

Dynamic Diffuse Global Illumination Resampling

Zander Majercik¹, Thomas Müller¹, Alexander Keller¹, Derek Nowrouzezahrai², and Morgan McGuire³

¹NVIDIA ²McGill University ³Roblox




Figure 1: The indirectly lit GREEK VILLA scene rendered at 1920×1080 with 1 sample per pixel (spp) on an i7 6800k CPU and RTX 3090 GPU. Direct illumination resampling, such as ReSTIR [BWP*20] (leftmost inset), reduces noise at the primary path vertex but does not affect the remainder of the path which is both noisy and expensive to trace. Therefore, we propose replacing the remainder of the path, starting from the secondary vertex, with a cheap, noise-free approximation: an extended variant of DDGI [MMSM21] (second inset). Our key observation is that DDGI thereby acts as a light source, allowing us to include it in the resampling algorithm (third inset). This unifies the sampling of direct and indirect illumination. Combined with denoising (fourth inset), scenes with complex indirect illumination such as this one can be rendered in real time from 1 spp, with quality approaching offline path tracing with hundreds of samples per pixel (fifth inset).

Abstract

Interactive global illumination remains a challenge in radiometrically and geometrically complex scenes. Specialized sampling strategies are effective for specular and near-specular transport because the scattering has relatively low directional variance per scattering event. In contrast, the high variance from transport paths comprising multiple rough glossy or diffuse scattering events remains notoriously difficult to resolve with a small number of samples. We extend unidirectional path tracing to address this by combining screen-space reservoir resampling and sparse world-space probes, significantly improving sample efficiency for transport contributions that terminate on diffuse scattering events. Our experiments demonstrate a clear improvement – at equal time and equal quality – over purely path traced and purely probe-based baselines. Moreover, when combined with commodity denoisers, we are able to interactively render global illumination in complex scenes.

CCS Concepts

- **Computing methodologies** → **Ray tracing**;

1. Introduction

Modern physically-based production renderers rely primarily on variants of the unidirectional path tracing algorithm [Bal18], in which paths are traced from the camera and scattered in the scene until they reach an emitter. Path tracing scenes with high geometric and radiometric complexity can result in high-variance (i.e., noisy) images unless they take a prohibitively large number of samples. To alleviate this, many Monte Carlo variance reduction-based solutions have been proposed, such as variants of importance sampling [VG95; OZ98; TCE05; BWP*20] and path guiding [VHH*19]. However, even with advanced sampling techniques in place, hundreds of samples per pixel are typically required to produce a converged image, which is prohibitive for real-time applications.

Even though recent improvements in GPU ray tracing and effective denoising have greatly improved path tracing performance for complex scenes, real-time path tracing budgets remain limited to only few indirect scattering events (often just one) and scenes with only modestly complex illumination. As such, interactive graphics methods necessarily resort to approximate global illumination methods *in addition* to importance sampling. One such recent approach combines world-space irradiance probes with visibility-aware interpolation to rapidly approximate multi-bounce indirect illumination (DDGI) [MGNM19]. Crucially, the normal-dependent irradiance can also be interpreted as cosine-prefiltered radiance [MMSM21] in the normal direction, allowing its use in both rough glossy as well as Lambertian diffuse transport. As with most approximate solutions, these *dynamic probes* introduce bias – in the form of the transport that is computed – in exchange for a noise-free and fast result.

In contrast to interactive indirect illumination sampling, the sample efficiency and approachable scene complexity for ray-traced direct illumination has been greatly improved by recent a spatio-temporal resampling scheme (ReSTIR) [BWP*20]. In the graphics pipeline, sampling is followed by full-screen post-processing, including denoising [BRM*16; BB17; SKW*17; VRM*18; BVM*17; CKS*17; MMBJ17; XZW*19; HY21]. So, the goal of a modern real-time sampling algorithm is not a fully converged image, but rather one with sufficiently low noise that after post-processing it is acceptably close to a converged image for the application. We address the efficient sampling problem in this work, assuming a downstream denoiser. We note that today’s commodity denoisers both desirably conceal sampling noise and also undesirably exhibit various residual temporal errors that are beyond the scope of this work. Hence, we show the impact of denoising in Fig. 1 for a static image to validate the unconverged targets, and then show the pre-denoising output of our method in all other results.

We present a new real-time global illumination method that combines the advantages of spatio-temporal resampling (improved sample efficiency for stochastic estimators) and dynamic diffuse probe volumes (smooth, interpolation-friendly transport proxies). Using ReSTIR for direct illumination and separately using DDGI for indirect illumination (see Fig. 8, “Primary DDGI” column) will neither eliminate the probe grid artifacts nor address the discrepancy between the two methods in terms of their respective artifact/noise characteristics: DDGI’s noise-free and biased indirect illumination, and ReSTIR’s noisy and unbiased direct illumination.

We alleviate this discrepancy by postponing DDGI queries by one

bounce (from the eye) – where sample contributions are evaluated during ReSTIR (see Fig. 8, “DDGI in ReSTIR” column) – similarly to *final gathering* with photon mapping [Jen96]. This hides a large portion of DDGI’s bias and allows us to combine direct and indirect samples during resampling, selecting each sample proportional to its *global* transport contribution. Where the original ReSTIR approach would normally resample direct illumination, it now treats both direct and indirect samples, each with similar noise characteristics due to the final gathering-like postponed DDGI sample, resulting in a sample efficient global shading estimate that is amenable to commodity denoising methods (see Fig. 1).

Fundamental to our technique is the combination of resampling (ReSTIR) with global illumination caching (DDGI). Specifically, by deferring the query of our cache by one bounce, we apply resampling to exitant radiance caches instead of directionally resolved caches at the primary vertex, which would be too costly at adequate resolution.

Our method takes a principled approach to balancing the performance-quality trade-off inherent to any combination of DDGI and ReSTIR, as evidenced in our analysis and benchmark against purely path traced and purely probe-based variants (Sec. 5). In contrast to unbiased resampling methods such as that of Ouyang et al. [OLK*21a], we intentionally trade bias for noise for a smoother result.

2. Background and Related Work

(Ir)radiance Caching. Most modern real-time global illumination caches have their roots in Ward’s seminal irradiance cache [WRC88] and other early approaches to precomputed light transport [Arv86; Hec90]. While initial work focused on adapting precomputed light transport to the harsh constraints of real-time rendering in *static* settings [Abr97; Oat05; ME10], many recent techniques offer at least partial dynamics, such as dynamic lighting through precomputed radiance transfer [SKS02]. Orthogonally, radiance caching techniques overcome the Lambertian assumption by including the directional domain [KGPB05]. Although today’s selection of high-quality *partially dynamic* techniques is vast [GS12; SL17; SSS*20; VPG14; SNRS12; RZD14; SJJ12], our goal is *fully dynamic* global illumination with as few constraints as possible.

To this end, we build upon the DDGI volume [MGNM19], which consists of a 3D grid of directionally resolved irradiance probes that is updated in real-time through hardware ray-tracing. Crucially, DDGI contains visibility information to prevent light leaking and it can approximate sufficiently rough *glossy* transport by re-purposing its directional dependence [MMSM21]. These features make it – with mild modifications (Sec. 3) – a reasonably versatile approximation of global illumination for most phenomena other than specular reflections. By querying DDGI *at secondary path vertices*, we employ a similar strategy to final gathering to conceal DDGI’s limitations (e.g., coarse spatial discretization) behind the blurring effect of the primary scattering interaction at the cost of sampling noise. We then address the sampling noise by applying spatio-temporal reservoir sampling [BWP*20].

Müller et al. [MRNK21] use a real-time trained neural radiance cache to approximate fully-dynamic global illumination. They evaluate their cache at later path vertices to hide the artifacts of their

neural network, but unlike this paper, they do not employ specialized importance sampling. Crucially, their method is compatible with ours in that the caches can be interchanged in future investigations.

Importance Sampling. In offline rendering, importance sampling is the dominant technique for reducing sampling noise: the closer the sampling distribution matches the distribution of light, the less noise there is [Vea97]. With the advent of hardware ray tracing, it is now important to adapt established techniques to the real-time setting [KVB*19].

Like with caching techniques, spatio-temporal reuse is key for high-quality importance sampling, as evidenced by bidirectional techniques [LW93; VG94; VG97; Kel97; GKDS12], resampling and mutation strategies [TCE05; VG97; KSAC02; KMA*15; HKD14], as well as learned distributions via neural networks [MMR*19; ZZ19; KD19] or path guiding [Jen95; LW95; VKŠ*14; MGN17; DK18]. Unfortunately, most reuse strategies come with a significant performance penalty, meaning that while they handle difficult illumination well, they are outperformed by naïve unidirectional path tracing under simple illumination. This disqualifies them from the real-time setting that we strive for. Path guiding has generally little overhead, which makes it a good choice for production rendering [VHH*19], however adapting its underlying data structures to animated content in real-time is non-trivial ongoing work [DHD20].

Our importance sampling method of choice is ReSTIR [BWP*20]: a recent combination of importance *resampling* [TCE05] and classic weighted reservoir sampling [Vit85; ES06; CHA82] that permits the reuse of a massive number of samples in constant time. For estimating *direct* lighting, ReSTIR generates many candidate samples on the light sources in the scene, resamples those candidates across space and time proportional to their predicted contribution, and then traces a shadow ray to determine the visibility of the selected sample(s). This scheme not only has negligible overhead compared to simple path tracing, but also gracefully handles dynamic content.

In contrast, we apply ReSTIR to *global* illumination by targeting it at the sum of direct illumination from light sources and indirect illumination stored in DDGI probes [MGNM19] of irradiance (for diffuse reflection) and angularly-filtered radiance [MMSM21] (for rough glossy reflection). This converts the global illumination problem into a purely direct illumination problem for the high variance case of low-frequency angular scattering, by treating DDGI as a light source. In contrast to multiple importance sampling [VG95], this means that we draw a single sample approximately proportional to the total transported light rather than multiple samples that are heuristically combined. Note that the result has two symbiotic levels of reuse: (i) screen-space resampling by ReSTIR for both direct and indirect light as well as (ii) world space DDGI probes which contain spatio-temporal aggregates of global illumination.

