global transport operator into two parts: a precomputed local transport matrix that accounts for near field effects between the receiver and the mutually visible surfaces seen by the probes, and a global probe transport operator that captures radiance (also far-field) in a sparse set of probes at runtime (Figure 3). In effect, our construction amounts to a *local*, custom sender basis for each receiver given in terms of the local probes only. This automatically decouples the effects of high-frequency near-field visibility and slowly-varying distant irradiance in a natural way without the need for increasing the number of hierarchical transport links (Figure 4), enabling high-quality indirect shadows with a small number of precomputed transport coefficients. The global probe transport operator can be implemented in a variety of ways (e.g. our precomputed gather from texture, rendering cube maps, hierarchical direct-to-indirect methods); all it needs to produce are the SH expansions of incident radiance at the probe locations.

3.4 Analysis

To analyze our interpolator, we study a flatland scenario depicted in Figure 5a. The scene consists of a planar (line) receiver at $y = 0$, and two textured surfaces at $y = 0.25$ (small, greenish segment) and $y = 2$ (larger, reddish segment). In addition, a bright area light is situated in the middle of the reddish surface. We study the light field incident onto points on the receiver as a function of space x and angle w using the angle parameterization of Chai et al. [2000]: the direction w of a ray emanating from a point x is specified by its intersection with the $y = 1$ line, such that an intersection point directly above x always means $w = 0$. The incident radiance field on the receiver is shown in Figure 5b, where the two blue example rays in Figure 5a correspond to the two marked points.

Figure 5c depicts a version of the light field that has been bandlimited (blurred) along the angular dimension, but with full resolution along the x axis. Note how the extremely bright light correspondingly “bleeds” along the angular dimension, but not along the spatial dimension: receiver points that do not see any part of the light source remain properly occluded. Ramamoorthi and Hanrahan’s [2001] and Basri and Jacobs’ [2003] deep results tell us that given a low-frequency BRDF at the receiver, we can reconstruct outgoing radiance from such bandlimited incident signal. *This permits us to seek reconstruction of this bandlimited signal instead of the original signal with full resolution over angle.*

Next, we introduce a sparse spatial sampling of the bandlimited light field (Figure 5d). This stage corresponds precisely to computing finite SH expansions over the angular domain at a sparse set of probe locations. (The sample columns have been fattened for visualization only; they are Diracs in x .)

Figure 5e illustrates naive interpolation from the samples using Equation 1. This corresponds to blending the probes’ SH coefficient vectors using the spatial weighting functions $w_i(\mathbf{x})$. This results in significant light leaking and blurring.

Our visibility-aware reconstruction is depicted in Figure 5f. For each ray (x, w) in the output light field, we consult nearby probes, and use their radiance value (which is bandlimited in angle) for reconstruction only if the point seen at the end of the ray (x, w) is visible to the probe too. This panel clearly illustrates our approximation: due to the bandlimited probes’ inability to represent discontinuities, the light source seen by the probe influences also surfaces that are nearby the light source in angle (diagonal light streaks along the edges of the green surface). However, the accuracy of the approximation increases with the angular bandwidth of the probes. This is seen in the final panel (Figure 5g) which illustrates reconstruction from the *same* set of 10 spatial probe locations, but with increased angular resolution.

Figure 6 depicts the interpolation in a realistic scenario. Radiance from two bandlimited probes is interpolated in a visibility-aware manner to a third receiver point and compared to the standard interpolation. In addition to the flatland example, this figure also includes a final convolution of the radiance field with the cosine kernel (panels j, k); this is the irradiance function seen by the receiver as a function of the surface normal. As can be seen, while the reconstruction of angular variation is not perfect, once the interpolated radiance function is convolved with the cosine kernel, the reconstruction can remain faithful.

4 IMPLEMENTATION

This section describes our practical global illumination technique using our factorized transport approach. We describe how we precompute and compress the local transport operator in Section 4.1. The runtime part, including details on how we evaluate the probe transport, is described in Section 4.2.

4.1 Precomputation

Probe Locations and Radii. Our goal is to find a sparse set of probe locations and support radii with the two properties: 1) every surface point is under the support of at least one probe, and, 2), every surface point is visible from at least one probe.




Fig. 5. 2D analysis of sampling and reconstruction of a diffuse radiance field from sparse spatial samples bandlimited in angle (cf. Section 3.4). a) Scene. The planar receiver sits at $y = 0$. Two surfaces at $y = 0.25$ and $y = 2$ shine radiance onto the receiver. A bright area light is situated in the middle of the further occluder. Rays emanating from the occluder are parameterized by their x coordinate and their intersection with the $y = 1$ line. b) The light field seen by the receiver over space x and angle w . The two rays from the previous panel are indicated. c) The light field, bandlimited (blurred) over angle, but not space. This bandlimited version suffices to reconstruct reflected radiance for low-frequency BRDFs at the receiver [Basri and Jacobs 2003; Durand et al. 2005; Ramamoorthi and Hanrahan 2001]. d) A sparse spatial sampling of the angle-bandlimited light field. e) Reconstruction from the samples by spatial blending only (Equation (1); no reprojection/shearing, no visibility). f) Our reconstruction from the same samples. g) Our reconstruction from the same samples, but with increased angular bandwidth (samples are sharper in w). As all surfaces are seen by a probe, the result converges to ground truth with increasing bandwidth — even with this fixed set of spatial samples.

