

# Path Space Regularization for Holistic and Robust Light Transport

Anton S. Kaplanyan and Carsten Dachsbacher

Karlsruhe Institute of Technology, Germany




Figure 1: Equal-time rendering (2 hours) of the BATHROOM scene with a pinhole camera and indoor illumination from an LED panel. The LED panel showing the EG logo is modelled of 4194 point light sources (not covered by a fixture). The reflection of these light sources in the mirror is missing with previous methods as neither path tracing, nor Metropolis light transport, nor progressive photon mapping can handle all light transport paths. With our regularization, applicable to all unbiased methods, even path tracing can sample such difficult paths in a consistent way – without introducing bias to the rest of the image.

## Abstract

*We propose a simple yet powerful regularization framework for robust light transport simulation. It builds on top of existing unbiased methods and resorts to a consistent estimation using regularization only for paths which cannot be sampled in an unbiased way. To introduce as little bias as possible, we selectively regularize individual interactions along paths, and also derive the regularization consistency conditions. Our approach is compatible with the majority of unbiased methods, e.g. (bidirectional) path tracing and Metropolis light transport (MLT), and only a simple modification is required to adapt existing renderers. We compare to recent unbiased and consistent methods and show examples of scenes with difficult light paths, where regularization is required to account for all illumination features. When coupled with MLT we are able to sample all phenomena, like recent consistent methods, while achieving superior convergence.*

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

**Keywords:** global illumination, light transport, regularization, Markov chain Monte-Carlo

## 1. Introduction

Ever since generating synthetic images with computers became possible, researchers in computer graphics were fascinated by the reproduction of realistic scenes. The key to photorealism is a physically-based global illumination computation. Although the problem can be formulated as the well-known rendering equation [Kaj86], not all computational models provide a full solution. High-quality rendering, which we target in this work, is nowadays based on Monte-Carlo or Markov chain Monte Carlo methods. However, it is often ignored that, as pointed out by Veach [Vea98], there are configurations of scene geometry, materials, and lighting that cannot be handled with existing (unbiased) meth-

ods. In particular, these include reflected caustics (specular-diffuse-specular chains) or multiple consecutive specular interactions and point light sources as shown in Fig. 1.

Global illumination methods are often used for predictive rendering in architecture, product concept design, and movie productions. With these applications, we can nowadays observe more and more involved scenarios for light transport simulation, e.g. complex materials or lighting from light-emitting diodes (LEDs) used in cars and interior lighting. The emitting surfaces in LEDs have very small area (ranging from a few square millimeters to a few square micrometers) and such luminaires are thus often treated as point or almost point light sources which raises a lot of problems in light

transport simulation with existing unbiased methods. Moreover, as pointed out by Hachisuka et al. [HOJ08], if a light source is enclosed in a glass fixture, the majority of unbiased methods fails to simulate the light transport carried by such paths in presence of mirrors. Many of these problems stem from singularities in the integrand of the rendering equation.

In this paper we propose a new regularization framework for handling this part of the light transport in a consistent (biased) way, while the remaining transport can still be computed in an unbiased way. The core of our idea is to identify paths which cannot be sampled because of singularities of the integrand during construction. We then use regularization at a minimum of the vertices of such a path to turn it into a samplable path, yet with as little bias as possible.

In particular our contributions are:

- A general regularization framework that easily integrates into existing renderers;
- Robust light transport simulation by selective regularization for complex illumination;
- Consistent estimation for paths which cannot be sampled using unbiased methods.

## 2. Related Work

Kajiya [Kaj86] formulated the light transport problem as an integral equation and was the first who proposed an unbiased algorithm for computing a image by sampling paths. Veach [Vea98] introduced the path integral form for light transport along with several important methods for variance reduction and sampling of the path space, e.g. Metropolis light transport. Important for our work is his observation that there exist configurations of geometry, light sources, and sensors, which cannot be sampled by any unbiased method.

Jakob and Marschner [JM12] focused on one type of such difficult light transport through specular manifolds. They pointed out that the integrand becomes ill-posed due to the manifold constraints of specular interactions and proposed a combination of a local linear solver and ray tracing for walking along manifolds. They also mentioned that their method cannot find such paths on itself and only allows for efficient exploration of the manifold from a known path on it.

Hachisuka et al. [HOJ08] and Hachisuka and Jensen [HJ09] recently showed that another types of difficult configurations can be handled by consistent methods, e.g. by progressive photon mapping, where the systematic error vanishes in the limit. However, these methods do not address the problem of difficult paths explicitly, which is usually done using random walk methods [Vea98]. Following this idea, Fan et al. [FCL05], Hachisuka and Jensen [HJ11] and Chen et al. [CWY11] introduced Metropolis sampling for photon shooting. However, camera subpaths are still sampled randomly as in path tracing and this, as pointed out by the authors, causes slow image convergence, e.g. with

glossy surfaces [HJ11]. Kniep et al. [KHM09] considered difficult configurations of directly visible light sources and extended photon mapping with angular smoothing in addition to the spatial density estimation. Our regularization of pure specular materials is similar to this work, however, in contrast, we smooth only a fraction of light paths and only in the angular domain which introduces bias in fewer places and leads to an overall faster convergence.

In general, photon mapping methods blur the radiance signal even where it can be well handled by unbiased methods. This problem was recently addressed by a series of works, e.g. Vorba [Vor11], Georgiev et al. [GKDS12] and Hachisuka et al. [HPJ12]. In order to reduce bias these techniques attempt to combine PPM with bidirectional path tracing (BPT) using multiple importance sampling (MIS) [VG94]. However, both methods have different convergence rates [KD13] and the unbiased method gradually outweighs the biased one. This might lead to redundant computation on the later stages of image convergence. The engineering complexity of these approaches is also very high [HPJ12]. Note that MIS in this case only accounts for the sampling probability, but not for the absolute amount of systematic error (bias) admixed. Recently, Kaplanyan and Dachsbacher [KD13] showed that photon mapping can be seen as a uniform regularization in path space. Our work elaborates on this concept, but comprises a general framework for selective regularization in path space.