Concurrent work of Ouyang et al. [OLK*21b] proposes an alternative mechanism for spatio-temporal resampling of global illumination: they target ReSTIR at single-sample Monte Carlo estimates – rather than a cache – which can be likened to importance sampling of virtual point lights. Compared with our approach, theirs is capable of less biased (optionally unbiased) rendering, but lacks world-space spatio-temporal reuse and early path termination, which in our case is handled by the DDGI volume.

As we show in our results, combining the underlying ideas of ReSTIR, DDGI, and path tracing into a new sampling strategy produces a less-biased and less noisy result than naïvely compositing the results of separately-computed ReSTIR direct, DDGI diffuse/rough-glossy indirect, and path traced near-specular contributions.

3. Algorithm

3.1. Notation and Original DDGI Algorithm

In order to render images with global illumination in real-time, we strive to efficiently simulate radiance transport. The rendering equation [Kaj86]

$$L = L_e + T_f L \quad (1)$$

describes the *outgoing* radiance $L: \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathbb{R}^3$ as the sum of the radiance L_e emitted by the light sources and the transported radiance $T_f L$, where T_f is a higher order operator that maps between two functions (similar to notation defined by Veach [Vea97]). The symbol f represents the bidirectional scattering distribution function (BSDF), describing how the surfaces in the scene transport radiance. For our derivation, it is useful to partition radiance transport

$$T_f = T_{f_d} + T_{f_g}$$

into diffuse T_{f_d} and glossy T_{f_g} transport, the latter of which including specular transport as the special case of small surface roughness.

With this operator notation at hand, the dynamic diffuse global illumination (DDGI) algorithm [MGNM19] is

$$L \approx L_e + T_{f_d} L_e + L_{\text{DDGI}} + T_{f_g} L, \quad (2)$$

in which $L_{\text{DDGI}} \approx \sum_{i=2}^{\infty} (T_{f_d})^i L_e$ is an approximation of the diffuse transported radiance without the directly visible light sources (i starts at 2). The other terms are estimated by sampling. Inserting the DDGI approximation into equation (2) and assuming equality results almost in the rendering equation (1), missing only higher order glossy transport.

Majercik et al. [MMSM21] extended DDGI to include this missing contribution for the special case of specular reflection, which we will build on in Section 3.4. The “Primary DDGI” column in Fig. 3 and 8 shows results of this extended DDGI algorithm. As compared to path tracing [Kaj86], the images are smooth and noise only results from sampling the direct diffuse illumination $T_{f_d} L_e$ and the recursion for glossy transport $T_{f_g} L$. Yet, the irradiance probe approximation of L_{DDGI} shows visible artifacts that we resolve by postponing its evaluation to secondary path vertices. We develop our new algorithm in the following sections.

3.2. DDGI Resampling Overview

As the DDGI approximation can be queried at any point of the scene surface, we consider the whole scene surface as a light source rather than only the actual light sources. Our new sampling algorithm performs the following steps:

1. Generate candidate samples on *all* surfaces of the scene. To this end, we uniformly sample a small number of positions on emissive surfaces and generate additional samples by tracing secondary rays in directions sampled from the BSDF.




Figure 2: We compute direct and indirect illumination in a unified way by targeting the ReSTIR algorithm at the sum of emitted light L_e and the DDGI approximation L_{DDGI} . Left: two points are sampled. One (red) is sampled on the emissive surfaces and another (green, y) by BSDF sampling at the primary vertex x . Right: with all samples across pixels and past frames converted to the area measure, a single sample z is selected with probability proportional to $L_e + L_{\text{DDGI}}$ and a shadow ray (red) is traced. The radiance leaving x towards the camera is estimated with low variance by combining the reflected light from y and z through the method of Kollig and Keller [KK06]. (Near-)specular transport at y and z is estimated through recursion (blue) and the remaining transport by L_{DDGI} .

2. Select one candidate by first sampling from a weighted reservoir proportional to the sample contribution, then resample spatially among neighboring pixels, and resample temporally among previous frames, proportional to the sum $L_e + L_{\text{DDGI}}$.
3. Trace one shadow ray to the selected sample point.
4. If visible, shade from that sample point, using the emissive contribution L_e and adding the contribution from the DDGI volume.

As with original ReSTIR, our unified sampling scheme traces only a single ray to the selected sample (see Figure 2).

However, as the original DDGI volume only contains the diffuse *indirectly reflected* light without direct reflections as shown in equation (2), we need to add these, as well as glossy contributions.

The following sections discuss augmentations to the DDGI volume, as well as optimizations to the algorithm as a whole to reduce variance and improve performance. In Section 3.3 and Section 3.4, we discuss adding secondary direct illumination and glossy illumination to the DDGI probe volume. Section 3.5 discusses changes to the volume query to reduce bias. The last two subsections describe optimizations to minimize the variance caused by the weak singularity in the radiance integral (Section 3.6) and to reduce the cost of the resampling step (Section 3.7).

3.3. Augmenting the DDGI Approximation

As noted in Section 3.2, the DDGI approximation includes only diffuse indirect radiance. Integrating L_{DDGI} from secondary vertices over the hemisphere at the primary vertex would result in too dark images, because the DDGI approximation lacks the diffuse direct illumination and the radiance from non-diffuse transport. In fact, adding these missing terms to the DDGI approximation recovers an approximation of the transport operator

$$L_{\text{DDGI}} + T_{f_d} L_e + T_{f_g} L \approx T_{f_d} L + T_{f_g} L = T_f L. \quad (3)$$




Figure 3: Using the DDGI approximation L_{DDGI} at the primary vertex (left) shows light leaking artifacts beneath the couch and does not correctly reproduce indirect contact shadows on the table, vases, and plate. Using our augmented $L_{\text{DDGI}+}$ at the *secondary* vertex (middle) exhibits no leaks, recovers some of the indirect contact shadows and is closer to the path traced reference (right).

Adding the missing terms by recursive path tracing would increase the noise and involve further shading and tracing cost.

Instead, we augment the DDGI approximation in two steps, which are most easily explained by examining the underlying irradiance approximation that DDGI makes. Given a 3D position and 2D normal vector, the DDGI probe volume stores irradiance from indirect diffuse reflections, which is approximated by repeatedly evaluating the following double-bounce transport [MGNM19]

$$E_{\text{DDGI}}(\mathbf{x}, \mathbf{n}) \approx \int (L_{\text{DDGI}} + T_{f_d} L_e)(h(\mathbf{x}, \omega_i), -\omega_i) \langle \mathbf{n}, \omega_i \rangle d\omega_i, \\ L_{\text{DDGI}}(\mathbf{x}, \omega_i) := \frac{\rho(\mathbf{x})}{\pi} E_{\text{DDGI}}(\mathbf{x}, \mathbf{n}), \quad (4)$$

where h is the ray tracing operation and ρ is the diffuse albedo. Repeated Monte Carlo estimation of this equation at the centers \mathbf{x} of all DDGI probes converges to the aforementioned approximation

$$L_{\text{DDGI}} \approx \sum_{i=2}^{\infty} (T_{f_d})^i L_e.$$

To include direct reflections and non-diffuse transport, we replace the double-bounce transport with simpler single-bounce recursion

$$E_{\text{DDGI}+}(\mathbf{x}, \mathbf{n}) \approx \int (L_{\text{DDGI}+} + L_e)(h(\mathbf{x}, \omega_i), -\omega_i) \langle \mathbf{n}, \omega_i \rangle d\omega_i, \quad (5)$$

where L_e no longer undergoes a diffuse interaction and $L_{\text{DDGI}+}$ will be adapted to include glossy transport in the following section. The approximation thus converges to the full transport operator

$$L_{\text{DDGI}+} \approx T_f L.$$





Figure 4: For shading, we recursively trace (near-)specular interactions until their footprint grows sufficiently large to blur away small-scale artifacts of the DDGI approximation (middle). This produces less biased results than always querying the DDGI approximation at the secondary vertex (left).


Including direct illumination (i.e. the source term L_e) in the probe update equation (5) produces results that are too bright whenever probe centers are closer to light sources than the surfaces that they shade. This causes surfaces to receive more (secondary) direct illumination than they should – a common problem in probe-based caches. To alleviate this, we approximate the true distance to the surface by querying the average visibility already computed by the probes in the *backwards* direction [MGNM19] (dashed red lines), clipping that value to the probe boundary (pink), and adding it to the length of the corresponding probe update rays (solid red arrows). Because the update rays are traced per probe, this costs just a texture read without any probe weight computation. We then use the inverse square of the *total* distance, i.e., solid + dashed, to attenuate L_e . In contrast to direct illumination computed at probe centers, this heuristic slightly underestimates direct illumination. We therefore expose a user-tunable slider to scale the additional attenuation per scene to correct for brightness differences that would otherwise be a significant source of bias.



3.4. Including Glossy Illumination

Recall that the reflected light from a diffuse material with albedo ρ can be cheaply approximated by looking up $\frac{\rho(\mathbf{x})}{\pi} E_{\text{DDGI}+}(\mathbf{x}, \mathbf{n})$.

Majercik et al. [MMSM21] make the observation that the irradiance approximation can *additionally* be used to very coarsely approximate specular transport by substituting the normal \mathbf{n} with the direction of specular reflection ω_r . The irradiance $E_{\text{DDGI}}(\mathbf{x}, \omega_r)$ can then be *re-interpreted* as a prefiltered (by the cosine term acting as the kernel) coarse approximation of incident radiance. To obtain units of (prefiltered) radiance, one must divide $E_{\text{DDGI}}(\mathbf{x}, \omega_r)$ by 2π – the mass of the cosine under the integration measure. In the illustration, the green lobes visualize cosine-weighted prefiltering of incident radiance and the red arrows indicate the query direction of E_{DDGI} .



