

Wave Tracing: Generalizing The Path Integral To Wave Optics

Shlomi Steinberg¹  and Matt Pharr² 

¹University of Waterloo, Canada

²NVIDIA, United States




Figure 1: In this paper we discuss the limitations of classical, ray-based path tracing, and rigorously study what optical phenomena that light transport formalism may reproduce. We show that some effects, like dispersion and scattering by a restricted class of statistical surface models (rendered in image **A**), fall under its regime. Then, we generalize the classical path integral to a formulation that is able to account for a much wider class of wave effects, and based on that present a unified framework that is able to: (**B**) simulate long-wave radiation and its propagation and diffraction in complex environments, for example to compute its signal strength (visualized color-coded); and (**C**) render optical wave effects, such as diffraction by arbitrary geometry.

Abstract

Modeling the wave nature of light and the propagation and diffraction of electromagnetic fields is crucial for the accurate simulation of many phenomena, yet wave simulations are significantly more computationally complex than classical ray-based models. In this work, we start by analyzing the classical path integral formulation of light transport and rigorously study which wave-optical phenomena can be reproduced by it. We then introduce a bilinear path integral generalization for wave-optical light transport that models the wave interference between paths. This formulation subsumes many existing methods that rely on shooting-bouncing rays or UTD-based diffractions, and serves to give insight into the challenges of such approaches and the difficulty of sampling good paths in a bilinear setting.

With this foundation, we develop a weakly-local path integral based on region-to-region transport using elliptical cones that allows sampling individual paths that still model wave effects accurately. As with the classic path integral form of the light transport equation, our path integral makes it possible to derive a variety of practical transport algorithms. We present a complete system for wave tracing with elliptical cones, with applications in light transport for rendering and efficient simulation of long-wavelength radiation propagation and diffraction in complex environments.

CCS Concepts

• **Computing methodologies** → **Rendering; Computer graphics; Scientific visualization;** • **Applied computing** → **Physics;**

1. Introduction

Ray optics has been the foundation of light transport algorithms in computer graphics since Whitted's introduction of recursive ray tracing. Especially with the development of path tracing, geometric optics light transport has become progressively more sophisticated and has come to be able to model a wide variety of scattering and illumination effects. Researchers have long sought to extend

such ray-based algorithms to model wave effects while maintaining the efficiency of ray models. One area of focus has been scattering models that account for wave effects when light interacts with surfaces or geometric edges; another has been the development of algorithms to simulate wave-based light transport. These areas are complementary in that the light transport depends on accurate scattering and diffraction models while scattering models require information that is not available in a classical ray-based model: for

example, the shape of the electromagnetic wavefronts or optical coherence. In this work we focus on the light transport problem.

Wave-based models are more computationally expensive than ray-based models, and so we begin by studying a fundamental question: precisely which wave effects can ray-based light transport accurately simulate, and which can they not? Under a few assumptions that are characteristic of rendering applications (e.g., sensors that average measurements over long times and a far-field setting), we rigorously characterize the conditions under which electromagnetic fields obey the Eikonal equation point-wise, and thus, the fields' wavefronts can be propagated using rays. We further discuss and formally show what classes of materials can be simulated with such light transport models.

With the boundaries of models based on independent rays established, we turn to determining what wave-based models may achieve. A number of wave simulation methods have been developed that use geometric rays as the underlying transport primitive, so we start by deriving a *bilinear path integral* that characterizes light transport with interfering ray-based paths. This work is inspired by Veach's introduction of the path integral formulation of the light transport equation that transformed the traditional recursive rendering equation into an infinite-dimensional integral over all light carrying paths in a scene [Vea97].

Our bilinear path integral spans the cases of pairs of paths that do not interact—corresponding to standard path tracing; pairs of *phase-carrying paths* that always interfere—corresponding to several wave-based light transport models (see Section 2); as well as pairs of paths that are partially-coherent, including approaches like those based on propagating the Wigner distribution function or the mutual coherence function. This model clearly shows the difficulties of path sampling with ray-based paths: although each path individually may have a meaningful contribution, the interaction of a pair of paths may lead to destructive interference and thus, no light transport after all. Therefore, the efficiency of *local path sampling* is greatly reduced, and path sampling effectively becomes a *global problem*, where all path pairs need to be considered. While these issues are well-known, our path integral makes the associated sampling challenges clear.