To find the probe locations, we use a greedy algorithm with a single parameter ρ_{probes} , the desired spacing. To avoid placing probes inside scene geometry, we first voxelize the scene and flood-fill the empty interior space. Next, we generate a vastly overcomplete set of candidate locations \mathbf{p}_i in empty voxels near the scene surfaces, i.e., in empty voxels with non-empty neighbors. We then compute a kernel density estimate from the candidates \mathbf{p}_i using the weight function $w_i(\mathbf{x}; \rho_{\text{probes}})$ (see below), and iteratively remove candidates from the densest regions until we have met our target probe count. To avoid introducing additional parameters, we set the target probe count to the number of points in a regular grid that covers the scene with grid spacing set to ρ_{probes} .

Every receiver should fall under the support of at least one probe, but smooth spatial interpolation necessitates coverage from multiple probes for each receiver. The user provides the desired number N_{overlap} of overlapping probes per receiver. Since the probe set obtained in previous step has approximately constant density, we use a fixed radius r for all probes, i.e., $r_i \equiv r$. To find r , we perform a search over possible radii and choose the one that best satisfies the overlap constraint on average. All our results use $N_{\text{overlap}} = 10$.

Spatial Weight Function. For spatial interpolation, we use a radial weight kernel $w_i(\mathbf{x}; r) = w(\|\mathbf{x} - \mathbf{p}_i\|_2/r)$, where r is the radius obtained in the previous step, and $w(t)$ is given by

$$w(t) = \begin{cases} 2t^3 - 3t^2 + 1, & \text{if } 0 \leq t \leq 1 \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

Precomputed Radiance Transport and Compression. We proceed to precompute transport coefficients $\alpha_{\mathbf{x}}$ or $\mathbf{B}_{\mathbf{x}}$ for all receivers. (Section 5 shows results for both irradiance (scalar) and 1st order SH transport.) Receivers are only linked to probes whose support radius covers the receiver: other probes contribute nothing as per Equation (2). The coefficients are then compressed using Clustered Principal Components Analysis [Sloan et al. 2003]: receivers are split to clusters, and the transport matrices in each cluster approximated separately. More precisely, the transport vectors/matrices of all receivers in a cluster c are stacked vertically into the matrix \mathbf{T}_c ,

and the Singular Value Decomposition applied:

$$\mathbf{T}_c = \mathbf{U}\Sigma\mathbf{V}^T \approx \mathbf{U}_c\Sigma_c\mathbf{V}_c^T. \quad (11)$$

Compression is obtained by replacing the diagonal matrix of singular values Σ with a truncated version Σ_c , where only the n_c largest entries have been kept. The cluster projection matrix $\Sigma_c\mathbf{V}_c^T$ has size $n_c \times n_i n_j$ (recall that n_i is the total number of probes and n_j the number of basis functions captured by the probes). As only few receivers contribute to a spatially coherent cluster of receivers, large blocks of the columns of $\Sigma_c\mathbf{V}_c^T$ are zeros — it suffices to store the non-zero columns. The receiver reconstruction coefficient matrix \mathbf{U}_c encodes the linear response of each receiver in the cluster to each of the n_c dimensions of the intermediate cluster-specific illumination basis; i.e., it contains $n_c n_k$ numbers per receiver, where n_k is the number of rows in the radiance transport matrix $\mathbf{B}_{\mathbf{x}}$ (in case of irradiance, $n_k = 1$). The cluster projection matrix $\Sigma_c\mathbf{V}_c^T$ and reconstruction coefficients \mathbf{U}_c are stored in 16-bit floating point.

As computing the SVD for very large matrices is inefficient, we build, in departure from Sloan et al. [2003], our clusters in two steps. An initial clustering is obtained by building an AABB-tree over the receiver positions, splitting the longest axis until the leaf nodes contain fewer than 1024 points. The clusters are then adaptively refined [Sloan et al. 2003] based on a fixed error threshold and a maximum number of SVD coefficients (we use a threshold of 0.005 and a maximum of 32 coefficients).

4.2 Runtime

Probe Transport. Our method places no restrictions on how to compute the probe transport, i.e., the probes' SH coefficients λ at runtime. We choose to perform the projection with the aid of a direct illumination light map that is updated every frame, accounting for the diffuse albedo and emissivity of each surface point. For each probe, we use a fixed set of uniformly distributed “relight rays”, and precompute the direct light map uv coordinates of the relight ray hit points ahead of time by tracing visibility rays (this does not preclude dynamic blockers; see section 5.1). At runtime, the probes merely