While Majercik et al. use this strategy for just specular transport, we generalize it to arbitrary glossy BSDFs – in our case microfacet models – by Monte Carlo estimation. More specifically, we approximate reflected radiance by our augmented DDGI model

$$L_{\text{DDGI}+}(\mathbf{x}, \omega_o) = \underbrace{\frac{\rho(\mathbf{x})}{\pi} E_{\text{DDGI}+}(\mathbf{x}, \mathbf{n})}_{\text{Diffuse}} + \underbrace{\frac{1}{2\pi} \int E_{\text{DDGI}+}(\mathbf{x}, \omega_i) f_g(\mathbf{x}, \omega_o, \omega_i) \langle \mathbf{n}, \omega_i \rangle d\omega_i}_{\text{Glossy}} \quad (6)$$

as the decomposition into diffuse and glossy transport, where we estimate the glossy integral with a single Monte Carlo sample drawn proportionally to the glossy BSDF component f_g .

Specular recursion for shading. We use our radiance approximation $L_{\text{DDGI}+}$ in three parts of the algorithm: (i) in the update rule of $E_{\text{DDGI}+}$ (5), (ii) to resample proportional to $(L_e + L_{\text{DDGI}+})$, and (iii) for shading. The latter use case – shading – requires special care to avoid exposing too much of the inherent bias of $L_{\text{DDGI}+}$ through (near-)specular interactions.

We thus trace shading paths recursively until their scattering footprint is sufficiently spread out to blur DDGI’s visual artifacts. To this end, we adopt the path termination strategy of Müller et al. [MRNK21; MRKN20], which is based on the isotropic path footprint approximation of Bekaert et al. [BSC*03].

More specifically, we approximate the footprint of a path with vertices $\mathbf{x}_1 \cdots \mathbf{x}_n$ as

$$a(\mathbf{x}_1 \cdots \mathbf{x}_n) = \left(\sum_{i=1}^{n-1} \sqrt{\frac{\|\mathbf{x}_i - \mathbf{x}_{i+1}\|^2}{p_i(\omega_{i,i} | \mathbf{x}_i, \omega_{o,i}) \langle \mathbf{n}_{i+1}, -\omega_{i,i} \rangle}} \right)^2, \quad (7)$$

where p_i is the BSDF sampling density at the i -th vertex. At the primary vertex, p_1 is proportional to the entire BSDF f , whereas $p_i; i > 1$ is proportional to only the glossy portion of the BSDF f_g , because we simply terminate the recursion early and query $\frac{\rho(\mathbf{x}_n)}{\pi} E_{\text{DDGI}+}(\mathbf{x}_n, \mathbf{n}_n)$ when the diffuse portion is sampled. When the

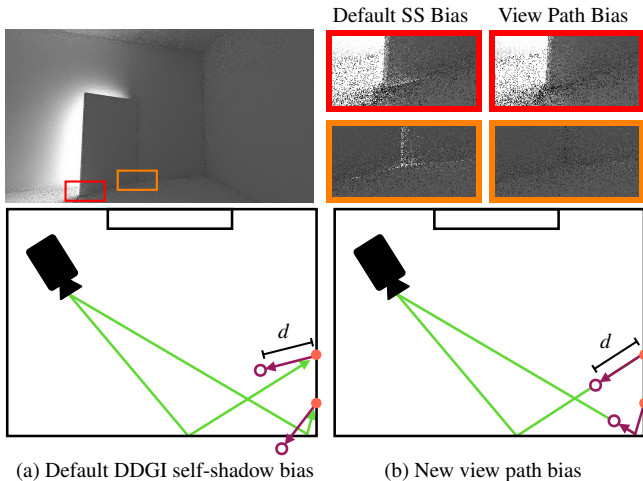


Figure 5: DDGI combats self-shadowing by pushing the evaluation point away from surfaces (violet arrow) by a fixed distance d that is proportional to the probe size. Left: DDGI uses a linear combination of the view direction and the surface normal. This may push the evaluation point through geometry, leading to light leaking. Right: we instead move the evaluation point a fixed distance back along the path, which is guaranteed to not penetrate surfaces.

glossy portion is sampled, we terminate the recursion by estimating the glossy integral of Eq. (6) as soon as $a(\mathbf{x}_1 \cdots \mathbf{x}_n) > a_0 \cdot c$, where

$$a_0 := \frac{4\pi \|\mathbf{x}_0 - \mathbf{x}_1\|^2}{\langle \mathbf{n}_1, -\omega_{i,j} \rangle} \quad (8)$$

is an approximation (up to constant factors) of the surface area viewed by the camera along direction $\omega_{i,j}$. \mathbf{x}_0 is the camera position, and c is a user-chosen threshold. See Bekaert et al. [BSC*03] and Müller et al. [MRNK21; MRKN20] for details.

We empirically set the threshold $c = 0.2$ to obtain a satisfactory noise versus bias trade-off. Fig. 4 confirms that this helps reduce artifacts on highly glossy surfaces without affecting rough glossy and diffuse surfaces.

3.5. Correcting Self-shadow Bias

The original DDGI algorithm applies an offset to the positions at which L_{DDGI} is queried to avoid computing visibility weights exactly at a surface where their variance is high. This offset is called the “self-shadow” bias, and is computed as a weighted combination of the view and normal vector at the queried point.

While the self-shadow bias was sufficient for the original primary-vertex DDGI algorithm, using it unmodified in our indirect sampling strategy showed light leaking artifacts (Figure 5). These artifacts arise from very short distances between the view point and the query point, as might arise from sharp corners. To address these artifacts, we bias the query point along the *view path* as opposed to the surface normal and view direction. Results using our view path bias against the original self-shadow bias are shown in Figure 5.

3.6. Taming the Weak Singularity in the Geometry Term

In order to perform spatio-temporal resampling among different shading locations, the geometry term is introduced as part of the change of the integration measure. Unfortunately, the weak singularity in the geometry term may lead to unbounded variance in the resampling procedure in geometric creases; see Fig. 6 (middle).

In typical path tracing applications, uniform multiple importance sampling [VG95] of light sources would bound the variance when light sources (almost) touch the surfaces that they shade. But this is an impractical solution in our case, where the entire scene is considered a light source. It may be feasible to *non-uniformly* sample all relevant scene surfaces according to a data structure that is constructed online, but we instead take a simpler approach that does not require multiple importance sampling: we partition the integrand to bound the variance [KK06]. This strategy of partitioning the integrand is generally applicable to other caching techniques, as it deals with the issue of resampling points close to the primary vertex which is a cache independent decision.

To avoid the weak singularity caused by samples very close to the primary vertex, we clamp the computed geometry term to a maximum constant value. Using a constant value of $G_{\text{max}} = 1$ ensures that the differential under solid angle measure ($d\omega$) is always less than or equal to the differential under area measure (dA), and works well for our test scenes. As scene scale increases, a constant G_{max} will bound a smaller portion of the integrand, meaning that keeping G_{max} constant is not scale invariant. To match scene scale, G_{max} may be scaled relative to surface area as viewed by the camera (in other words, proportional to a_0 in Eq. 8).

We use the notation $\lceil T_f \rceil$ for the transport with the bounded geometry term. The residual transport is denoted by $\lfloor T_f \rfloor$, such that

$$T_f = \lceil T_f \rceil + \lfloor T_f \rfloor.$$

$\lfloor T_f \rfloor$ is evaluated by an auxiliary BSDF sample in the sampling step described in Sec. 3.2. The improved variance is shown in Fig. 6.

In summary, our new algorithm to approximate the rendering equation amounts to summing separate estimators of each term of


$$L \approx L_e + \underbrace{\lceil T_f \rceil (L_e + L_{\text{DDGI}+})}_{\text{bounded geometry term}} + \underbrace{\lfloor T_f \rfloor (L_e + L_{\text{DDGI}+})}_{\text{residual transport}}, \quad (9)$$

where

- L_e is evaluated at the primary vertex,
- $\lfloor T_f \rfloor$ by importance sampling the BSDF, and
- $\lceil T_f \rceil$ by spatio-temporal reservoir resampling.

3.7. Improved Performance by Caching DDGI Queries

During the spatio-temporal resampling step, our algorithm requires numerous DDGI queries: one for each candidate’s contribution to the shaded pixel, which, in our implementation, amounts to 1 temporal and 3 spatial candidates per pixel. In the illustration, we depict the simpler case of resampling among just 2 spatial candidates y_2 and y_3 for shading at the center






Figure 6: Even though spatio-temporal resampling of DDGI (middle) improves the noise of path tracing (left), outliers (“fireflies”) cause distracting artifacts in geometric concavities. These fireflies are due to the weak singularity in the geometry term of the rendering equation. Using the method of Kollig and Keller [KK06], we bound the geometry term during resampling and then trace an auxiliary ray to estimate the residual transport (right). This eliminates fireflies and preserves a significant noise reduction. The images were rendered at 128 samples per pixel.




Figure 7: Caching the DDGI query for resampling is more efficient: while the increase in noise level and MAPE is negligible in query caching, the reduction in frame rendering time is substantial.

vertex \mathbf{x}_1 . For each candidate, we must query its contribution $L_{\text{DDGI}+}(\mathbf{y}_i, \mathbf{y}_i \rightarrow \mathbf{x}_1)$ to the center vertex, where $\mathbf{y} \rightarrow \mathbf{x}$ is the direction vector pointing from \mathbf{y} towards \mathbf{x} (dashed arrows).