## 3. Preliminaries

In this section we review the path integral formulation of light transport and full-path regular expressions. We also introduce the concept of the regularization with mollification.

### 3.1. Path Integral of Light Transport

The path integral formulation [Vea98] expresses the function of camera response measurement in the form of the integral

$$I = \int_{\Omega} f(\bar{x}) d\mu(\bar{x}), \quad (1)$$

where  $\Omega = \Omega(\mathcal{M})$  is the unified space of all possible light paths of all lengths;  $\bar{x} \in \Omega(\mathcal{M})$  is a complete path from light to camera in path space, which is represented as a vector of points on the scene manifold  $\mathcal{M}$ , i.e.  $\bar{x} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k)_{k=1 \dots \infty}$ ;  $I$  is a camera sensor measurement;  $f(\bar{x})$  is a measurement contribution function, which includes a product of reflection operators at all points of an argument path  $\bar{x}$  of length  $k$ :

$$f(\bar{x}) = L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1) G(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1) \left( \prod_{j=1}^{k-1} f_s(\mathbf{x}_{j-1} \rightarrow \mathbf{x}_j \rightarrow \mathbf{x}_{j+1}) G(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1}) \right) W_e(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) \quad (2)$$

where  $G(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1})$  is the geometry factor between the points  $\mathbf{x}_j$  and  $\mathbf{x}_{j+1}$ ;  $f_s(\mathbf{x}_{j-1} \rightarrow \mathbf{x}_j \rightarrow \mathbf{x}_{j+1})$  is the bidirectional reflectance distribution function (BRDF) at the point

$\mathbf{x}_j$ ;  $L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1)$  is the outgoing radiance of the light source and  $W_e(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k)$  the sensor sensitivity. Veach [Vea98] provides more details on the path integral formulation.

### 3.2. Full-Path Regular Expressions

We will use Veach's extended notation [Vea98] of the Heckbert notation [Hec90] which also describes the properties of light sources and sensors: a light path begins with  $L(S|D)(S|D)$ , where the first letter after  $L$  (indicating the begin on the light source) denotes a finite-area source  $D$  or source  $S$  with zero area, and the second letter denotes emission over a finite solid angle  $D$  or a set of directions with measure zero  $S$ . Analogously, the path ends with  $(S|D)(S|D)E$  where the first letter denotes the directional sensitivity of the sensor and the second whether the sensor has a finite area or is a point. Hereafter we denote any of  $D$  or  $S$  types with the letter  $A \equiv (S|D)$ . Letters between this prefix and suffix correspond to path vertices as in Heckbert's original notation. Throughout the paper, we will use highlighting to more easily distinguish the vertices corresponding to light sources (green), interactions (black), and camera (brown). For example,  $LSD(S)^+DDSE$  describes paths producing caustics cast by a point light source, going through one or more mirrors to a pinhole camera.

### 3.3. Paths that Cannot Be Sampled

In his thesis, Veach [Vea98, Theorem 8.2] defines the problem of samplable paths:

**Theorem 1.** *Let  $\bar{x}$  be a path generated by a local sampling algorithm for which the measurement contribution function is non-zero. Then  $\bar{x}$  necessarily has the form  $LA^*DDA^*E$ , i.e. it must contain the substring  $DD$ . Furthermore, it is possible to generate any path of this form using local sampling strategies.*

In other words, the probability of sampling paths different from Theorem 1 is zero for unbiased methods with local sampling, including path tracing (PT) and bidirectional path tracing (BPT), which are also required to initialize Metropolis light transport (MLT). Note that path tracing can only sample a subset of the paths which are possible with BPT, e.g.  $LSDSD^+SDE$  can only be sampled with a complete set of bidirectional estimators. The mutation strategies for MLT proposed by Veach [Vea98] are also based on local sampling and thus obey this theorem. Important to note that the recent manifold exploration mutation strategy [JM12] could theoretically explore such non-samplable paths. However, the algorithm cannot find such paths on its own, i.e. finding these paths remains an unsolved problem (see Fig 5 for example).


In contrast, consistent methods, such as photon mapping and progressive photon mapping [HOJ08, HJ09], relax the condition of Theorem 1, enabling the sampling of complex paths of the form  $LAAA^*DA^*AAE$ . This is achieved by generating light subpaths  $LAAA^*D$  and camera subpaths

$DA^*AAE$  independently, and then merging the  $D$ -vertices based on the spatial proximity condition. That is equivalent to regularizing one connection vertex. However these methods regularize for each path they sample and yet cannot construct all potential paths, e.g. paths of the form  $LSA(S)^+AAE$ , which correspond to a path going from a point light source to a pinhole camera through a pure specular manifold (see Fig. 1).

### 3.4. Mollification

In this section we discuss the approximation of delta distributions by integrable functions. This process will be used in our work to enable the sampling of difficult parts of the path space. The Dirac delta distribution, which introduces the singularities is zero everywhere except at  $x = 0$ , where  $\delta(0) = +\infty$ , and obeys  $\int_{\mathbb{R}} f(x)\delta(x)dx = f(0)$  for any absolutely integrable function  $f \in L_1$ . Throughout the paper, we distinguish between delta functions in spatial or angular domain by the argument.