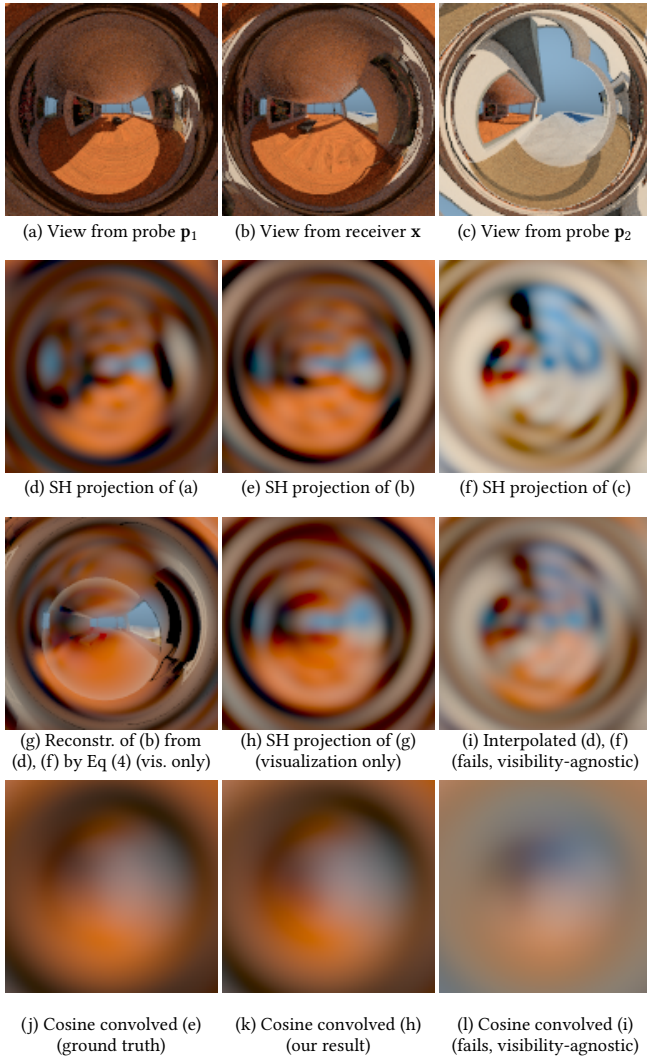


Fig. 6. Illustration of visibility-aware interpolation. (a)-(c) Fisheye views from probes p_1 , p_2 , and receiver x . (d)-(f) Spherical harmonic approximations of the angular radiance field at the probes and the receiver. (d) and (f) are computed at runtime; the expansion (e) is for visualization only, and is never formed explicitly. (g), (h) visualize the interpolated radiance field seen by the receiver, as evaluated by Equation (4), and its SH projection. In a perfect reconstruction, (g) = (b) and (h) = (e). This spherical image is also never explicitly formed. (i) visualizes the SH expansion obtained by directly interpolating the SH coefficients $\lambda_{1,2}$ (cf. Figure 2a). This fails due to lack of visibility awareness. (j), (k) The diffuse BRDF acts as a low-pass filter to the radiance field [Ramamoorthi and Hanrahan 2001]. Even though the radiance field is not matched precisely, the low frequencies remain relatively intact through our entire pipeline. This is apparent in the match between (j) and (k) (fourth row).

loop over all their relight rays, consult the current radiance at the appropriate uv coordinates in the direct illumination texture, and multiply by the spherical harmonics at the appropriate direction. To avoid evaluating spherical harmonics at runtime, we precompute their values together with the proper PDF-weights in the directions of the relight rays.

Cluster Radiance Transport. Once obtained, the probe radiance vector λ is transformed by the PCA clusters' projection matrices $\Sigma_c V_c^T$ to yield the n_c dimensional light basis vector for each cluster. These coefficients are finally multiplied by the cluster's receiver reconstruction coefficients U_c to obtain the n_k -dimensional transported radiance coefficients at each receiver. The result is then used to compute the reflected radiance by evaluating the receiver BRDF. For scalar irradiance transport, this reduces simply to multiplication with the local albedo and division by π . Higher-order transport can be evaluated in various ways. In our results, we use an order-1 version of the irradiance convolution of Ramamoorthi and Hanrahan [2001] to support normal mapping. This means scaling the transported SH coefficients by fixed constants and evaluating the resulting 4-term expansion in the direction of the surface normal.

To account for infinite light bounces, the resulting diffuse component of the outgoing radiance is fed back to the system and used as input for the next direct-to-indirect transport iteration.

5 RESULTS

We have implemented our method using C++ and DirectX 11, using compute shaders to evaluate the runtime radiance transport. All images in this paper contain the raw output of our method without any additional screen-space ambient occlusion or reflection effects. In all results, we define the receivers as the sample points of a light map's texels for ease of final reconstruction from the camera's point of view (for resolutions, see Table 1; note that any other set of receivers can be used just as well). The accompanying video features animations with dynamic lights. In addition to the standard CORNELL and SPONZA, we demonstrate the method on two novel scenes, GALLERY from a "triple-A" Xbox One title, and BRUTALIST HALL from an unannounced title. All results are computed on a Intel Core i7-3770K CPU PC with 32 GB memory and NVIDIA Titan X Pascal GPU. Unless otherwise noted, all results are computed with probes of SH order 7 (64 basis functions) and irradiance (scalar) transport (Equation 7). All scenes use a direct illumination light map of size 2048^2 to gather direct illumination to the surfaces prior to computing the probe radiance vectors. We use 8000 relight rays per probe.

Comparison to Ground Truth. A sampling of our results is shown in Figure 7. The figure compares our irradiance transport against a converged reference, where the irradiance at the same light map texels is computed using path tracing. In order to reveal all the intricate lighting details, we show our comparisons without albedo maps, as they tend to hide lighting imperfections. Please see the accompanying video for dynamic illumination in the same scenes. While using only less than 90MB of GPU memory and $< 5ms$ of GPU time for indirect lighting, we observe a generally good match with the reference with the chosen parameters. Section 5.2 presents analysis of errors. Table 1 provides detailed statistics of our scenes and




Fig. 7. We compare our results (left and middle) to path traced reference (right) using three scenes, (GALLERY, SPONZA, and BRUTALIST HALL) and two different lighting configurations in each scene. The probe positions are shown in insets in the left column. We show the comparisons as lighting-only images assuming a constant albedo of (0.5, 0.5, 0.5). Note that the albedo maps still contribute to the indirect lighting. The supplemental material contains all images in uncompressed form. Please see the supplemental video for dynamic illumination.