These $\#\text{vertices} \times \#\text{candidates}$ queries can be avoided by approximating them with the $\#\text{vertices}$ values $L_{\text{DDGI}+}(\mathbf{y}_i, \mathbf{y}_i \rightarrow \mathbf{x}_i)$, which have to be evaluated for shading each vertex \mathbf{x}_i in any case. This saves a factor of $\#\text{candidates}$ queries. Importantly, using approximate values in the resampling step does not introduce additional bias in the rendered image, because we include BSDF samples in our initial sample generation, thus guaranteeing a positive probability density for all directions that reflect radiance. This results in additional noise, but no additional bias relative to using the DDGI+ approximation without resampling.

Nonetheless, we analyze the error of this approximation by considering its three sources:

1. the scene and the DDGI volume may change each frame,
2. the shading location \mathbf{x} is different, and
3. the view-path bias (Sec. 3.5) depends on the path prefix

each of which become small if we apply the assumptions of our method: local pixel neighborhoods, near-diffuse scattering, and short frame times. The approximation error is thus reasonably small in practice (see Fig. 7) and we use this optimization in all our results.

4. Results

We implemented DDGI resampling in Direct3D 12 using the Falcor rendering framework [BYC*20]. All results were generated on a high-end desktop machine with an Intel i7-6800K CPU and an NVIDIA RTX 3090 GPU.

Our configuration of DDGI follows Majercik et al. [MGNM19; MMSM21]: irradiance is stored per-probe in a 8×8 octahedral representation, updated per frame with a spherical raycast from probe centers (256 rays/probe). Each probe also stores a filtered average of distance and distance squared (16×16 octahedral) which is used to compute a rough estimate of probe visibility using the Chebyshev statistical test. Probe contributions to an irradiance query at any point and direction within the volume are weighted by combining probe visibility weights, backface weighting, and trilinear filtering. Our scenes used a grid of $16 \times 8 \times 16$ probes spread evenly over the scene bounding box.

Fig. 8 summarizes our results, comparing visuals, render time, and mean absolute percentage error (MAPE) of 5 different rendering techniques at 1 sample per pixel. MAPE is defined as $\frac{1}{N} \sum_{i=1}^N |v_i - \hat{v}_i| / (\hat{v}_i + \epsilon)$, where \hat{v}_i is the value of the i -th pixel in the reference image, v_i is the value of the i -th rendered pixel, and $\epsilon = 0.01$ prevents near-black pixels from dominating the metric.

All rendering techniques are based on an unbiased path tracer that uses ReSTIR [BWP*20] to estimate direct illumination at the primary vertex.

1. **PT + ReSTIR:** The aforementioned baseline path tracer.
2. **Primary DDGI:** The original DDGI algorithm [MGNM19; MMSM21], which approximates *indirect* illumination at the *primary* vertex of the path.

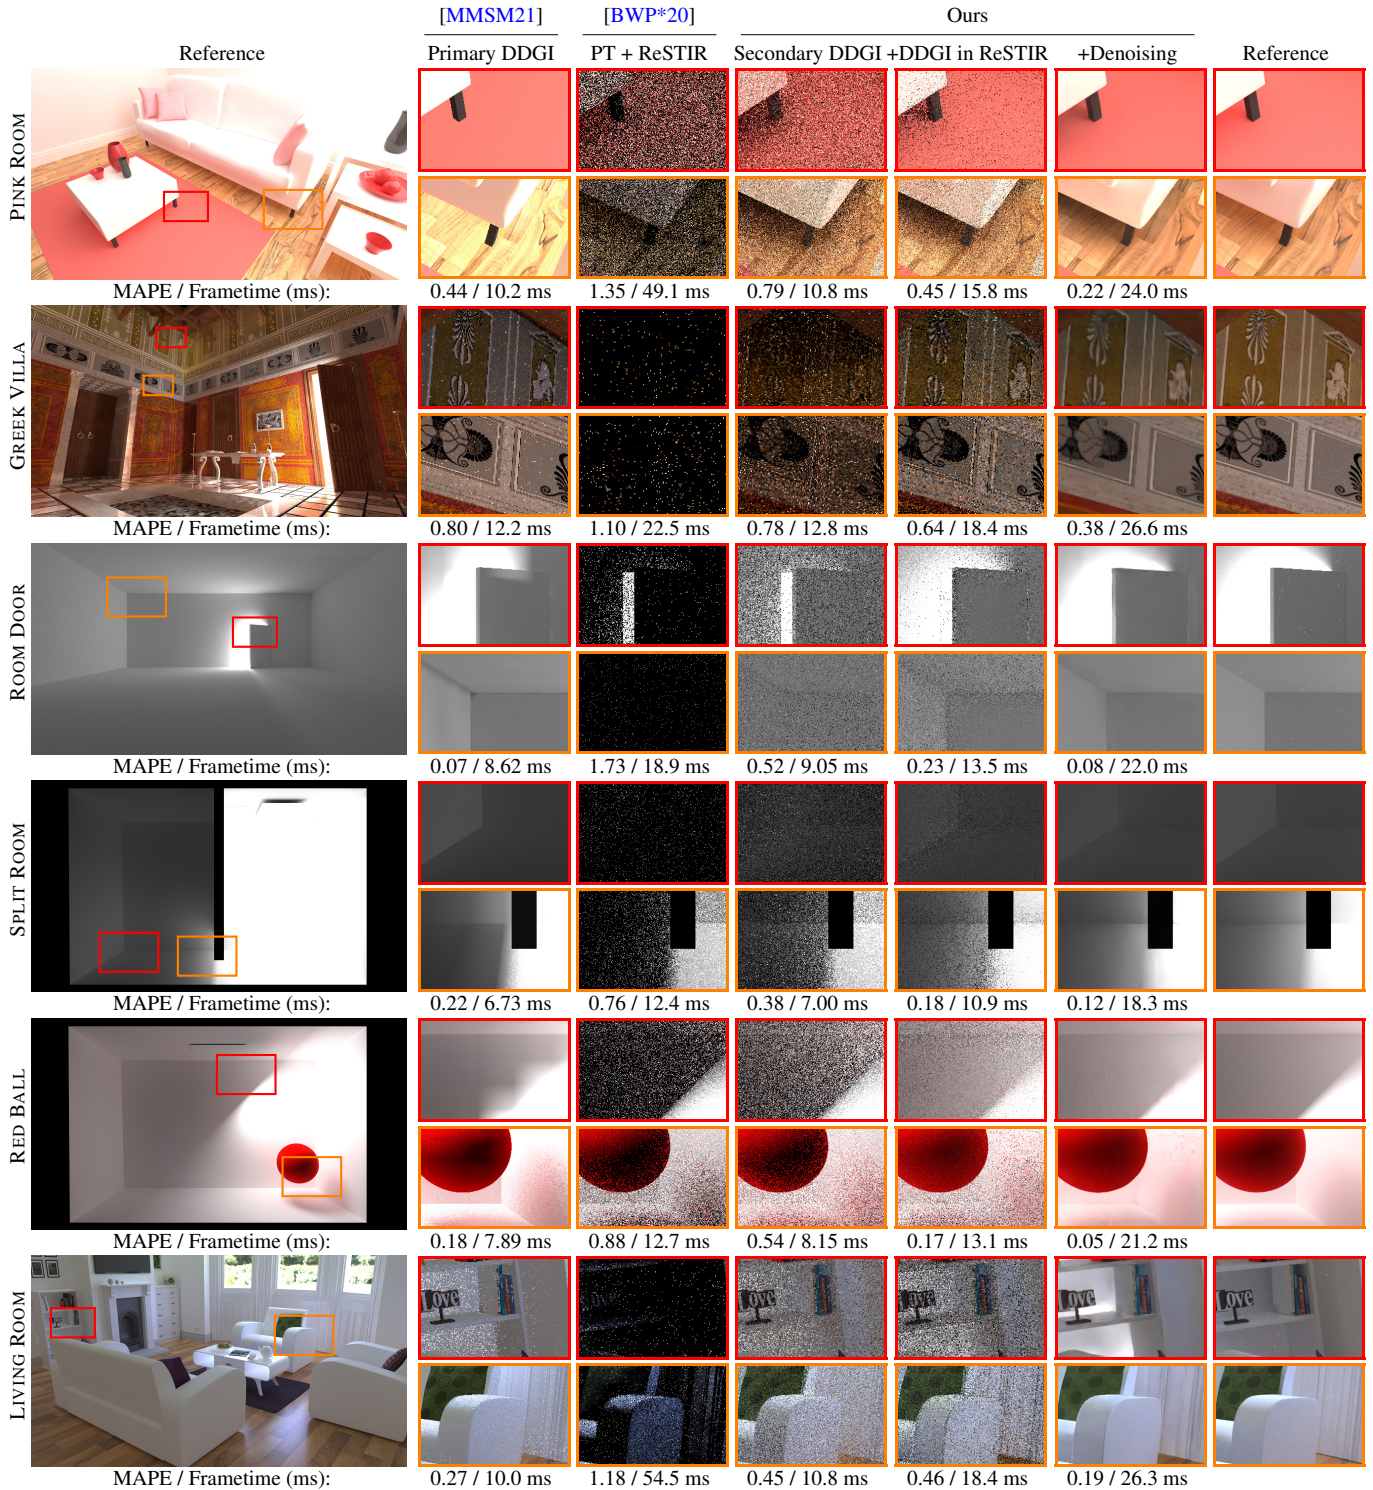


Figure 8: Comparison of rendering algorithms at 1 sample per pixel with respect to noise and bias. All algorithms are based on an unbiased path tracer that uses ReSTIR [BWP*20] to estimate direct illumination at the primary vertex. As compared to this baseline, dynamic diffuse global illumination (Primary DDGI) renders smooth indirect illumination. Bias artifacts inherent with the DDGI approximation can be ameliorated by pushing the query of DDGI one bounce further into the light path (Secondary DDGI). Facilitating the ReSTIR sampling mechanism to jointly importance sample these DDGI queries with direct illumination, the induced noise is reduced (+DDGI in ReSTIR), and is further amenable to commodity denoising methods (+Denoising). We also show a failure case (LIVING ROOM) where the outer walls consist of flat planes that receive strong sunlight from the outside. In this difficult scene, our method is unable to hide DDGI light leaking artifacts (red inset) and BSDF importance sampling (Secondary DDGI) performs slightly better than ReSTIR sampling.