Typically, in order to approximate a  $d$ -dimensional delta distribution  $\delta$  by integrable functions, one constructs a sequence of smooth positive functions  $\{\phi_r\} \subset L_1$ , such that



$$\left\{ \begin{array}{l} \|\text{supp}(\phi_r)\| \propto O(r^d) \\ \int_{\mathbb{R}^d} \phi_r(x) dx = 1 \end{array} \right\} \Rightarrow \lim_{r \rightarrow 0} \phi_r(x) = \delta(x), \quad (3)$$

where  $\|\text{supp}(\phi_r)\|$  is a mass of the function support. In other words, a normalized sequence of smooth functions, whose supports are vanishing as  $O(r^d)$ , implies that it approximates the  $d$ -dimensional delta function. The sequence  $\{\phi_r\}$  is called a *mollifier* of the delta function.

It is usually expressed in a form of  $\phi_r(x) = 1/r^d \phi(x/r)$ , where  $\phi(x)$  is a smooth normalized canonic kernel. The parameter  $r$  is called the *mollification bandwidth* and has to be gradually reduced to zero in the limit. We will discuss the reduction condition for the bandwidth for different integration methods in Sect. 5. Hegland and Anderssen [HA96] provide further information on mollifiers and their properties.

A simple example of mollification for specular reflections would be to replace the delta distribution in BSDFs by a function with finite support (see Sect. 4.2).

## 4. Path Space Regularization

In this section, we introduce our concept of selective regularization of the path space which will enable us to sample arbitrary paths with local sampling methods. The core idea is to turn a path which cannot be sampled into a path of the type from Th. 1 by selectively mollifying some interactions, i.e. by turning them from  $S$  into  $D$ . There are two major reasons for  $S$  vertices to occur in a path: (1) irregular sensors




Figure 2: A path that cannot be sampled in an unbiased way. Note that a regularization of either of the  $S$ -vertices on the lower blue surface (here illustrated for the right vertex) enables us to sample the path (see Sect. 4).

and light sources; and (2) pure specular materials causing delta responses in the BRDF.

#### 4.1. Singularities in the Integrand of the Path Integral

Light transport can exhibit very narrow modes and singularities due to delta, or near-delta, distributions in the integrand of the path integral in Eq. 2. Some singularities cannot be explicitly sampled with unbiased integrators (e.g. the double-reflection in Fig. 2) and thus some parts of the light transport are not accounted for in the global illumination solution. In the following we discuss when such modes and singularities appear during the path construction process.

**Singularities in BRDFs.** Perfect and almost perfect mirrors cause arbitrarily narrow modes or even delta responses, resulting in configurations that provably cannot be sampled [RTY94]. A perfect specular reflection, for example, can be formulated as an angular delta distribution [Vea98]:

$$f_s(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k \rightarrow \mathbf{x}_{k+1}) = \delta(\omega - \omega_i) / \langle \omega_i \cdot n \rangle, \quad (4)$$

where  $\omega_i = (\mathbf{x}_{k-1} - \mathbf{x}_k) / \|\mathbf{x}_{k-1} - \mathbf{x}_k\|$  is the incident direction and  $n$  is a surface normal at the point  $\mathbf{x}_k$ . The term  $\delta(\omega - \omega_i)$  is a two-dimensional delta distribution on the sphere of directions  $\mathbb{S}^2$  and is equal to zero everywhere, except for the direction of perfect specular reflection  $\omega$ . The BSDF of a specular refraction can be described analogously by taking into account the indices of refraction.

**Singularities in Sensors and Emitters.** Other models used in computer graphics also introduce singularities: a pinhole camera or a point light source, for example, have zero area, while directional light sources or orthographic cameras have a fixed direction. This can lead to paths which cannot be sampled with stochastic light transport methods. Let  $\mathbf{x}_i$  denote the vertices of a light path, then a pinhole camera and a point light source models can be defined as:

$$\begin{aligned} W(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) &= \delta(\mathbf{x}_k - \mathbf{x}_{camera}) W_0(\omega); \\ L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1) &= \delta(\mathbf{x}_0 - \mathbf{x}_{light}) L_0(\omega), \end{aligned}$$

where  $W(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k)$  is the sensor response function of a pinhole camera;  $L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1)$  is the radiance distribution of the point light source; and  $W_0(\omega)$  and  $L_0(\omega)$  are the angular sensitivity and radiant intensity correspondingly. Both models have a spatial delta distribution which makes sampling

them stochastically impossible as the probability of hitting a point (light or pinhole camera) by randomly shooting rays is zero. Another example is a thin laser beam which could be modelled as both spatial and angular delta distributions.

#### 4.2. Example: Regularization of Specular Interactions

If a path cannot be sampled due to a delta function in the BSDF, as in Eq. 4, we can regularize it by applying mollification in the angular domain, e.g. for next event estimation in path tracing on specular surfaces. The simplest mollifier, meeting the conditions in Eq. 3, to replace  $\delta(\omega - \omega_i)$  is the constant angular mollifier:

$$\varphi_\epsilon(\omega, \omega') = \frac{1}{2\pi(1 - \cos \epsilon)} \mathbf{1}_{\angle(\omega, \omega') < \epsilon}, \quad (5)$$

where  $\mathbf{1}$  is an indicator function, i.e. this function takes non-zero values only if the angle between  $\omega$  and  $\omega'$  is less than  $\epsilon$ ; normalization of the mollifier is achieved using the factor  $1/(2\pi(1 - \cos \epsilon))$ . Normalized smooth kernels with a vanishing first moment can improve the convergence slightly, but we decided to use indicator for brevity.

Instead of letting the user choose an initial mollification angle, we opt for making the parameter choice similar to photon mapping. To this end, we compute  $\epsilon$  based on the distance  $l$  between the two involved path vertices and a user-specified spatial radius  $r$  as  $\epsilon = \arctan(r/l)$  (see Fig. 3, left). Note that for  $r \ll l$  the term  $\arctan(r/l) \approx r/l$ . That allows us to use  $r$  instead of  $\epsilon$  in the asymptotic convergence conditions in Sect. 5.