We then turn to develop a *weakly-local path integral* that is based on region-to-region transport. This model is inspired by *physical light transport* (PLT) [SSY22; SRB*24b], which introduced *weakly-local*—meaning confined to a small spatial region—light transport primitives. Our formulations generalize their wave-based light transport equation and maintain the crucial characteristics of modeling transport over small regions of space such that only observable wave effects are accounted for. As transport is done region-to-region (and not point-to-point), no interference between pairs of paths may arise, making it possible to efficiently construct individual paths via local sampling. Our formulations further allow the derivation of advanced light transport algorithms for wave models, including a bidirectional model.

We show that elliptical cones form tight geometric envelopes for PLT's light transport primitives. Then, based upon the weakly-local path integral formulation, we develop a practical, general-purpose wave transport model. At its core, our wave transport algorithm can be understood as replacing the classical ray with an elliptical cone.

We discuss how to traverse the scene with elliptical cones, simulate the interaction of their underlying electromagnetic fields with the geometry and materials that fall within an interaction region, and introduce a general importance sampling strategy. We further show that not only is this algorithm able to reproduce accurate diffraction effects, but it can also do accurate wave simulations for applications in different parts of the electromagnetic spectrum, see Fig. 1.

2. Related Work

Our work builds on two classic foundations of ray-based light transport. First is the path integral formulation of the light transport equation [VG97; Vea97]. This expression of the rendering equation made it clear how to apply non-local sampling techniques to light transport (i.e., that paths do not necessarily need to be sampled incrementally from the sensor) and has provided a rigorous foundation for deriving more advanced light transport simulation algorithms. We also build upon the operator expression of light transport developed by ARVO [Arv95] and VEACH [Vea97]. By expressing light transport and scattering as linear operators on functions describing light emission and sensor response, this formulation allows compact expression of light transport and has enabled analysis of its convergence.

Starting with Stam's pioneering work on modeling diffraction from rough metallic surfaces [Sta99], there has been considerable work on BSDFs that model wave effects. Examples include diffraction due to scratches [WVJH17; VWH18], iridescent and pearlescent materials [GMC*20], thin-film interference [HIK*20], and dispersion via thin dielectric layers above conductors [BB17; KKG19]. Scattering models have also been developed based on measured scattering from diffractive surfaces [TG17], diffractive surfaces modeled with explicit microgeometry [FJF20; YXW*23], and statistical models [HP17; Kry06; SY22]. Also related is recent work on modeling free space diffraction as a BSDF [SRB*24a].

Many techniques have been developed to simulate wave-based light transport. Approaches based on the Wigner distribution function have been developed both in graphics [CHB*12] and in optics [JB91; MWB*18; MJ21]. Outside of graphics, algorithms based on different shooting-bouncing rays methodologies [Wei06] or the geometric/uniform theory of diffraction (UTD) [SM99; YHM*22; BVD20] are often used. All these approaches rely on ray tracing for transport, and share a fundamental difference compared with the more common ray or path tracing seen in computer graphics: rays *mutually interfere* with each other. In this work, we generalize the classical path integral to such *bilinear* transport; this generalization highlights the additional sampling difficulties that arise with mutually-interfering rays, and in the future may serve as a theoretical foundation for developing better path sampling techniques for such bilinear light transport models.

Recently, physical light transport [SY21; SSY22; SRB*24b] has introduced *weakly-local* light transport primitives, which can be understood as tracing beams of light instead of singular rays—enabling the reproduction of a much wider class of wave effects, compared with classical light transport. Previous implementations of it have approximated the transport over regions with singular rays. We develop the underlying path tracing theory by generalizing the classical path integral to *region-to-region* transport, and

present a solver that models wave transport over volumes using elliptical cones.

We employ the wide bounding volume hierarchy (BVH) designed for cone traversal proposed by EMRE et al. [EKS25], as well as their intersection routines, for our method. Cone queries have also been used for differential visibility [WMB*25].