Table 1: Time to converge to equal MAPE

Scene	Method	Frames	Time	MAPE
GREEK VILLA	PT+ReSTIR	129	2908.00 ms	0.573
	+Secondary DDGI	8	113.52 ms	0.578
	+DDGI Resampling	1	19.83 ms	0.572
LIVING ROOM	PT+ReSTIR	315	17157.04 ms	0.366
	+Secondary DDGI	1	10.82 ms	0.366
	+DDGI Resampling	3	56.57 ms	0.365
PINK ROOM	PT+ReSTIR	29	1423.70 ms	0.384
	+Secondary DDGI	9	97.21 ms	0.390
	+DDGI Resampling	1	16.85 ms	0.383
RED BALL	PT+ReSTIR	109	1381.71 ms	0.140
	+Secondary DDGI	29	254.17 ms	0.142
	+DDGI Resampling	1	13.37 ms	0.140
ROOM DOOR	PT+ReSTIR	4095	77278.05 ms	0.206
	+Secondary DDGI	79	748.51 ms	0.200
	+DDGI Resampling	1	13.28 ms	0.199
SPLIT ROOM	PT+ReSTIR	150	1854.92 ms	0.162
	+Secondary DDGI	27	202.42 ms	0.162
	+DDGI Resampling	1	10.92 ms	0.161
Average	PT+ReSTIR	804.50	17000.57 ms	0.305
	+Secondary DDGI	25.50	237.77 ms	0.306
	+DDGI Resampling	1.33	21.81 ms	0.303

Table 2: Rendering cost by component

Scene	Method	Trace & shade	Resampling	Update
GREEK VILLA	PT+ReSTIR	20.38 ms	1.10 ms	—
	+Secondary DDGI	10.65 ms	1.20 ms	1.54 ms
	+DDGI Resampling	14.03 ms	3.13 ms	1.55 ms
LIVING ROOM	PT+ReSTIR	52.72 ms	1.42 ms	—
	+Secondary DDGI	8.18 ms	1.62 ms	0.57 ms
	+DDGI Resampling	13.32 ms	4.47 ms	0.57 ms
PINK ROOM	PT+ReSTIR	46.60 ms	1.08 ms	—
	+Secondary DDGI	8.08 ms	1.26 ms	1.13 ms
	+DDGI Resampling	11.99 ms	2.41 ms	1.14 ms
RED BALL	PT+ReSTIR	10.71 ms	0.95 ms	—
	+Secondary DDGI	5.47 ms	1.27 ms	1.02 ms
	+DDGI Resampling	8.57 ms	3.23 ms	1.01 ms
ROOM DOOR	PT+ReSTIR	16.91 ms	1.10 ms	—
	+Secondary DDGI	6.71 ms	1.50 ms	1.07 ms
	+DDGI Resampling	9.64 ms	2.33 ms	1.08 ms
SPLIT ROOM	PT+ReSTIR	10.90 ms	0.88 ms	—
	+Secondary DDGI	5.33 ms	1.24 ms	0.44 ms
	+DDGI Resampling	7.63 ms	2.65 ms	0.44 ms
Average	PT+ReSTIR	26.37 ms	1.09 ms	—
	+Secondary DDGI	7.40 ms	1.35 ms	0.96 ms
	+DDGI Resampling	10.86 ms	3.04 ms	0.97 ms

- 3. Secondary DDGI:** Our augmented DDGI algorithm approximates *global* illumination at the *secondary* vertex of the path (Sec. 3.3–3.5). Near-specular transport is traced recursively.
- 4. + DDGI in ReSTIR:** The same as before, but the DDGI approximation is included in the ReSTIR algorithm (Sec. 3.2). The diffuse weak singularity is separately sampled (Sec. 3.6).
- 5. + Denoising:** The same as before, but with OptiX denoising.

As expected, the DDGI algorithm [MGNM19; MMSM21] has the least noise with the greatest performance (“Primary DDGI” column). However, it also exhibits significant visual artifacts such as flat shading (e.g., in the GREEK VILLA), a lack of ambient occlusion (e.g., in the PINK ROOM), and light leaking (e.g., in the LIVING ROOM). At the other extreme, an unbiased path tracer produces highly accurate results but requires a long time for its noise to converge away in globally illuminated scenes, even when using ReSTIR [BWP*20] for direct lighting (“PT + ReSTIR” column).

Our proposed *indirect* use of DDGI probes – provided that the probes contain the energy of the full transport (Sec. 3.3 and 3.4) – reduces noise while introducing only little bias, because the inaccuracies of the DDGI probes are hidden behind the first scattering event (“Secondary DDGI” column). Additionally, the performance is improved, because paths are terminated early into DDGI probes. Incorporating the indirect lighting injected by the DDGI probes in the resampling step further reduces noise (“DDGI in ReSTIR” column). Overall, the noise is now low enough that adding the OptiX denoiser results in more visually pleasing images than the original DDGI algorithm, albeit at an on average 1.5–2× greater cost.

We also show a failure case (LIVING ROOM) where the outer walls consist of flat planes that receive strong sunlight from the outside. In this difficult scene, secondary DDGI queries are unable to hide DDGI’s light leaking artifacts (red inset) and BSDF importance sampling (“Secondary DDGI” column) performs slightly better than ReSTIR sampling. This shows that further work is needed in making DDGI more resilient to thin geometry as well as in understanding

the circumstances in which ReSTIR can be outperformed by its candidate generation strategies.

Quantitative evaluation. To gauge the convergence improvements offered by using DDGI at the secondary path vertex (with or without resampling), we list the time it takes for all noisy methods to reach equal error in Tab. 1. In all scenes, except for the LIVING ROOM, DDGI resampling converges the fastest – in some scenes by a significant margin – even when taking into account the slightly increased cost of rendering.

To better understand this cost, we break it down by component in Tab. 2. As expected, terminating paths into DDGI reduces the ray tracing cost because paths are shorter, even when taking into account the additional cost of querying DDGI during the shading step, as well as the cost of updating the DDGI volume every frame. However, perhaps unexpectedly, resampling only becomes marginally more expensive by incorporating DDGI in it (by 1.7 ms on average) thanks to the caching of previous DDGI queries (Sec. 3.7). The remaining overhead of DDGI resampling arises in tracing & shading of (near-)specular path suffixes, which must now be performed twice per pixel: once for the BSDF sample that estimates $[T_f]$ and once for the resampled vertex that estimates $[T_f]$ (see Sec. 3.6).

5. Discussion and Future Work

The following discussion focuses on future avenues for potential performance and noise improvements.

Subsampling. To reduce the remaining resampling cost, it might be feasible to subsample the generation of new candidates: for example, only every n -th pixel may generate novel samples from the BSDF and from light sources. All other pixels would merely resample among spatio-temporal neighbors. Since this strategy is rather orthogonal to our work – it would benefit all ReSTIR-type algorithms – we did not include it in our technique. However, we feel that it warrants a detailed analysis in the future.




Figure 9: Analysis of DDGI Resampling (left split) augmented by DDGI used as a control variate at the primary vertex (right split). The transient artifacts visible at low sampling rates are due to the underlying probe volume as can be seen by the visualization of DDGI. At 1 spp, these are too salient to ignore, although with higher sample counts both DDGI Resampling and DDGI Resampling with control variate converge to the same result.

Alternative Caching Strategies. While we focused on the specific use of DDGI, we believe other caching strategies, such as the recent neural radiance caching [MRNK21], can similarly benefit from importance sampling through ReSTIR. Another example is the concurrent work of Ouyang et al. [OLK*21b] who apply ReSTIR to a cache that is effectively a collection of spatio-directional point light sources. Their approach represents a different point on the Pareto front that trades off noise, bias, and performance – one that has lower bias (and can even be unbiased) at the cost of tracing longer paths, leading to extra noise.

5.1. Alternative Variance Reduction Strategies

In real-time rendering, sampling with low variance is key due to the limited budget of rays to be traced. In this context, we also analyzed the efficiency of utilizing DDGI as a control variate, as well as that of using it for path guiding instead of resampling. Both approaches were overall less efficient than the presented DDGI resampling, but we will nonetheless describe them in the hope that future algorithms may build on top of them.

DDGI as Control Variate. For scenes with simple occlusion, the original DDGI data structure provides a low-bias, noise-free result. This can sometimes be desired over our new resampling, which will introduce noise even in cases where the resulting lower bias is not visually noticeable. To combat this, we explored extending our solution by using primary hitpoint irradiance $E_{\text{DDGI}+}$ as a control variate of diffuse transport. That is, we assumed that $\frac{\rho}{\pi} E_{\text{DDGI}+}$ at

primary vertices yields the correct indirect illumination and we re-sampled proportional to the remaining non-diffuse transport *plus* (neglecting constants for brevity) the absolute difference between $(L_{\text{DDGI}+} + L_e)$ at the secondary vertex and the spatio-directionally nearest DDGI update ray among those that were traced while rendering the last frame. This amounts to using the DDGI update rays as a spatio-directionally piecewise constant control variate and their average, $E_{\text{DDGI}+}$, as its exact integral.

When shading, we then applied DDGI to the primary vertex (as described in [MGNM19]) and added the *difference* (not the absolute value) between $L_e + L_{\text{DDGI}+}$ at the secondary vertex and the nearest DDGI update ray. Adding this difference in the shading step may result in negative radiance values, which frequently manifest as negative fireflies that cause black pixels. We omitted spatio-temporal probe interpolation for the DDGI control variate, as otherwise the value of $L_{\text{DDGI}+}$ at the primary vertex would not be an exact integral of the control variate.