Fig. 8. We compare our interpolation (middle right) to spatial interpolation (middle left) and to radiance caching with irradiance decomposition [Arikan et al. 2005] (middle). The images show indirect illumination from two point light sources with varying number of probes (left). Our method is closer to the path traced reference (right) in both cases and yields superior results in the sparse probe setting.

Table 1. Statistics and memory consumption for all test scenes. Precomputation time contains both radiance and irradiance transport. For performance figures, see Table 2.

Scene	Precomp. time	Transport statistics				Compression statistics and memory		
		Lightmap resolution	#Probes	#Receivers	#Clusters	Avg. #SVD coeffs./receiver	Avg. #Probes / SVD cluster	Total memory
GALLERY	52 min	1216 ²	275	321 745	2311	15.56	17.55	86MB
SPONZA	45 min	1216 ²	150	246 325	935	24.41	20.17	72MB
BRUTALIST HALL	72 min	1088 ²	260	567 370	2223	18.89	13.40	89MB

Table 2. Timing breakdowns for our test scenes.

Scene	Scene relight (ms)	Probe relight (ms)	Cluster SVD projection (ms)	Receiver irradiance (ms)	Total GPU (ms)
GALLERY	0.70	0.84	2.06	0.30	3.90
SPONZA	0.36	1.01	1.00	0.13	2.50
BRUTALIST HALL	1.28	1.55	1.80	0.24	4.87

transport parameters, including the number of probes and memory usage. See Table 2 for a timing breakdown for the indirect lighting. In all scenes, the full rendering is performed at 90 FPS.

Dynamic Material Properties. Since the precomputed transport coefficients do not encode any material information, we support dynamic surface material properties, such as textured diffuse albedo, emissivity, normal maps and surface roughness, all of which can be freely changed or animated at runtime. Figure 9d demonstrates area shadows from a door with an emissive surface material. See the accompanying video for an interactive session.

Temporal Stability. The accompanying result video demonstrates the our method is temporally stable and does not suffer from flickering artifacts. Temporally stable results are achieved by maintaining a sufficient sample rate when sampling the input illumination.

Comparison to Previous Methods. Figure 8 shows a comparison between our method and the local-global irradiance separation of Arikan et al. [Arikan et al. 2005]. Like ours, their method interpolates spherical harmonics illumination vectors from nearby probes. In contrast to our technique, they proceed to heuristically subtract radiance from local occluders processed one triangle at a time, resulting in double occlusion from overlaid occluders due to lack of

proper occluder fusion. To compare, we implemented an idealized version of their technique using accurately sampled visibility instead of their approximation; this effectively compares against the best-case output of their method. Spatial-only interpolation (Equation 1) is included as a reference in Figure 8; it does not yield usable results with a small number of probes. Using the same sparse set of probes, the interpolation of Arikan et al. contains visible artifacts. When the probe count is increased, both comparison methods start to resemble the ground truth, but our method is closest to the path traced reference in both cases. Notably, our interpolator yields a high-quality result also with the sparse probe set.

Discussion. The uncompressed local transport matrix has dimensions $Receivers \times Probes \cdot SHBasisFunctions$ and the probe transport matrix has dimensions $Probes \cdot SHBasisFunctions \times Probes \cdot RelightRays$. Typical numbers are 500K receivers, 200 probes, 64 SH basis functions and 8000 relight rays. Compared to the full transport matrix consisting of all senders and receivers, the total number of coefficients in our factorized transport correspond to an effective sparsity factor of the order 10^{-4} for uncompressed matrices. With SVD compression, the observed effective sparsity factor is of order 10^{-5} . Previously, Lehtinen et al. [2008] reported sparsity factors of order 10^{-3} . However, the results are not directly comparable, as the reported numbers are from different scenes and using a different receiver/sender parametrization.

5.1 Extensions

Normal Mapping. Normal maps are used to add high-frequency surface details without increasing the triangle count of the scene. Unfortunately, scalar transport is insufficient to support normal mapping at the receivers, as the irradiance varies with the local surface normal. To enable normal mapping, we instead project the interpolated radiance signal to order-1 spherical harmonics (4 coefficients) at the receivers (Equation 8), allowing us to apply an order-1 version of Ramamoorthi and Hanrahan’s [2001] efficient irradiance convolution in the SH domain. To counter the increased memory use due to the richer representation (4 instead of 1 transport coefficients per receiver), we simultaneously lower the probes’ SH order to 4 (25 coefficients). These changes result in approximately the same memory use as scalar transport using order-7 probes. Figure 9a shows an example. Naturally, the overall accuracy suffers due to the probes’ lower angular bandwidth, but results remain plausible; increasing the order at both receivers and the probes allows more faithful results at the cost of more memory and GPU time.