3. Generalizing the Path Integral Formulation

The classic path integral formulation of light transport [Vea97] is given by:

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

where $\bar{\mathbf{x}} \in \Omega$ is a scene path—connecting an emitter to the sensor over a finite sequence of points—in the space of all paths Ω , and f is the real-valued, non-negative *measurement contribution function*. The result of f is the non-negative power measured by the sensor.

It is understood that light transport under the formulation above is governed by the laws of geometric optics. Nevertheless, some optical effects beyond geometric optics may be reproduced accurately by it. It is of interest then to ask the question—*to what extent is classical path tracing able to reproduce wave-optical phenomena?* In our supplemental material, we study the formal conditions under which classical path tracing is able to simulate wave-optical phenomena, and discuss which optical effects can be reproduced.

We also show that some effects may not be reproduced classically. As a motivating example, consider a ray passing close to

geometry, but not intersecting it. Under wave optics, part of the energy carried by the ray should have diffracted around the geometry. However, we may not detect this situation with classical ray tracing (a “does this ray travel close to geometry” query would be needed), so the energy that is propagated straight on is overestimated. To amend the error, one may consider two approaches: we could (i) allow other paths to fix these erroneous contributions, by facilitating interference between different paths, as is done with ray-based techniques such as UTD; or, (ii) replace ray queries with a volumetric query, in order to sample *regions*—and not singular points—where wave-optical interactions occur. In this section we generalize the path integral to both of these approaches.

3.1. Bilinear Path Integral

We extend the classical path integral with a map $F : \Omega \times \Omega \rightarrow \mathbb{R}$, the *mutual contribution function*, producing a *bilinear path integral*:

$$I = \int_{\Omega \times \Omega} F(\bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mu(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{y}}) . \quad (2)$$

We only consider maps F that fulfill the following properties:

1. symmetry, $F(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = F(\bar{\mathbf{y}}, \bar{\mathbf{x}})$;
2. non-negative contribution over a single path, $F(\bar{\mathbf{x}}, \bar{\mathbf{x}}) \geq 0$;
3. Cauchy–Schwarz inequality, $|F(\bar{\mathbf{x}}, \bar{\mathbf{y}})|^2 \leq F(\bar{\mathbf{x}}, \bar{\mathbf{x}})F(\bar{\mathbf{y}}, \bar{\mathbf{y}})$.

Note, when $\bar{\mathbf{x}} \neq \bar{\mathbf{y}}$, the map $F(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ may take negative values. It is easy to see that the properties above mandate that the path integral in Eq. (2) integrates to a non-negative I , as desired (see Section S2 in the supplemental for a proof).

As a statistical distribution To see how the mutual contribution function F connects to optics, let ψ be some realization of a statistical wave ensemble (i.e., ψ is some electric field). Given a geometric path $\bar{\mathbf{x}}$, we will let $\psi(\bar{\mathbf{x}})$ denote the field strength that is transported over that path. (We briefly mention a few formalisms that enable such field transport over a path later). Then, the measurement in expectation becomes:

$$I = \left\langle \left| \int_{\Omega} \psi(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) \right|^2 \right\rangle = \int_{\Omega \times \Omega} \langle \psi(\bar{\mathbf{x}}) \psi^*(\bar{\mathbf{y}}) \rangle d\mu(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{y}}) , \quad (3)$$

where $\langle \cdot \rangle$ denotes ensemble averaging and ψ^* is the complex conjugate of ψ . We formally interchanged the order of ensemble averaging and integration. The above takes an identical form to Eq. (2), with $F(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = \langle \psi(\bar{\mathbf{x}}) \psi^*(\bar{\mathbf{y}}) \rangle$ quantifying the ensemble’s second-order statistics (its cross-correlation), and is the *mutual coherence function* in the optical context.

The statistical formulation above may also account for temporal fluctuations, e.g., when F operates on paths with different wavelengths (under spectral rendering). This arises in practice in simulations where sensors do not time average (or averaging times are too short to induce temporal decoherence), for example for acoustic simulations with observable interference across spectral samples. For ergodic ensembles ψ , the ensemble averages can indeed be understood as time averages, though we keep the formulation general.