Fig. 9 shows our results. Though this algorithm converges to the same (biased) solution as plain DDGI resampling, at one sample per pixel, the control variate exposes light leaks similar to the naïve ReSTIR + DDGI combination. At the same time, the noise level does not improve upon plain DDGI resampling, which produces less prominent light leaks.

While using DDGI as a control variate may be beneficial in some settings, in general, it appears not sufficiently efficient and remains an avenue for future research.

DDGI for Path Guiding. Based on the DDGI data structure, guiding rays towards where the radiance comes from is straightforward. Given a position in space, we can query the closest DDGI probe and build the cumulative distribution function (CDF) from the results of its update rays to sample a guiding direction. While the results of the probe rays are stored, building the CDF and sampling from it was too slow to increase the overall efficiency of our algorithm.

In addition, using $L_{\text{DDGI}+}$ to normalize the PDF and sample the guiding direction in an unbiased way requires the same modifications to be made to the DDGI volume that were made to the control variate. These make the volume less stable and thus less effective for path guiding. An alternative would be porting efficient path guiding data structures [VHH*19] to the GPU.

Combining Primary and Secondary DDGI. Even though we focus on hiding the bias of DDGI behind scattering interactions, there are situations in which DDGI’s bias is very low in the first place and does not warrant the extra noise caused by secondary queries; see the orange inset of the RED BALL scene in Fig. 8. It likely pays off to design an automatic mechanism for detecting such situations and to select among primary and secondary DDGI queries accordingly.

Neural Denoising Recently, much work has been done to apply neural networks to denoising [MZV*20; FWHB21; HMS*20; MH20]. While our work specifically focuses on sample generation rather than denoising, many neural denoising approaches are directly compatible with our resampling method. Further, neural networks may be used to learn a scene or view dependent probability distribution for selecting initial samples, further improving the sample efficiency of our work in concert with denoising.

Path Space Filtering. Lastly, alternatively to screen-space denoising, we believe that additional spatio-temporal reuse through path space filtering may be beneficial. While ReSTIR amounts to spatio-temporal filtering of sampling probabilities and DDGI amounts to spatio-temporal filtering of radiance at *secondary* vertices, path space filtering would provide additional spatio-temporal filtering of radiance at *primary* vertices. Recent hash-based implementations [BFK22] can run with little overhead on modern GPUs.

5.2. Limitations

The limitations of our approach are largely inherited from its components: ReSTIR and DDGI. ReSTIR is a screen-space technique whose reuse capability depends highly on the availability of motion vectors, slow camera motion, and transparency. These downsides can be mitigated by transitioning to a world-space representation [JBW21], but another limitation remains: through resampling, samples become spatio-temporally correlated, which reduces the total amount of information available to modern denoisers. Effective denoising in the presence of ReSTIR samples is still ongoing work.

The limitations induced by DDGI are two-fold. First, being grid-based, DDGI does not scale easily to expansive scenes where the camera may only focus on small regions. And second, DDGI is incapable of resolving fine spatio-directional detail such as intricate caustics, shadows, ambient occlusion or near-specular reflections. These limitations are hidden to a degree when DDGI is queried after a primary scattering interaction that is sufficiently diffuse or after undergoing sufficiently many secondary interactions, but better quality can be achieved by using a more accurate cache [MRNK21].

6. Conclusion

We have combined DDGI [MMSM21] and ReSTIR [BWP*20], which, at first glance, might seem orthogonal. DDGI renders indirect illumination and ReSTIR direct illumination. However, the appearance is deceiving and a small modification to the DDGI algorithm – querying the probe volume at secondary path vertices rather than primary ones – unifies the two approaches into an algorithm that performs much better than its constituent parts. By querying DDGI at secondary vertices, it acts as a scene-spanning light source, effectively turning the global illumination problem into a much simpler direct illumination problem, which ReSTIR can holistically attack. This observation is a general one that has inspired rendering algorithms since the 90s, making us excited about future work in this direction.

Quantitatively, the combined algorithm is about 60% more expensive than either of the two prior works, but makes up for its cost in terms of low noise and bias. It achieves equal error over an order of magnitude more quickly. Together with denoising, real-time rendering of global illumination in dynamic scenes is feasible.

Acknowledgements

We would like to thank David Luebke and Aaron Lefohn for supporting this work. We also thank Benedikt Bitterli and Chris Wyman for providing the initial codebase in Falcor [BYC*20], Chris Wyman and Peter Shirley for many early ideas and fruitful discussion, and

Nikolaus Binder and Christoph Schied for proofreading the article and many valuable comments. The Pink Room scene was modified from “The Modern Living Room” by Wig42 and the Living Room scene was modified from “The White Room” by Jay-Artist. The Greek Villa was modified from “3D Greek Villa” by holy_diver.

References

- [Abr97] ABRASH, MICHAEL. “Quake’s lighting model: Surface caching”. *Graphics Programming Black Book*. Coriolis Group, 1997. Chap. 68, 1245–1256 2.
- [Arv86] ARVO, JAMES. “Backward Ray Tracing”. In *ACM SIGGRAPH ’86 Course Notes - Developments in Ray Tracing*. 1986, 259–263 2.
- [Bal18] “Special Issue on Production Rendering”. *ACM Trans. Graph.* 37.3 (2018). Ed. by BALA, KAVITA. ISSN: 0730-0301 2.
- [BB17] BOUGHIDA, MALIK and BOUBEKEUR, TAMY. “Bayesian Collaborative Denoising for Monte Carlo Rendering”. *Computer Graphics Forum* 36.4 (2017), 137–153. DOI: <https://doi.org/10.1111/cgf.13231>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13231>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13231> 2.
- [BFK22] BINDER, NIKOLAUS, FRICKE, SASCHA, and KELLER, ALEXANDER. “Massively Parallel Path Space Filtering”. *Monte Carlo and Quasi-Monte Carlo Methods 2020*. Ed. by KELLER, ALEXANDER. Springer, 2022, to appear 11.
- [BRM*16] BITTERLI, BENEDIKT, ROUSSELLE, FABRICE, MOON, BOCHANG, et al. “Nonlinearly Weighted First-Order Regression for Denoising Monte Carlo Renderings”. *Comput. Graph. Forum* 35.4 (July 2016), 107–117. ISSN: 0167-7055 2.
- [BSC*03] BEKAERT, PHILIPPE, SLUSALLEK, PHILIPP, COOLS, RONALD, et al. *A custom designed Density Estimation Method for Light Transport*. Tech. rep. Max-Planck-Institut für Informatik, Saarbrücken, Germany, Apr. 2003 5, 6.
- [BVM*17] BAKO, STEVE, VOGELS, THIJS, MCWILLIAMS, BRIAN, et al. “Kernel-Predicting Convolutional Networks for Denoising Monte Carlo Renderings”. *ACM Trans. Graph.* 36.4 (July 2017). ISSN: 0730-0301. DOI: [10.1145/3072959.3073708](https://doi.org/10.1145/3072959.3073708). URL: <https://doi.org/10.1145/3072959.3073708> 2.
- [BWP*20] BITTERLI, BENEDIKT, WYMAN, CHRIS, PHARR, MATT, et al. “Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting”. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39.4 (July 2020). DOI: [10.1145/3386569.3392481](https://doi.org/10.1145/3386569.3392481) 1–3, 7–9, 11.
- [BYC*20] BENTY, NIR, YAO, KAI-HWA, CLARBERG, PETRIK, et al. *The Falcor Rendering Framework*. <https://github.com/NVIDIAGameWorks/Falcor>. Aug. 2020. URL: <https://github.com/NVIDIAGameWorks/Falcor> 7, 11.
- [CHA82] CHAO, M. T. “A general purpose unequal probability sampling plan”. *Biometrika* 69.3 (Dec. 1982), 653–656. ISSN: 0006-3444. DOI: [10.1093/biomet/69.3.653](https://doi.org/10.1093/biomet/69.3.653). eprint: <https://academic.oup.com/biomet/article-pdf/69/3/653/591311/69-3-653.pdf>. URL: <https://doi.org/10.1093/biomet/69.3.653> 3.
- [CKS*17] CHAITANYA, CHAKRAVARTY R. ALLA, KAPLANYAN, ANTON S., SCHIED, CHRISTOPH, et al. “Interactive Reconstruction of Monte Carlo Image Sequences Using a Recurrent Denoising Autoencoder”. *ACM Trans. Graph.* 36.4 (July 2017). ISSN: 0730-0301. DOI: [10.1145/3072959.3073601](https://doi.org/10.1145/3072959.3073601). URL: <https://doi.org/10.1145/3072959.3073601> 2.
- [DHD20] DITTEBRANDT, ADDIS, HANIKA, JOHANNES, and DACHSBACHER, CARSTEN. “Temporal Sample Reuse for Next Event Estimation and Path Guiding for Real-Time Path Tracing”. *Eurographics Symposium on Rendering - DL-only Track*. Ed. by DACHSBACHER, CARSTEN and PHARR, MATT. The Eurographics Association, 2020. DOI: [10.2312/sr.20201135](https://doi.org/10.2312/sr.20201135) 3.