Glossy Materials. While a low-order expansion allows normal mapping on diffuse surfaces, the low bandwidth of the angular variation at the receivers is not sufficient to enable accurate rendering of glossy or specular materials; in practice, increasing the receiver SH order sufficiently quickly ends up in a storage and runtime problem. We can, however, support an approximate form of specular transport by interpreting the linear terms in the SH expansion as a directional light source [Sloan 2008], with the diffuse component given by the DC term, and apply any standard BRDF model to this approximation. Figure 9b depicts an example (only specular illumination shown). Here, we use a Cook-Torrance BRDF with GGX normal distribution and Smith shadowing-masking term [Walter et al. 2007]. Although

this is a crude approximation, it provides visually pleasing results at negligible additional cost.

Dynamic Occluders. While our basic technique is fixed to a scene through precomputation, the local supports of the probes allows insertion of dynamic blockers; intuitively, as probes only affect receivers close by, changes in radiance seen by the probes due to moving objects have at least some hope of being captured in the local interactions. Figure 9c demonstrates a soft area-light shadows from a dynamic sphere occluder blocking the relight rays. Naturally, the accuracy of the approximation depends on the size of the occluder relative to the probe density: detailed occlusion from small occluders is generally not possible with a sparse probe set. However, larger objects, such as dynamic doors and windows, yield usable results in many cases, and we expect this to increase the applicability of our method in practical applications. Please see the accompanying video for an interactive session.

5.2 Error Analysis

Reconstruction Error. We analyze the expected error in irradiance reconstruction by studying how the incident light field changes at a fixed receiver by a round trip to the probes and back. That is, taking a ground truth angular radiance field at the receiver, we use our interpolation technique to project it *backwards* to the probes, then back to the receiver, and comparing the two. The better the reconstruction enabled by the probe set, the smaller the difference.

To enable numerical estimation, we represent the incident radiance at a fixed receiver \mathbf{x} using a order 9 (100-coefficient) SH expansion, called the *receiver radiance vector* $\lambda_{\mathbf{x}}$. The receiver radiance is first projected to the probes’ SH basis via a dual projection operator \mathcal{P}^* defined analogously to the forward one: each probe takes the role of a receiver, and the receiver becomes the only probe and has infinite support. This gives us the probe radiance vector $\lambda = \mathcal{P}^* \lambda_{\mathbf{x}}$ that represents the receiver light field projected to the probes’ perspectives. To pull the radiance back to the receiver, we apply our reconstruction operator $\mathcal{P}_{\mathbf{x}}\{\lambda\}$ to the probe radiance vector λ , and measure the reconstructed irradiance with a cosine measurement functional. Putting this together, the squared irradiance reconstruction error $E(\mathbf{x})$ for a fixed receiver radiance vector $\lambda_{\mathbf{x}}$ is

$$E(\mathbf{x}; \lambda_{\mathbf{x}}) = \|\langle \cos \theta, (\mathbf{I} - \mathcal{P}_{\mathbf{x}}\{\lambda\} \mathcal{P}^*) \lambda_{\mathbf{x}} \rangle\|_2^2, \quad (12)$$

where \mathbf{I} is the identity operator. We note that the error is quadratic in the probe radiance vector, and that we can consequently measure the average round-trip error over all possible unit radiance vectors $\lambda_{\mathbf{x}}$ by computing the L_2 operator norm of the linear operator $\langle \cos \theta, (\mathbf{I} - \mathcal{P}_{\mathbf{x}}\{\lambda\} \mathcal{P}^*) \cdot \rangle$. The inclusion of the cosine measurement term ensures that we measure error in irradiance reconstruction (the final result we seek), not the intermediate radiance reconstruction. The norm is visualized in the bottom row of Figure 10 in CORNELL. As can be seen, increasing the density of probes causes a reduction in the expected error.

Probe Density vs. Quality. Figure 10 visualizes the effect of probe density to reconstruction error. The first rows visualize the transported irradiance in CORNELL with an increasing probe count in a fixed lighting condition and using order-7 probes (64 coefficients).

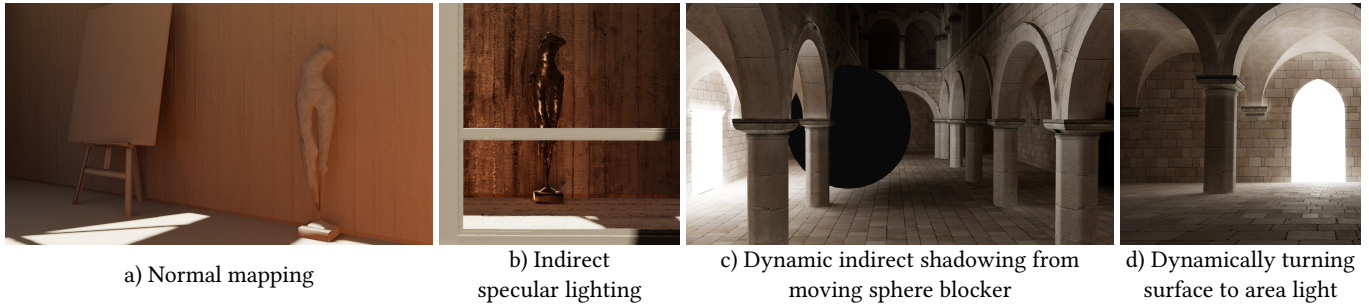


Fig. 9. Extensions (see Section 5.1 and the accompanying video). a) Normal mapping through 1st-order transport to receivers. b) Approximate indirect specular lighting by interpreting 1st order radiance at receivers as a directional light. c) A dynamic sphere blocker used for modulating the visibility during probe re-light. d) The emission of any surface may be changed at will at no extra runtime cost. Here the polygons in a door in SPONZA have been turned emissive; soft shadows automatically result.