SYMBOL		UNITS
<i>paths and their contributions</i>		
$\bar{\mathbf{x}} = \bar{\mathbf{x}}_1 \dots \bar{\mathbf{x}}_n$	a <i>path</i> : a finite sequence of n spatial positions $\bar{\mathbf{x}}_j \in \mathbb{R}^3$	
$\mathbf{R}$	a <i>region</i> : a bounded, closed subset $\mathbf{R}_j \subset \mathbb{R}^3$	
$\bar{\mathbf{R}} = \mathbf{R}_1 \dots \mathbf{R}_n$	a <i>weakly-local path</i> : a finite sequence of n regions (regions may overlap)	
$f(\bar{\mathbf{x}}), g(\bar{\mathbf{R}})$	the <i>measurement contribution functions</i> over a path	w
$F(\bar{\mathbf{x}}, \bar{\mathbf{y}})$	the <i>mutual contribution function</i> , given over a pair of paths $\bar{\mathbf{x}}, \bar{\mathbf{y}}$	w
<i>electrodynamics</i>		
λ	wavelength of electromagnetic radiation	mm
$k = \frac{2\pi}{\lambda}$	wavenumber of electromagnetic radiation	mm ^{−1}
$\vec{k}$	wavevector	mm ^{−1}
ψ	wavefunction (scalar electric field)	v m ^{−1}
<i>light transport and Wigner (space-frequency) distributions</i>		
$\phi(\vec{r}, \vec{k})$	the Wigner distribution of a <i>Gaussian beam</i>	W or 1
$L(\vec{r}, \vec{k})$	Wigner distribution of an emitter’s emission	w
$W(\vec{r}, \vec{k})$	Wigner distribution of a sensor’s sensitivity	1
$T, \mathcal{T}$	context-dependant interaction or scattering operators acting upon Wigner distributions	

Table 1: List of symbols and notation in this paper.




Figure 2: Path integrals for wave optics. We develop two generalizations of the classical light transport path integral [Vea97]—capable of reproducing a wide range of wave effects. (A) The bilinear path integral, [Subsection 3.1](#) (implicitly employed by common methods that use phase-carrying rays, like those that rely on UTD-based diffractions) accounts for the interference between all pairs of paths. Bilinearity frustrates local path sampling: assume that a path was sampled; afterwards, another path might be sampled that destructively interferes with the former path, resulting in overall negligible contribution, even with good local path sampling. (B) The weakly-local path integral, [Subsection 3.2](#), formalizes region-to-region transport. In-place of rays, we employ elliptical cones for transport, and formulate the aggregated interaction of light with all materials that fall within a region (illustrated in orange). This requires more sophisticated machinery—tracing of elliptical cones and aggregated wave-optical interactions—but allows crucial local path sampling. Ballistic segments are illustrated using thick lines, and an elliptical conic frustum within which intersected geometry is found in orange (see [Subsection 4.2](#)).

Examples In the limiting case where the mutual contribution function is proportional to a Dirac delta, $F(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = |\psi(\bar{\mathbf{x}})|^2 \delta(\bar{\mathbf{x}}, \bar{\mathbf{y}})$, the above reduces to the classical unitary case, with $f = |\psi(\bar{\mathbf{x}})|^2$ being the classical measurement contribution function. The Dirac delta formalizes the fact that all paths are assumed to be perfectly mutually uncorrelated—*incoherent*—and no interference arises. Therefore, the bilinear path integral generalizes the classical path integral.

Another important special case is path tracing with *phase-carrying rays*, where every pair of paths are fully coherent in their measurement and interfere. Let the map F take the form:

$$F(\bar{\mathbf{x}}, \bar{\mathbf{y}}) = |\psi(\bar{\mathbf{x}})| |\psi(\bar{\mathbf{y}})| \operatorname{Re} e^{-ik(|\bar{\mathbf{x}}| - |\bar{\mathbf{y}}|)}, \quad (4)$$

where $|\bar{\mathbf{x}}|$ is the path’s total distance (in world units), and $k = \frac{2\pi}{\lambda}$ is the wavenumber with wavelength λ . Clearly F fulfills the desired properties of the mutual contribution function. The terms $k|\bar{\mathbf{x}}|$ are then the phases carried by each ray, and the complex exponent term quantifies the mutual interference between every pair of paths. This formalism of phase-carrying rays encompasses several common light transport methods, including UTD-based simulations—very common in acoustics and RF simulations—and other methods that rely on mutually-interfering phase-carrying rays.