- [DK18] DAHM, KEN and KELLER, ALEXANDER. “Learning Light Transport the Reinforced Way”. *Monte Carlo and Quasi-Monte Carlo Methods*. Ed. by OWEN, ART B. and GLYNN, PETER W. Springer International Publishing, 2018, 181–195 3.
- [ES06] EFRAIMIDIS, PAVLOS S. and SPIRAKIS, PAUL G. “Weighted random sampling with a reservoir”. *Information Processing Letters* 97.5 (2006), 181–185. ISSN: 0020-0190. DOI: <https://doi.org/10.1016/j.ipl.2005.11.003>. URL: <https://www.sciencedirect.com/science/article/pii/S002001900500298X>.
- [FVHB21] FAN, HANGMING, WANG, RUI, HUO, YUCHI, and BAO, HUN. “Real-time Monte Carlo Denoising with Weight Sharing Kernel Prediction Network”. *Computer Graphics Forum* 40.4 (2021), 15–27. DOI: <https://doi.org/10.1111/cgf.14338>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14338>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14338>.
- [GKDS12] GEORGIEV, ILIYAN, KRÍVÁNEK, JAROSLAV, DAVIDOVIČ, TOMÁŠ, and SLUSALLEK, PHILIPP. “Light transport simulation with vertex connection and merging”. *ACM Trans. Graph.* 31.6 (Nov. 2012), 192:1–192:10. ISSN: 0730-0301. DOI: [10.1145/2366145.2366211](https://doi.org/10.1145/2366145.2366211). URL: <http://doi.acm.org/10.1145/2366145.2366211>.
- [GS12] GILBERT, NS MICKAEL and STEFANOV, NIKOLAY. “Deferred radiance transfer volumes, global illumination in Far Cry 3”. *Game Developers Conference*. 2012 2.
- [Hec90] HECKBERT, PAUL S. “Adaptive Radiosity Textures for Bidirectional Ray Tracing”. *SIGGRAPH Comput. Graph.* 24.4 (Sept. 1990), 145–154. ISSN: 0097-8930. DOI: [10.1145/97880.97895](https://doi.org/10.1145/97880.97895). URL: <https://doi.org/10.1145/97880.97895>.
- [HKD14] HACHISUKA, TOSHIYA, KAPLANYAN, ANTON S., and DACHSBACHER, CARSTEN. “Multiplexed Metropolis Light Transport”. *ACM Trans. Graph.* 33.4 (July 2014), 100:1–100:10. ISSN: 0730-0301. DOI: [10.1145/2601097.2601138](https://doi.org/10.1145/2601097.2601138). URL: <http://doi.acm.org/10.1145/2601097.2601138>.
- [HMS*20] HASSELGREN, JON, MUNKBERG, J., SALVI, MARCO, et al. “Neural Temporal Adaptive Sampling and Denoising”. *Computer Graphics Forum* 39 (May 2020), 147–155. DOI: [10.1111/cgf.13919](https://doi.org/10.1111/cgf.13919).
- [HY21] HUO, YUCHI and YOON, SUNG-EUI. “A survey on deep learning-based Monte Carlo denoising”. *Computational Visual Media* 7 (Mar. 2021). DOI: [10.1007/s41095-021-0209-9](https://doi.org/10.1007/s41095-021-0209-9).
- [JBW21] JUKARAINEN, PAULA, BOKSANSKY, JAKUB, and WYMAN, CHRIS. “Rendering of Many Lights with Grid-Based Reservoirs”. I3D ’21 Posters. Apr. 2021, 194–195 11.
- [Jen95] JENSEN, HENRIK WANN. “Importance Driven Path Tracing using the Photon Map”. Vienna: Springer Vienna, 1995, 326–335. ISBN: 978-3-7091-9430-0. DOI: [10.1007/978-3-7091-9430-0_31](https://doi.org/10.1007/978-3-7091-9430-0_31). URL: http://dx.doi.org/10.1007/978-3-7091-9430-0_31.
- [Jen96] JENSEN, HENRIK WANN. “Global Illumination Using Photon Maps”. *Proceedings of the Eurographics Workshop on Rendering Techniques ’96*. Porto, Portugal: Springer-Verlag, 1996, 21–30. ISBN: 3211828834 2.
- [Kaj86] KAJIYA, JAMES. “The Rendering Equation”. *Computer Graphics (Proc. SIGGRAPH 1986)*. 1986, 143–150 3.
- [KD19] KELLER, ALEXANDER and DAHM, KEN. “Integral Equations and Machine Learning”. *Mathematics and Computers in Simulation* 161 (2019), 2–12 3.
- [Kel97] KELLER, A. “Instant Radiosity”. *SIGGRAPH ’97: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. 1997, 49–56 3.
- [KGPB05] KRÍVÁNEK, JAROSLAV, GAUTRON, PASCAL, PATTANAIK, SUMANTA, and BOUATOUCH, KADI. “Radiance caching for efficient global illumination computation”. *IEEE Transactions on Visualization and Computer Graphics* 11.5 (2005), 550–561 2.
- [KK06] KOLLIG, THOMAS and KELLER, ALEXANDER. “Illumination in the Presence of Weak Singularities”. *Monte Carlo and Quasi-Monte Carlo Methods 2004*. Ed. by NIEDERREITER, HARALD and TALAY, DENIS. Springer, 2006, 245–257 4, 6, 7.
- [KMA*15] KETTUNEN, MARKUS, MANZI, MARCO, AITTALA, MIKA, et al. “Gradient-Domain Path Tracing”. *ACM Trans. Graph.* 34.4 (July 2015). ISSN: 0730-0301. DOI: [10.1145/2766997](https://doi.org/10.1145/2766997). URL: <https://doi.org/10.1145/2766997>.
- [KSAC02] KELEMEN, CSABA, SZIRMAI-KALOS, LÁSZLÓ, ANTAL, GYÖRGY, and CSONKA, FERENC. “A Simple and Robust Mutation Strategy for the Metropolis Light Transport Algorithm”. *Computer Graphics Forum* 21.3 (2002), 531–540. ISSN: 1467-8659. DOI: [10.1111/1467-8659.t01-1-00703](https://doi.org/10.1111/1467-8659.t01-1-00703). URL: <http://dx.doi.org/10.1111/1467-8659.t01-1-00703>.
- [KVB*19] KELLER, ALEXANDER, VIITANEN, TIMO, BARRÉ-BRISEBOIS, COLIN, et al. “Are We Done with Ray Tracing?”. *ACM SIGGRAPH 2019 Courses*. SIGGRAPH ’19. Los Angeles, California: Association for Computing Machinery, 2019. ISBN: 9781450363075. DOI: [10.1145/3305366.3329896](https://doi.org/10.1145/3305366.3329896). URL: <https://doi.org/10.1145/3305366.3329896>.
- [LW93] LAFORTUNE, ERIC P. and WILLEMS, YVES D. “Bi-Directional Path Tracing”. *Compugraphics ’93*. Alvor, Portugal, 1993, 145–153 3.
- [LW95] LAFORTUNE, ERIC P. and WILLEMS, YVES D. “A 5D tree to reduce the variance of Monte Carlo ray tracing”. *Rendering Techniques ’95 (Proc. of the 6th Eurographics Workshop on Rendering)*. 1995, 11–20. DOI: [10.1007/978-3-7091-9430-0_23](https://doi.org/10.1007/978-3-7091-9430-0_23).
- [ME10] MARTIN, SAM and EINARSSON, PER. *A Real-Time Radiosity Architecture for Video Games*. SIGGRAPH 2010 Course: Advances in Real-Time Rendering in 3D Graphics and Games. 2010 2.
- [MGN17] MÜLLER, THOMAS, GROSS, MARKUS, and NOVÁK, JAN. “Practical Path Guiding for Efficient Light-Transport Simulation”. *Computer Graphics Forum* 36.4 (June 2017), 91–100. ISSN: 1467-8659. DOI: [10.1111/cgf.13227](https://doi.org/10.1111/cgf.13227). URL: <http://dx.doi.org/10.1111/cgf.13227>.
- [MGNM19] MAJERCIK, ZANDER, GUERTIN, JEAN-PHILIPPE, NOWROUZEZAHRAI, DEREK, and MCGUIRE, MORGAN. “Dynamic Diffuse Global Illumination with Ray-Traced Irradiance Fields”. *Journal of Computer Graphics Techniques (JCGT)* 8.2 (2019), 1–30. ISSN: 2331-7418. URL: <http://jcg.t.org/published/0008/02/01/2-5,7,9,10>.
- [MH20] MUNKBERG, JACOB and HASSELGREN, JON. “Neural Denoising with Layer Embeddings”. *Computer Graphics Forum* 39.4 (2020), 1–12. DOI: <https://doi.org/10.1111/cgf.14049>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14049>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14049>.
- [MMBJ17] MARA, MICHAEL, MCGUIRE, MORGAN, BITTERLI, BENEDIKT, and JAROSZ, WOJCIECH. “An Efficient Denoising Algorithm for Global Illumination”. *ACM SIGGRAPH / Eurographics High Performance Graphics*. HPG 2017. July 2017, 7. URL: <https://casual-effects.com/research/Mara2017Denoise/index.html>.
- [MMR*19] MÜLLER, THOMAS, MCGUIRE, MORGAN, BITTERLI, BENEDIKT, and JAROSZ, WOJCIECH. “An Efficient Denoising Algorithm for Global Illumination”. *ACM SIGGRAPH / Eurographics High Performance Graphics*. HPG 2017. July 2017, 7. URL: <https://casual-effects.com/research/Mara2017Denoise/index.html>.
- [MMR*19] MÜLLER, THOMAS, MCGUIRE, MORGAN, BITTERLI, BENEDIKT, and JAROSZ, WOJCIECH. “An Efficient Denoising Algorithm for Global Illumination”. *ACM SIGGRAPH / Eurographics High Performance Graphics*. HPG 2017. July 2017, 7. URL: <https://casual-effects.com/research/Mara2017Denoise/index.html>.
- [MMSM21] MAJERCIK, ZANDER, MARRS, ADAM, SPJUT, JOSEF, and MCGUIRE, MORGAN. “Scaling Probe-Based Real-Time Dynamic Global Illumination for Production”. *Journal of Computer Graphics Techniques (JCGT)* 10.2 (May 2021), 1–29. ISSN: 2331-7418. URL: <http://jcg.t.org/published/0010/02/01/1-3,5,7-9,11>.
- [MRKN20] MÜLLER, THOMAS, ROUSSELLE, FABRICE, KELLER, ALEXANDER, and NOVÁK, JAN. “Neural Control Variates”. *ACM Trans. Graph.* 39.6 (Nov. 2020), 243:1–243:19. ISSN: 0730-0301. DOI: [10.1145/3414685.3417804](https://doi.org/10.1145/3414685.3417804). URL: <https://doi.org/10.1145/3414685.3417804>.