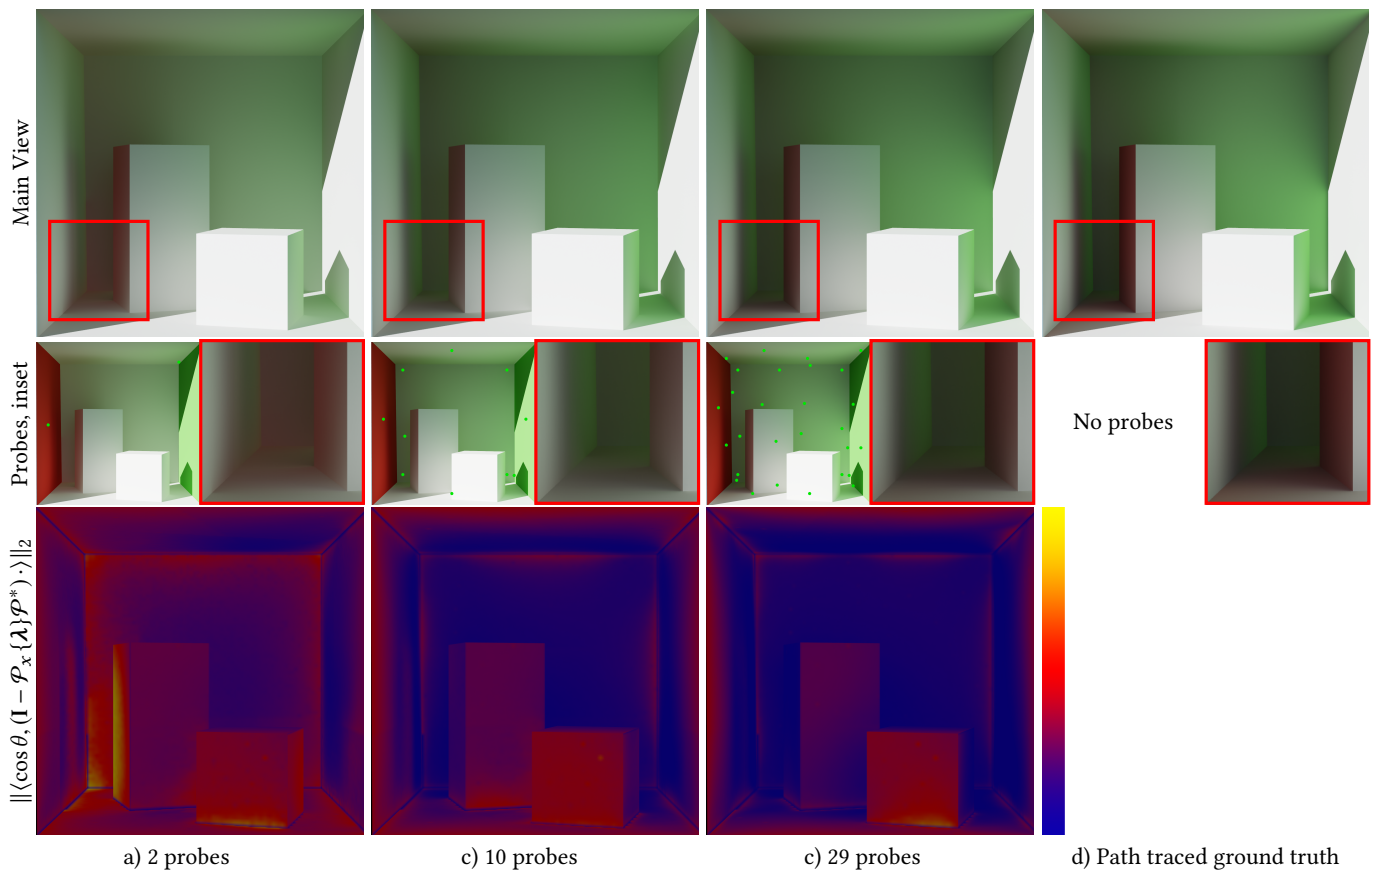


Fig. 10. Effect of the density of light probes on the accuracy of illumination using order-7 probes (64 coefficients). Top two rows: The Cornell box in fixed illumination. Only incident irradiance is visualized for added clarity (surface albedo not shown, but causes color bleeding). The thumbnails on the middle rows show the position of the probes. Even with only two probes, the indirect illumination is captured relatively faithfully. Loss of detail in small-scale local indirect shadows can be noticed, but additional probes remedy the situation. Notice, in particular, the erroneous red color bleeding in the inset in column a). Bottom row: Visualization of the operator norm of the forward-and-backward projection operator (Section 5.2). The color denotes the average L_2 error over all possible illumination conditions. Note how areas of high error correspond to the regions of poorer fidelity in the result images above.




Fig. 11. Effect of the order of the probes' SH expansions on the final result. Lighting only (top). Difference to ground truth (16x magnified, bottom). Differences between successive approximations quickly taper off to the point it is hard to see them in a side-by-side visualization; the reader is encouraged to flip back and forth between the full-resolution images provided in the supplemental material.