In between the above two extremes—perfect incoherence and perfect coherence—partially-coherent formalisms arise. Such formalisms include light transport formalisms that employ the Wigner distribution function or the mutual coherence function directly.

Local path sampling Local path sampling [Kaj86] is a very common technique to construct complete paths $\bar{\mathbf{x}}$, where each point in the path is recursively sampled *locally* by sampling the local interaction function and then ray tracing. Practical light transport often relies on the ability to perform good local path sampling.

With bilinear light transport, local path sampling is hampered by the fact that the mutual contribution $F(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ usually cannot be evaluated without full knowledge of the paths $\bar{\mathbf{x}}, \bar{\mathbf{y}}$, and may annihilate the paths’ contributions. That is

$$\langle |\psi(\bar{\mathbf{x}}) + \psi(\bar{\mathbf{y}})|^2 \rangle = \langle |\psi(\bar{\mathbf{x}})|^2 \rangle + \langle |\psi(\bar{\mathbf{y}})|^2 \rangle + 2F(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \quad (5)$$

may vanish entirely. That is, even if a path $\bar{\mathbf{x}}$ was constructed

with effective local path sampling, another path $\bar{\mathbf{y}}$ might annihilate its contribution (illustrated in [Fig. 2a](#)). Under the bilinear setting, *path sampling always becomes a global problem*. The classical path integral, [Eq. \(1\)](#), does not suffer from this issue, as its integrand is always non-negative. It is also noteworthy that constructive/destructive interference happens at a high frequency across the scene, making techniques like path guiding ineffective. The difficulty with local path construction is a fundamental challenge in path tracing simulations that employ typical UTD, or other phase-carrying rays, formalisms.

Other forms of non-local path construction, like Metropolis light transport, could be applied to the bilinear path integral. We leave this to be explored in future work.

3.2. Weakly-Local Path Integral

The classical path integral, [Eq. \(1\)](#), formalizes point-to-point light transport. We now generalize that formulation to region-to-region light transport (illustrated in [Fig. 2b](#)):

$$I = \int_{\Omega} g(\bar{\mathbf{R}}) d\mu(\bar{\mathbf{R}}), \quad (6)$$

where $\bar{\mathbf{R}} = \mathbf{R}_0 \mathbf{R}_1 \dots \mathbf{R}_n$ is a generalized path defined as sequence of bounded regions $\mathbf{R}_j \subset \mathbb{R}^3$, Ω is now the set of all such sequences of all bounded subsets of space, and μ is then the appropriate product measure: $\mu(\bar{\mathbf{R}}) = \mu(\mathbf{R}_0) \dots \mu(\mathbf{R}_n)$. g is the path contribution function, which can be written in operator notation:

$$g(\bar{\mathbf{R}}) = \left\langle W, T_{\mathbf{R}_{n-2} \rightarrow \mathbf{R}_{n-1} \rightarrow \mathbf{R}_n} T_{\mathbf{R}_{n-3} \rightarrow \mathbf{R}_{n-2} \rightarrow \mathbf{R}_{n-1}} \dots T_{\mathbf{R}_0 \rightarrow \mathbf{R}_1 \rightarrow \mathbf{R}_2} L_{\mathbf{R}_0 \rightarrow \mathbf{R}_1}(\mathbf{R}_0) \right\rangle. \quad (7)$$

T are generic transport operators, generalizing the classical BSDF to region-to-region transport. The emission function L sources a distribution of light from a given initial region $\mathbf{R}_0$:

$$\phi_1 = L_{\mathbf{R}_0 \rightarrow \mathbf{R}_1}(\mathbf{R}_0). \quad (8)$$

The transport operators then transform these distributions, as follows, thereby propagating them from region to region:

$$\phi_{j+1} = T_{\mathbf{R}_{j-1} \rightarrow \mathbf{R}_j \rightarrow \mathbf{R}_{j+1}} \phi_j. \quad (9)$$