- [MRNK21] MÜLLER, THOMAS, ROUSSELLE, FABRICE, NOVÁK, JAN, and KELLER, ALEXANDER. “Real-time Neural Radiance Caching for Path Tracing”. *ACM Trans. Graph.* 40.4 (Aug. 2021), 36:1–36:16. DOI: [10.1145/3450626.3459812](https://doi.org/10.1145/3450626.3459812). URL: <https://doi.org/10.1145/3450626.3459812> 2, 5, 6, 10, 11.
- [MZV*20] MENG, XIAOXU, ZHENG, QUAN, VARSHNEY, AMITABH, et al. “Real-time Monte Carlo Denoising with the Neural Bilateral Grid”. *Eurographics Symposium on Rendering - DL-only Track*. Ed. by DACHSBACHER, CARSTEN and PHARR, MATT. The Eurographics Association, 2020. ISBN: 978-3-03868-117-5. DOI: [10.2312/sr.20201133](https://doi.org/10.2312/sr.20201133) 10.
- [Oat05] OAT, CHRISTOPHER. “Irradiance Volumes for Games”. *Game Developers Conference*. 2005 2.
- [OLK*21a] OUYANG, YAOBIN, LIU, SHIQIU, KETTUNEN, MARKUS, et al. “ReSTIR GI: Path Resampling for Real-Time Path Tracing”. *Computer Graphics Forum* (2021). ISSN: 1467-8659. DOI: [10.1111/cgf.14378](https://doi.org/10.1111/cgf.14378) 2.
- [OLK*21b] OUYANG, YAOBIN, LIU, SHIQIU, KETTUNEN, MARKUS, et al. “ReSTIR GI: Path Resampling for Real-Time Path Tracing”. *Computer Graphics Forum* (2021). ISSN: xxx. DOI: [10.1111/xxx.3](https://doi.org/10.1111/xxx.3) 10.
- [OZ98] OWEN, ART and ZHOU, YI. “Safe and Effective Importance Sampling”. *Journal of the American Statistical Association* 95 (1998), 135–143 2.
- [RZD14] REHFELD, HAUKE, ZIRR, TOBIAS, and DACHSBACHER, CARSTEN. “Clustered Pre-Convolved Radiance Caching”. *Proceedings of the 14th Eurographics Symposium on Parallel Graphics and Visualization*. PGV '14. Swansea, Wales, United Kingdom: Eurographics Association, 2014, 25–32. ISBN: 9783905674590 2.
- [SJJ12] SCHWARZHaupt, JORGE, JENSEN, HENRIK WANN, and JAROSZ, WOJCIECH. “Practical Hessian-Based Error Control for Irradiance Caching”. *ACM Trans. Graph.* 31.6 (Nov. 2012). ISSN: 0730-0301. DOI: [10.1145/2366145.2366212](https://doi.org/10.1145/2366145.2366212). URL: <https://doi.org/10.1145/2366145.2366212> 2.
- [SKS02] SLOAN, PETER-PIKE, KAUTZ, JAN, and SNYDER, JOHN. “Pre-computed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments”. *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '02. San Antonio, Texas: Association for Computing Machinery, 2002, 527–536. ISBN: 1581135211. DOI: [10.1145/566570.566612](https://doi.org/10.1145/566570.566612). URL: <https://doi.org/10.1145/566570.566612> 2.
- [SKW*17] SCHIED, CHRISTOPH, KAPLANYAN, ANTON, WYMAN, CHRIS, et al. “Spatiotemporal Variance-Guided Filtering: Real-Time Reconstruction for Path-Traced Global Illumination”. *Proceedings of High Performance Graphics*. HPG '17. Los Angeles, California: Association for Computing Machinery, 2017. ISBN: 9781450351010. DOI: [10.1145/3105762.3105770](https://doi.org/10.1145/3105762.3105770). URL: <https://doi.org/10.1145/3105762.3105770> 2.
- [SL17] SILVENNOINEN, ARI and LEHTINEN, JAAKKO. “Real-Time Global Illumination by Precomputed Local Reconstruction from Sparse Radiance Probes”. *ACM Trans. Graph.* 36.6 (Nov. 2017). ISSN: 0730-0301. DOI: [10.1145/3130800.3130852](https://doi.org/10.1145/3130800.3130852). URL: <https://doi.org/10.1145/3130800.3130852> 2.
- [SNRS12] SCHERZER, DANIEL, NGUYEN, CHUONG H., RITSCHEL, TOBIAS, and SEIDEL, HANS-PETER. “Pre-Convolved Radiance Caching”. *Comput. Graph. Forum* 31.4 (June 2012), 1391–1397. ISSN: 0167-7055. DOI: [10.1111/j.1467-8659.2012.03134.x](https://doi.org/10.1111/j.1467-8659.2012.03134.x). URL: <https://doi.org/10.1111/j.1467-8659.2012.03134.x> 2.
- [SSS*20] SEYB, DARIO, SLOAN, PETER-PIKE, SILVENNOINEN, ARI, et al. “The design and evolution of the UberBake light baking system”. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39.4 (July 2020). awaiting publication 2.
- [TCE05] TALBOT, JUSTIN F., CLINE, DAVID, and EGBERT, PARRIS. “Importance Resampling for Global Illumination”. *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*. EGSR '05. Konstanz, Germany: Eurographics Association, 2005, 139–146. ISBN: 3-905673-23-1. DOI: [10.2312/EGWR/EGSR05/139-146](https://doi.org/10.2312/EGWR/EGSR05/139-146). URL: <https://doi.org/10.2312/EGWR/EGSR05/139-146> 2, 3.
- [Vea97] VEACH, E. “Robust Monte Carlo Methods for Light Transport Simulation”. PhD thesis. Stanford University, 1997 3.
- [VG94] VEACH, ERIC and GUIBAS, LEONIDAS J. “Bidirectional estimators for light transport”. *EG Rendering Workshop*. 1994 3.
- [VG95] VEACH, ERIC and GUIBAS, LEONIDAS J. “Optimally Combining Sampling Techniques for Monte Carlo Rendering”. *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '95. New York, NY, USA: Association for Computing Machinery, 1995, 419–428. ISBN: 0897917014. DOI: [10.1145/218380.218498](https://doi.org/10.1145/218380.218498). URL: <https://doi.org/10.1145/218380.218498> 2, 3, 6.
- [VG97] VEACH, ERIC and GUIBAS, LEONIDAS J. “Metropolis Light Transport”. *Proc. SIGGRAPH 1997*. Annual Conference Series. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, 65–76 3.
- [VHH*19] VORBA, JIŘÍ, HANIKA, JOHANNES, HERHOLZ, SEBASTIAN, et al. “Path Guiding in Production”. *ACM SIGGRAPH 2019 Courses*. SIGGRAPH '19. Los Angeles, California: ACM, 2019, 18:1–18:77. ISBN: 978-1-4503-6307-5. DOI: [10.1145/3305366.3328091](https://doi.org/10.1145/3305366.3328091). URL: <https://doi.org/10.1145/3305366.3328091> 2, 3, 10.
- [Vit85] VITTER, JEFFREY S. “Random Sampling with a Reservoir”. *ACM Trans. Math. Softw.* 11.1 (Mar. 1985), 37–57. ISSN: 0098-3500. DOI: [10.1145/3147.3165](https://doi.org/10.1145/3147.3165). URL: <https://doi.org/10.1145/3147.3165> 3.
- [VKŠ*14] VORBA, JIŘÍ, KARLÍK, ONDŘEJ, ŠIK, MARTIN, et al. “On-line Learning of Parametric Mixture Models for Light Transport Simulation”. *ACM Trans. Graph.* 33.4 (Aug. 2014) 3.
- [VPG14] VARDIS, KOSTAS, PAPAIOANNOU, GEORGIOS, and GKARAVELIS, ANASTASIOS. “Real-time Radiance Caching using Chrominance Compression”. *Journal of Computer Graphics Techniques (JCGT)* 3.4 (Dec. 2014), 111–131. ISSN: 2331-7418. URL: <http://jcg.t.org/published/0003/04/06/> 2.
- [VRM*18] VOGELS, THIJS, ROUSSELLE, FABRICE, MCWILLIAMS, BRIAN, et al. “Denoising with Kernel Prediction and Asymmetric Loss Functions”. *ACM Trans. Graph.* 37.4 (July 2018). ISSN: 0730-0301. DOI: [10.1145/3197517.3201388](https://doi.org/10.1145/3197517.3201388). URL: <https://doi.org/10.1145/3197517.3201388> 2.
- [WRC88] WARD, GREGORY J., RUBINSTEIN, FRANCIS M., and CLEAR, ROBERT D. “A Ray Tracing Solution for Diffuse Interreflection”. *SIGGRAPH Comput. Graph.* 22.4 (June 1988), 85–92. ISSN: 0097-8930. DOI: [10.1145/378456.378490](https://doi.org/10.1145/378456.378490). URL: <https://doi.org/10.1145/378456.378490> 2.
- [XZW*19] XU, BING, ZHANG, JUNFEI, WANG, RUI, et al. “Adversarial Monte Carlo Denoising with Conditioned Auxiliary Feature Modulation”. *ACM Trans. Graph.* 38.6 (Nov. 2019). ISSN: 0730-0301. DOI: [10.1145/3355089.3356547](https://doi.org/10.1145/3355089.3356547). URL: <https://doi.org/10.1145/3355089.3356547> 2.
- [ZZ19] ZHENG, QUAN and ZWICKER, MATTHIAS. “Learning to Importance Sample in Primary Sample Space”. *Computer Graphics Forum* 38.2 (2019), 169–179. DOI: [10.1111/cgf.13628](https://doi.org/10.1111/cgf.13628) 3.