#### 4.3. Other Types of Regularization

**Spatial Regularization** can be easily introduced, e.g. using a spherical spatial mollifier  $\varphi_r(\mathbf{x}, \mathbf{x}_0) = \frac{3}{4\pi r^3} \mathbf{1}_{\|\mathbf{x} - \mathbf{x}_0\| < r}$ . This mollifier also obeys the conditions in Eq. 3. Spatial mollification can be useful in participating media rendering, e.g. in spirit of volumetric photon mapping [JC98]. However, we recommend to avoid spatial mollification if possible, as the spatial delta function is 3-dimensional (compared to the 2-dimensional angular delta function of specular BRDF) and, as we will see in Sect. 5, the dimensionality of the delta functions directly affects the convergence of the integration.

**Near-Singular Interactions.** In general, most of the described non-samplable paths do not exist in the real world where no perfect point light sources or pinhole cameras exist. However, in terms of sampling, near-singular functions, e.g. highly-glossy BRDFs, are almost as problematic as singularities, as the sampling probability can be very close to zero. In general, mollification can be applied on top of any near-singular function leading to consistent integration.

Regularization for BRDFs depends on the representation: for measured BRDFs it can be easily achieved using a numerical convolution. For the regularization of narrow lobes of analytic BRDFs we suggest the following options:



- Importance-sample the BRDF and then connect only if the connecting vertex lies within the mollification cone of the BRDF sample. This is statistically correct and equivalent to the procedure that photon mapping methods perform.
- We could extract the major asymptotic term of the solid angle quantile [JM12] of the analytic model w.r.t. its roughness parameters. Using this term it is possible to adjust the roughness of the model according to the mollification bandwidth when regularization is required.

In general, the decision on when to regularize near-specular paths should depend on various factors, such as performance-accuracy trade-offs or measures like the curvature of the local light field. The problem of near-singular transport (due to narrow BRDF lobes and complex illumination) can also be addressed using the manifold exploration mutation strategy [JM12]. Note that our work focuses on an orthogonal problem: finding *non-samplable* transport paths.

#### 4.4. Selective Regularization

An ad-hoc regularization approach – to mollify all delta distributions in the integrand – would lead to a large bias, in particular at the early stages of rendering. To keep the bias to a minimum, we mollify as few interactions as possible (e.g. only one  $S$ -vertex in Fig. 2). For a selective regularization we always have to determine all interactions along a path first. Prior to that, it is not possible to determine whether it can be constructed in an unbiased way or not. And thus we can also not decide earlier whether it is necessary, and where to regularize the path.

**Example: Regularization with Path Tracing.** If the paths are generated using unidirectional path tracing, we decide about the regularization only at the next event estimation (directly sampling light sources), because the path type is completely formed only at this stage. Consequently, the only vertex we can mollify at this stage, without regenerating the path, is either the last interaction type or the position on the light source (or both if necessary). Note that for a unidirectional path tracer, every path that ends with "S" is non-samplable. Thus there are more types of paths that cannot be handled by this technique since PT is a subset of BPT.

Consider an example where a subpath from a pinhole camera to a specular surface  $SDSE$ , which is to be connected to a point light source  $LSD$ , can be turned into a samplable path by mollifying the last-generated interaction type yielding  $LSDDDSE$  (Fig. 3, left). Note that the mollification of the light source position and/or of the singular BSDF at this interaction vertex would enable us to construct the path.

#### 4.5. Regularization with Bidirectional Estimators

When using a bidirectional estimator, the decision can be made only when connecting the two subpaths, as the type of the full path is not known before. And at this point, we




Figure 3: Left: regularization with mollification at the next event estimation in unidirectional sampling. Right: regularization at two interactions with bidirectional sampling.

can regularize either of the two connection vertices or both to eventually turn it into a path of the type in Theorem 1. Fig. 3 (right) shows an example where the regularization of interactions at both connection vertices is required.

In bidirectional methods, e.g. BPT, we typically use a set of bidirectional estimators which contains all possibilities to construct the respective path. If the resulting path can be sampled in an unbiased way, there is no need to apply regularization and we handle it in an unbiased way as in original BPT, weighting the respective contributions of the bidirectional estimators in the set using multiple importance sampling (MIS) [VG94]. However, if the path cannot be sampled in an unbiased way, we regularize it. It is important to note that the set of bidirectional estimators for such regularized path always consists solely of biased estimators. In the following we introduce a weighting strategy for regularized paths which can be seen as a counterpart of MIS for a set of biased estimators.

**Weighting Strategy for Biased Estimators.** Consider an example of a path  $LSDSDSE$  (as in Fig. 2) which can be turned into a samplable path by mollifying the BSDF at the first or second specular vertex in the middle part, which yields two different estimators  $E_1$  and  $E_2$ . In order to achieve consistent results we have to choose the weights for combining these estimators,  $w_1$  and  $w_2$ , such that  $w_1 + w_2 = 1$ , analogously to MIS. However, in contrast to MIS, both estimators with regularization are biased and the goal of the combining strategy should be bias minimization.

For this we propose to use a *maximum distance heuristic* which assigns  $w_i = 1$  if the estimator  $E_i$  corresponds to the path with the longest distance  $l$  between the two connection vertices among the estimators with the lowest possible number of regularized dimensions; otherwise we assign  $w_i = 0$ . This is equivalent to a minimization of the mollification bandwidth  $r$ . Hegland and Anderssen [HA96] also show that the mollification error (bias) directly depends on the bandwidth.

**Discussion.** When using regularization for near-singular interactions (Sect. 4.3), we would face the problem of combining consistent with unbiased estimators of the same path. In this case one can use the original MIS with the bias considerations as in [GKDS12].