Fig. 12. This figure demonstrates a failure case due to insufficient probe density and bandwidth. Path-traced reference (middle) contains a high-frequency secondary area light source (green), and our reconstruction (left) has blurred the incident radiance field in the near vicinity (red). The error is localized to the bottom-right region (right), where the probe density and bandwidth are too low to represent the incident illumination, leaving most of the image unaffected.

Only irradiance is shown, as final multiplication by the albedo hides problems even in this simple scene. The small thumbnails in the middle row show the positions of the probes. Generally, the irradiance is reconstructed remarkably well even from just 2 probes. However, erroneous color bleeding and washed-out indirect shadows can be observed (see inset). The bottom row shows a heat-map visualization of the L_2 operator norm of the back-and-forth transport operator that attempts to reconstruct the receivers' original radiance field by projecting to the probes and back; this visualization is not tied to a particular lighting condition (see above). We note how the areas of worst mismatches between our result and the reference can be found in the areas of higher error. As expected, increasing the density of the probes reduces the expected error.

Probe SH Order vs. Quality. Figure 11 visualizes the quality of transported irradiance at a fixed lighting condition and fixed probe set as a function of the probes' bandwidth, i.e., the number of basis functions n_j . Results remain plausible even at very low orders, and fidelity quickly increases to the point it is difficult to see the differences side by side. The reader is encouraged to flip back-and-forth in the full-resolution images found in the supplemental material. Interestingly, the quality difference resulting from lowering the probes' SH order can be countered by a denser spatial sampling: the closer the probes are to the receivers, the less bandwidth they need. Previous analysis show that in the limit, when probes coincide with

receivers, order 2 is sufficient for high-fidelity reconstruction [Ramamoorthi and Hanrahan 2001]. A full study of the 2D parameter space remains future work.

A Typical Failure Case. As shown by the various analyses above, the reconstruction error of our method is directly linked to probe density and bandwidth. Figure 12 shows a typical failure where the angular bandwidth at the probes is insufficient to capture the local light field in enough detail given the sparsity of the probes.

6 CONCLUSION

We have described a technique for interpolating radiance from sparse samplings of a diffuse light field, where the angular bandwidth (resolution) of the samples is limited due to their representation as spherical harmonics expansions. Despite these limitations (sparsity, low bandwidth), our method allows a much more faithful reconstruction of the incident light field than previous techniques that make use of such sparse sample sets.

It remains an interesting challenge to incorporate light field analysis tools in the sampling stage to drive the positions and bandwidths of the light probes. Furthermore, it is intriguing to see if our interpolation on surfaces seen by samples and receivers (as opposed to locally in the space-direction parameterization near the receiver) enables sampling of also non-diffuse light fields using spatial samplings that are sparser than previous local analyses of sampling and reconstruction would indicate.

REFERENCES

- Okan Arikan, David A. Forsyth, and James F. O'Brien. 2005. Fast and Detailed Approximate Global Illumination by Irradiance Decomposition. *ACM Trans. Graph.* 24, 3 (July 2005), 1108–1114. <https://doi.org/10.1145/1073204.1073319>
- R. Basri and D. W. Jacobs. 2003. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25, 2 (Feb 2003), 218–233. <https://doi.org/10.1109/TPAMI.2003.1177153>
- Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. 2001. Unstructured Lumigraph Rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. ACM, New York, NY, USA, 425–432. <https://doi.org/10.1145/383259.383309>
- Jin-Xiang Chai, Xin Tong, Shing-Chow Chan, and Heung-Yeung Shum. 2000. Plenoptic sampling. In *Proc. ACM SIGGRAPH 2000*. 307–318.
- Per Christensen. 2008. Point-Based Approximate Color Bleeding. Pixar Technical Memo #08-01. (2008).
- Per H. Christensen and Wojciech Jarosz. 2016. The Path to Path-Traced Movies. *Foundations and Trends? in Computer Graphics and Vision* 10, 2 (2016), 103–175. <https://doi.org/10.1561/06000000073>