## 5. Regularization with Integration Methods

As we have seen, it is straightforward to apply regularization to both unidirectional and bidirectional sampling. This enables us to apply our regularization framework to all existing unbiased methods, including bidirectional path tracing and Metropolis light transport.

Note that the regularization with mollification is transparent to the integration method: the integration method does not require changes, instead the regularization selectively smoothes singular modes in the underlying integrand. Thus all existing variance reduction techniques, sampling methods, and mutation strategies can be used together with this type of regularization without any changes.

We will describe how to gradually reduce the mollification bandwidth with each integration iteration to achieve consistent results in the limit. We derive the shrinkage rate conditions to compute the bandwidth  $r_n$  for the  $n$ -th iteration from the user-specified initial mollification bandwidth  $r_0$ .

### 5.1. Consistent MC Integration with Regularization

In order to achieve a consistent estimation of the path integral in Eq. 1 with Monte-Carlo (MC) integration methods, such as PT and BPT, the mollification bandwidth has to be gradually reduced after each step of the integration. Recall that one or both interactions might require regularization. For the latter case, the two delta distributions can be combined into one using a combined product measure. In this case the dimensionality of the new-formed delta distribution is the sum of dimensionalities of the two original ones. If the mollification is applied to a  $d$ -dimensional delta distribution, the reduction rate has to be in the following boundaries (see App. A for a proof):

$$O(n^{-1/d}) < r_n < O(1), \quad (6)$$

where  $n$  is the index of the sample taken for the integration. As a practical and simple sequence we suggest to use  $r_n = r_0 n^{-\lambda}$ , where  $r_0$  is the user-specified initial mollification radius and  $\lambda \in (0; 1/d)$  for a  $d$ -dimensional mollification. Note that this range evaluated for a 2D specular mollification matches the one in PPM for  $r_n = r_0 n^{(\alpha-1)/2}$  with  $\alpha \in (0; 1)$  [HOJ08, KD13]. Similarly to the optimal parameter for PPM [KD13], we choose  $\lambda = 1/6$  for a single-vertex specular mollification and  $\lambda = 1/12$  for two vertices. This mollification parameter  $\lambda$  is defined globally and the bandwidth shrinking causes the regularized parts of the integrand to become sharper with every integration step, yielding a consistent estimation in the limit.

Note that even path tracing with regularization of specular interactions is already able to handle all light transport paths. In the supplementary material we provide a modified version of “SmallPT”, a minimalistic path tracer by Kevin Beason, demonstrating the simplicity of our method. How-

ever, PT and also BPT are significantly less efficient for sampling regularized paths than MCMC methods introduced in the next section (see Sect. 7 for a discussion).

### 5.2. Consistent MCMC Methods with Regularization

Markov chain Monte-Carlo (MCMC) methods, such as Metropolis light transport (MLT), greatly improve on the slow convergence of MC methods: once a narrow mode is found, it is explored thoroughly by the Markov process. Regularization can also be applied to MCMC methods to enable a consistent estimation of paths which cannot be handled with the original methods.

Our proposed method combines the simulated annealing optimization method with MCMC integration. Intuitively, the “sharp peaks” of the integrand are regularized (smoothed out) in the beginning of the integration such that the Markov chain can easily find them. In order to achieve consistency, the regularization is reduced throughout the integration. The important condition is to “sharpen” these peaks slower than the integrand converges, as otherwise the Markov chain might get stuck in the sharp peak for a long time, jeopardizing the convergence of other parts of the image. When set up correctly, MCMC methods solve the slow convergence problem of MC methods along with the problem of difficult paths in consistent methods (see Fig. 4). Note that the Markov chain finds regularized paths even if it is not initialized with any of them, i.e. there is no need to initialize with non-samplable paths.

We start by creating initial paths using unbiased local sampling (BPT in our implementation). While mutating them, we apply selective regularization to non-samplable paths by mollifying the interactions, as in Sect. 4, and using the decreasing mollification bandwidth  $r_n$  at mutation  $n$ . The target distribution (altered by the shrinkage rate) should change slower than the Markov chain mixing rate. The shrinkage rate of the mollification bandwidth for a single Markov chain

$$O(\gamma^n) \leq r_n < O(1), \quad (7)$$

ensures that the change in the target distribution for mutation  $n$  is small enough (see App. B). We use an ad-hoc sequence  $r_n = r_0 \gamma^n$ , where the constant  $r_0$  should be appropriately selected by the user to be large enough to allow the Markov chain to easily reach any of the mollified paths from any state during the first steps. Also note that the parameter  $\gamma$  is problem-dependent and has a different meaning than in Monte Carlo integration:  $\gamma \in (0; 1)$  should be more than the spectral gap of the transition kernel  $K(\cdot, \cdot)$ . It has to be chosen close enough to 1 to guarantee the practical convergence.

**Normalization Constant.** The initialization of MCMC methods is done with unbiased methods to estimate the image normalization constant correctly. Veach’s original normalization [Vea98] also works in our case, as it is based on a

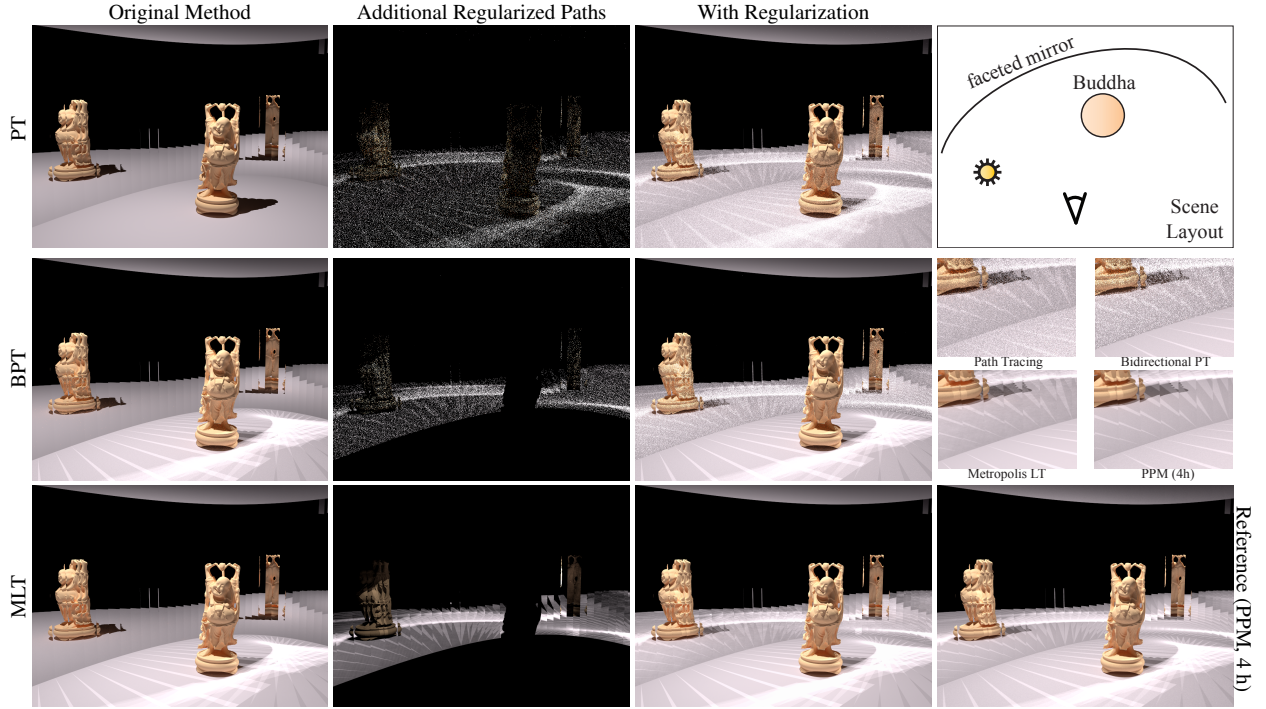


Figure 4: Equal time rendering (30 min.) of the BUDDHA scene at 1024×768. This simple scene illustrates the additional light transport paths enabled by the regularization. It is “simple enough” such that it can be handled with progressive photon mapping which we used to compute the “reference” (4 hours). Please zoom in the electronic version.

scaling ratio to the alternative sampling method (BPT in our case). However, there might be cases when regularized features dominate the resulting image and thus it is hard, or even impossible, to determine the normalization constant solely based on paths that can be sampled in an unbiased way. In order to avoid such cases, we propose to initialize MLT with a selectively regularized integrand (using regularized BPT). We then refine the normalization constant sparsely, e.g. once per several hundred mutations, by sampling some additional paths using regularized BPT. During this process both mollification bandwidths for MLT and BPT are shrunk as described in Sect. 5.1 and 5.2. This way the normalization constant is estimated in a consistent way and practically converges earlier than the actual image (see Fig. 5).

## 6. Implementation Details

For our experiments and comparisons we use the following rendering methods: standard path tracing (PT) [Kaj86] with next event estimation; bidirectional path tracing (BPT) [VG94, LW93] with multiple importance sampling [VG94] (power heuristic,  $\beta = 2$ ); Metropolis light transport (MLT) with the original set of mutations [Vea98] as well as with simple mutations in the unit hypercube and bidirectional connections [KSKAC02]; stochastic progressive photon mapping (PPM) [HOJ08, HJ09]; and its vari-

ant with Metropolis photon tracing (MCPPM) [HJ11]; and recent vertex connection and merging (VCM) [GKDS12]. Note that our method is orthogonal to the sampling strategy, i.e. one could use the original mutations proposed by Veach [Vea98] as well as the recent additional mutation through specular manifolds [JM12].

Our renderer is implemented using the Nvidia OptiX ray tracing platform and all results have been measured using a GeForce 580 GTX GPU. Unless stated otherwise, all images have been rendered at 1024 × 768 resolution.

In Alg. 1 we provide a simple listing for Monte-Carlo integration, which can replace the evaluation routine for the specular material in an existing unbiased renderer. Note that

**Algorithm 1** Selective Angular Mollification of a Single Specular Interaction for Bidirectional Path Tracing; reflection direction  $\omega$  and direction to vertex  $\omega'$ .

```

if current vertex is "S" then
  if path not contains "DD" then
     $r_N \leftarrow r_0 \cdot N^{-\lambda}$   $\triangleright$  Shrink bandwidth with iterations
     $\epsilon_N \leftarrow \arctan \frac{r_N}{\text{connection\_distance}}$   $\triangleright$  Compute angle
    return  $\phi_{\epsilon_N}(\omega, \omega') / \langle \omega \cdot n \rangle$   $\triangleright$  Evaluate mollifier
  else return 0  $\triangleright$  Mollification is not required

```

the same code can be reused for MCMC integration with the difference in shrinkage evaluation of  $r_n$ . Special attention should be paid to the machine precision during the integration. Since the mollification angle can become very narrow after many iterations, it can run out of floating point precision. We suggest to avoid dividing by the connection distance in Alg. 1, but to work on a scaled sphere, which contains the vertex to be connected to. In the additional material we provide a modified version of 'SmallPT', a minimalistic path tracer by Kevin Beason, demonstrating the few changes our method requires.