- Cyril Crassin, Fabrice Neyret, Miguel Sainz, Simon Green, and Elmar Eisemann. 2011. Interactive Indirect Illumination Using Voxel Cone Tracing: A Preview. In *Symposium on Interactive 3D Graphics and Games (I3D '11)*. ACM, New York, NY, USA, 207–207. <https://doi.org/10.1145/1944745.1944787>
- Holger Dammert, Daniel Sewtz, Johannes Hanika, and Hendrik P. A. Lensch. 2010. Edge-avoiding À-Trous wavelet transform for fast global illumination filtering. In *Proc. High Performance Graphics 2010*. 67–75.
- Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan, and François X. Sillion. 2005. A Frequency Analysis of Light Transport. *ACM Trans. Graph.* 24, 3 (2005), 1115–1126.
- Kevin Egan, Florian Hecht, Frédo Durand, and Ravi Ramamoorthi. 2011. Frequency Analysis and Sheared Filtering for Shadow Light Fields of Complex Occluders. *ACM Trans. Graph.* 30, 2 (2011), 9:1–9:13.
- Kevin Egan, Yu-Ting Tseng, Nicolas Holzschuch, Frédo Durand, and Ravi Ramamoorthi. 2009. Frequency Analysis and Sheared Reconstruction for Rendering Motion Blur. *ACM Trans. Graph.* 28, 3 (2009), 93:1–93:13.
- Steven Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael Cohen. 1996. The Lumigraph. In *Proc. SIGGRAPH '96*. 43–54.
- Gene Greger, Peter Shirley, Philip M. Hubbard, and Donald P. Greenberg. 1998. The Irradiance Volume. *IEEE Comput. Graph. Appl.* 18, 2 (March 1998), 32–43. <https://doi.org/10.1109/38.656788>
- Miloš Hašan, Fabio Pellacini, and Kavita Bala. 2006. Direct-to-indirect Transfer for Cinematic Relighting. *ACM Trans. Graph.* 25, 3 (July 2006), 1089–1097. <https://doi.org/10.1145/1141911.1141998>
- J Hooker. 2016. Volumetric Global Illumination at Treyarch. *Advances in Real-Time Rendering Course*.
- Johannes Jendersie, David Kuri, and Thorsten Grosch. 2016. Precomputed Illuminance Composition for Real-time Global Illumination. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D '16)*. ACM, New York, NY, USA, 129–137. <https://doi.org/10.1145/2856400.2856407>
- Janne Kontkanen, Jussi Räsänen, and Alexander Keller. 2004. Irradiance Filtering for Monte Carlo Ray Tracing. In *Monte Carlo and Quasi-Monte Carlo Methods 2004*. Springer, 259–272.
- Janne Kontkanen, Emmanuel Turquin, Nicolas Holzschuch, and François X. Sillion. 2006. Wavelet Radiance Transport for Interactive Indirect Lighting. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques (EGSR '06)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 161–171. <https://doi.org/10.2312/EGWR/EGSR06/161-171>
- Jaroslav Křivánek, Pascal Gautron, Sumanta Pattanaik, and Kadi Bouatouch. 2005. Radiance caching for efficient global illumination computation. *IEEE Trans. Vis. and Compu. Graph.*, 11, 5 (2005), 550–561.
- Jaroslav Křivánek, Pascal Gautron, Kadi Bouatouch, and Sumanta Pattanaik. 2005. Improved Radiance Gradient Computation. In *Proceedings of the 21st Spring Conference on Computer Graphics (SCCG '05)*. 155–159. <https://doi.org/10.1145/1090122.1090148>
- Jaakko Lehtinen, Timo Aila, Jiawen Chen, Samuli Laine, and Frédo Durand. 2011. Temporal Light Field Reconstruction for Rendering Distribution Effects. *ACM Trans. Graph.* 30, 4 (2011), 55:1–55:12.
- Jaakko Lehtinen, Matthias Zwicker, Emmanuel Turquin, Janne Kontkanen, Frédo Durand, François X. Sillion, and Timo Aila. 2008. A meshless hierarchical representation for light transport. *ACM Trans. Graph.* 27, Article 37 (August 2008), 37:1–37:9 pages. Issue 3.
- Bradford J. Loos, Lakulish Antani, Kenny Mitchell, Derek Nowrouzezahrai, Wojciech Jarosz, and Peter-Pike Sloan. 2011. Modular Radiance Transfer. In *Proceedings of the 2011 SIGGRAPH Asia Conference (SA '11)*. ACM, New York, NY, USA, Article 178, 10 pages. <https://doi.org/10.1145/2024156.2024212>
- Sam Martin. 2010. A Real Time Radiosity Architecture for Video Games. *Advances in Real-Time Rendering Course*.
- Morgan McGuire, Michael Mara, Derek Nowrouzezahrai, and David Luebke. 2017. Real-Time Global Illumination using Precomputed Light Field Probes. In *I3D 2017*. 11. <http://graphics.cs.williams.edu/papers/LightField3D17>
- Ravi Ramamoorthi and Pat Hanrahan. 2001. An Efficient Representation for Irradiance Environment Maps. In *Proceedings of the ACM SIGGRAPH conference*. ACM.
- Tobias Ritschel, Thomas Engelhardt, Thorsten Grosch, Hans-Peter Seidel, Jan Kautz, and Carsten Dachsbacher. 2009. Micro-Rendering for Scalable, Parallel Final Gathering. *ACM Trans. Graph.* 28, 5 (2009), 132:1–132:8.
- Jorge Schwarzhaupt, Henrik Wann Jensen, and Wojciech Jarosz. 2012. Practical Hessian-Based Error Control for Irradiance Caching. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31, 6 (Nov. 2012). <https://doi.org/10.1145/2366145.2366212>
- Peter-Pike Sloan. 2008. Stupid Spherical Harmonics (SH) Tricks. (2008). www.ppsloan.org/publications/StupidSH36.pdf
- Peter-Pike Sloan, Jesse Hall, John Hart, and John Snyder. 2003. Clustered Principal Components for Precomputed Radiance Transfer. *ACM Transactions on Graphics* 22, 3 (July 2003). Proceedings of the ACM SIGGRAPH conference.
- Peter-Pike Sloan, Jan Kautz, and John Snyder. 2002. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. *ACM Transactions on Graphics* 21, 3 (2002). Proceedings of the ACM SIGGRAPH conference.
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 195–206. <https://doi.org/10.2312/EGWR/EGSR07/195-206>
- Gregory J. Ward and Paul Heckbert. 1992. Irradiance gradients. In *Proc. Eurographics Workshop on Rendering '92*. 85–98.
- Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. 1988. A ray tracing solution for diffuse interreflection. In *Proc. ACM SIGGRAPH '88*. 85–92.