**Parameters.** Both  $r_0$  and  $\lambda$  (for MC integration) /  $\gamma$  (for MCMC integration) are user-defined parameters;  $r_0$  controls the mollification blurring in spatial units. In the case of MC integration  $\lambda$  controls the shrinkage rate and is equivalent to this in PPM [HPJ12, KD13]. For MCMC, in contrast,  $\gamma$  is the second largest eigenvalue of the transition kernel and denotes the mixing rate of the Markov chain. It depends on the scene and cannot be known in advance. It is usually hard or impossible to obtain from a given transition kernel, thus we suggest to select it empirically close enough to 1 in order to achieve the practical convergence. We use the value of 1% of the scene bounding radius for the initial on-surface bandwidth  $r_0$ ,  $\lambda = 1/6$  for MC integration with 2D regularization and  $\gamma = 1 - 10^{-4}$  for MCMC integration in all experiments.

## 7. Results and Discussion

Fig. 4 shows PT, BPT and MLT with and without regularization to demonstrate the difference when accounting for all light transport paths (which in this case can also be handled with progressive photon mapping).

Fig. 5 shows a comparison of MLT with regularization and Markov chain PPM (MCPMM) demonstrating the convergence of our method. Note that MCPMM as well as the original MLT method does not handle all light transport paths (see also Fig. 1). The bias introduced by the regularization is apparent as blurring of the floor caustics reflected in the wall; due to a spatially-controlled mollification it is similar to the bias of PPM. However, the majority of lighting in regularized MLT is rendered in an unbiased way, which can be seen by comparing the sharpness of the edge of the directly visible caustic on the floor. Also note that the bright splotches in both regularized and original MLT images are due to Kelemen et al.'s simple mutation strategy [KSKAC02]; the original mutation set with the recent manifold exploration (the last two images in the bottom row) reduce these artifacts.

Fig. 6 demonstrates the strength of our method in finding difficult paths. Due to the specular BRDF and the pinhole camera, unbiased methods typically have difficulties handling the reflection of the LEDs on the curved surface.

**Future Work and Limitations.** An interesting study for future work would be an analysis of the spectral gap of the




Figure 6: The EG logo is made of 4194 LEDs and placed over a curved specular water surface with a subtle diffuse component. The image was rendered using MLT with regularization in 10 minutes at a resolution of  $256 \times 256$  pixels.

transition kernel of MCMC methods, since all marginal distributions are known. This would allow us to assess the lower bound of the convergence rate as well as the more accurate choice of the bandwidth shrinkage rate parameter  $\gamma$ .

One limitation of our method can be observed with energy redistribution path tracing (ERPT) [CTE05]. It can also benefit from regularization, but due to the design of the algorithm the bias can be high: ERPT computes a predetermined number of mutations for many short chains and thus the mollification bandwidth has to be reset to  $r_0$  for every new chain.

**Other Consistent Methods as a Special Case.** Photon mapping methods can be considered as a special case of regularization. These methods do not connect eye-subpaths with randomly shot light-subpaths (as BPT), but one more ray is shot from the eye-subpath according to the local importance function and then a range query gathers from proximate light-subpaths at the intersection point. This can be interpreted as angular regularization after sampling, leading to the same regularization procedure that we propose for specular BRDFs. However, this regularization is *applied to all paths* regardless whether they can be sampled or not.

On the other hand, PPM better samples certain types of illumination than BPT. This is due to the efficient caching of light-subpaths using range queries which connect an eye-subpath to multiple light-subpaths (stored in a photon map) [GKDS12]; in contrast, BPT connects to random instead of relevant paths. However, our method can also be combined with more advanced sampling methods such as






Figure 5: Equal time comparison: a point light illuminates the KITCHEN scene from outside through glass windows. Note how MCPPM introduces bias where MLT does not: here bias becomes noticeable as blurred edges of the floor caustic (red inset). The right inset shows regions where both MCPPM and regularized MLT introduce bias; original MLT cannot handle the caustic reflected off the mirror wall. For MCPPM we update eye subpaths once per 2000 mutations of the Markov chain. We use 16384 Markov chains mutating in parallel for both MLT and MCPPM. VCM [GKDS12], one of the most advanced non-Metropolis techniques, generates noisy results which is due to the inefficiency of random path sampling (compared to exploring the small subset of light paths through the window contributing to the lighting in the image). The two rightmost figures in the bottom row were rendered using original Metropolis light transport [Vea98] with the recent manifold exploration mutation [JM12] using the Mitsuba renderer [Jak10]. We use roughly the same number of mutations per pixel as for other images rendered with our renderer in 3 hours. The slight difference to the other images is due to the different material descriptions used. Note how manifold exploration is able to efficiently explore regularized paths.

MLT, which has a superior convergence than PPM in non-trivial illumination scenarios (Fig. 5) and does not suffer from the absence of the cache as BPT (Fig. 4).

Virtual spherical lights (VSL) [HKWB09] are used to replace virtual point lights for faster rendering of glossy scenes. This is equivalent to an angular mollification of specular BRDFs (at the connection point) and convergence is also achieved by gradually shrinking the VSL radii (as in Eq. 6).

**Conclusions.** We have shown that selective regularization in path space using mollification enables us to overcome the ill-posedness of the light transport integrand, orthogonally to the integration method used. The regularization is

only required for a small subset of the path space, avoiding bias for the majority of light transport, while handling difficult paths in a consistent way. This provides a simple yet powerful framework, which incorporates existing consistent methods as special cases. It enables consistent sampling of paths, hitherto non-samplable by advanced techniques, such as Metropolis light transport.

**Acknowledgements.** This work has been partially funded by the Intel Visual Computing Institute, Saarbrücken. The KITCHEN and the BATHROOM scenes are the part of Evermotion Archinteriors package. We thank the anonymous reviewers for helpful suggestions and improvements.

## References

- [CTE05] CLINE D., TALBOT J., EGBERT P. K.: Energy redistribution path tracing. *ACM Transactions on Graphics* 24, 3 (2005), 1186–1195. 8
- [CWY11] CHEN J., WANG B., YONG J.-H.: Improved stochastic progressive photon mapping with metropolis sampling. *Computer Graphics Forum* 30, 4 (2011), 1205–1213. 2
- [FCL05] FAN S., CHENNEY S., LAI Y.-C.: Metropolis photon sampling with optional user guidance. In *Proceedings of the 16th Eurographics Symposium on Rendering* (2005), pp. 127–138. 2
- [GKDS12] GEORGIEV I., KŘIVÁNEK J., DAVIDOVIČ T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* 31, 6 (2012), 192:1–192:10. 2, 5, 7, 8, 9
- [HA96] HEGLAND M., ANDERSSON R.: A mollification framework for improperly posed problems. *Numer. Math* 78 (1996), 549–575. 3, 5
- [Hec90] HECKBERT P. S.: Adaptive radiosity textures for bidirectional ray tracing. *Computer Graphics (Proc. of SIGGRAPH)* 24, 4 (1990), 145–154. 3
- [HJ09] HACHISUKA T., JENSEN H. W.: Stochastic progressive photon mapping. *ACM Transactions on Graphics* 28, 5 (2009), 141:1–141:8. 2, 3, 7
- [HJ11] HACHISUKA T., JENSEN H. W.: Robust adaptive photon tracing using photon path visibility. *ACM Transactions on Graphics* 30, 5 (2011), 114:1–114:11. 2, 7
- [HKWB09] HAŠAN M., KŘIVÁNEK J., WALTER B., BALÁ K.: Virtual spherical lights for many-light rendering of glossy scenes. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* 28, 5 (2009), 143:1–143:6. 9
- [HOJ08] HACHISUKA T., OGAKI S., JENSEN H. W.: Progressive photon mapping. *ACM Transactions on Graphics* 27, 5 (2008), 130:1–130:8. 2, 3, 6, 7
- [HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* 31, 6 (2012), 191:1–191:10. 2, 8
- [Jak10] JAKOB W.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. 9
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), SIGGRAPH '98, pp. 311–320. 4
- [JM12] JAKOB W., MARSCHNER S.: Manifold exploration: a markov chain monte carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 31, 4 (2012), 58:1–58:13. 2, 3, 5, 7, 9
- [Kaj86] KAJIYA J. T.: The rendering equation. *Computer Graphics (Proc. of SIGGRAPH)* (1986), 143–150. 1, 2, 7
- [KD13] KAPLANYAN A. S., DACHSBACHER C.: Adaptive progressive photon mapping. *ACM Transactions on Graphics* 32, 2 (2013). 2, 6, 8
- [KHM09] KNIPE S., HÄRING S., MAGNOR M.: Efficient and accurate rendering of complex light sources. *Computer Graphics Forum* 28, 4 (2009), 1073–1081. 2
- [KSKAC02] KELEMEN C., SZIRMAY-KALOS L., ANTAL G., CSOKA F.: A simple and robust mutation strategy for the metropolis light transport algorithm. *Computer Graphics Forum* 21, 3 (2002), 531–540. 7, 8
- [LW93] LAFORTUNE E. P., WILLEMS Y. D.: Bi-directional path tracing. In *Proc. of Conference on Computational Graphics and Visualization Techniques* (1993), pp. 145–153. 7
- [RTY94] REIF J. H., TYGAR J. D., YOSHIDA A.: Computability and complexity of ray tracing. *Discrete Computational Geometry* 11, 1 (1994), 265–288. 4
- [Vea98] VEACH E.: *Robust monte carlo methods for light transport simulation*. PhD thesis, Stanford University, 1998. AAI9837162. 1, 2, 3, 4, 6, 7, 9
- [VG94] VEACH E., GUIBAS L.: Bidirectional estimators for light transport. In *Proc. of Eurographics Rendering Workshop* (1994), pp. 147–162. 2, 5, 7
- [Vor11] VORBA J.: Bidirectional photon mapping. In *15th Central European Seminar on Computer Graphics* (2011). 2

## Appendix A: Consistency of Regularized Monte-Carlo

If we consider only a specular BRDF regularization as proposed in Sect. 4.2, one could claim equivalence to progressive photon mapping and conclude the consistency. Here we show that the regularization framework, including arbitrary (e.g. spatial) mollifications, is also consistent under general conditions. Due to size constraints we provide all proofs in the supplementary material.

**Theorem 2.** *Given the unified space of all light transport paths  $\Omega$ , the Monte-Carlo estimation of an integral as in Eq. 1 with a measurement function  $f(\bar{x})$  as an integrand, converges consistently with integrand  $f_r(\bar{x})$  with  $d$ -dimensional selective regularization (as described in Sect. 4), i.e.*

$$\mathbb{E}[\hat{f}] = \mathbb{E}\left[\frac{1}{N} \sum_{n=1}^N \frac{f_{r_n}(\bar{x}_n)}{p(\bar{x}_n)}\right] = I$$

*if and only if the sequence of mollification bandwidths  $\{r_n\}$  decreases as  $O(n^{-1/d}) < r_n < O(1)$ .*

## Appendix B: Consistency of Regularized MCMC

The following theorem states that given an ergodic mutation strategy, selective regularization leads to a consistent estimation of the integral in Eq. 1.

**Theorem 3.** *Given the unified space of all light transport paths  $\Omega$  as a state space, and a Harris-recurrent transition kernel  $K(\cdot, \cdot)$  on it, the Markov chain Monte-Carlo estimation of an integral in Eq. 1 with selective regularization of the integrand  $f_r(\bar{x})$  (as in Sect. 4), converges almost surely if bandwidth shrinkage rate is  $O(\gamma^n) \leq r_n < O(1)$  for some  $\gamma \in (0, 1)$ .*

Note that the conditions of the theorem assert that the mutation strategy should not only cover the whole path space, but the probability of returning to the same path has to be non-zero. All sets of existing strategies, including Kelemen's and Veach's mutation sets, comply.