Finally, the inner product of the sensor sensitivity function W over the distribution of light quantifies the sensor's response to a distribution and yields the measurement contribution of the path $\bar{\mathbf{R}}$:

$$g(\bar{\mathbf{R}}) = \langle W, \phi_n \rangle, \quad (10)$$

which must be real and non negative, $g(\bar{\mathbf{R}}) \geq 0$. We discuss the above further and prove the non-negativity of Eq. (10) in our supplemental material. $\langle \cdot, \cdot \rangle$ denotes the inner functional product (not to be confused with the ensemble average).

Our weakly-local formulation of the path integral makes explicit the regions $\mathbf{R}_j$ over which light transport takes place. It is easy to see that Eq. (6) generalizes the classical formulation: when the regions are reduced to points, T to classical BSDFs, and ϕ_j to radiance point samples, the above reduces to the classical point-to-point transport. However, Eq. (6) is also more powerful than the classical formulation: we now have a built-in mechanism to quantify the shapes of the wavefronts of the underlying electromagnetic field (via ϕ_j), and to quantify the regions over which weakly-local interactions of light with matter may take place (via $\mathbf{R}_j$).

We keep the definition of the distributions ϕ abstract to make our formulations in this Section applicable to different light transport formalisms: for example, ϕ may be correlation functions, quantifying a wave ensemble's second order statistics, or realizations from a wave ensemble. In the following section, we restrict the distributions ϕ to the Gaussian space-frequency (Wigner) distributions that are used by PLT, and then formalize our weakly-local path tracing framework.

Local path sampling Unlike the bilinear generalization in Subsection 3.1, the contribution g is crucially always non negative—no interference across samples may arise. Local path sampling may then proceed in a manner essentially identically to the classical case: sample an initial region $\mathbf{R}_0$, and then subsequent regions, using the emission function L and the transport operators T , respectively.

4. Wave Tracing

The ability to construct paths via local path sampling with our weakly-local path integral formulation, discussed in Subsection 3.2, is a major advantage over its bilinear counterpart. It comes at the cost of increased analytic and algorithmic complexities of the simulation: light transport is now formulated in terms of regions and not points, and our intersection queries are no longer ray-geometry intersections but volumetric queries. In this section we address these difficulties and present a *wave tracing* framework: a method for general weakly-local light transport targeting wave-optics simulations.

4.1. Background: Physical Light Transport

We build upon the physical light transport (PLT) framework for wave-optics simulation [SRB*24b; SSY22]. PLT works with Gaussian wave functions, which are the most compact physically-realizable wave functions, and generalize many of the useful properties of the classical ray of geometric optics to wave optics. For these reasons, previous work refers to such wave functions as *generalized rays*. However, these are not *rays*, per se, i.e. they are not

infinitesimal rays with zero cross section, therefore, to make the distinction from rays explicit, we prefer to refer to such wave functions simply as *Gaussian beams*.

The space-frequency (Wigner) distribution of a Gaussian beam is

$$\begin{aligned} \phi(\vec{\mathbf{r}}, \vec{\mathbf{k}}) = & \frac{1}{\pi^3} \exp \left[-(\vec{\mathbf{r}} - \vec{\mathbf{r}}_0)^\top \Sigma^{-1} (\vec{\mathbf{r}} - \vec{\mathbf{r}}_0) \right] \\ & \times \exp \left[-(\vec{\mathbf{k}} - \vec{\mathbf{k}}_0)^\top (\rho^2 \Sigma^{-1} + \Sigma) (\vec{\mathbf{k}} - \vec{\mathbf{k}}_0) \right] \\ & \times \exp \left[2(\vec{\mathbf{r}} - \vec{\mathbf{r}}_0)^\top \rho \Sigma^{-1} (\vec{\mathbf{k}} - \vec{\mathbf{k}}_0) \right], \end{aligned} \quad (11)$$

and its corresponding (spatial) wave function takes the form:

$$\begin{aligned} \psi(\vec{\mathbf{r}}) = & \frac{|\text{Re}(\Sigma + i\rho \mathbf{I})^{-1}|}{\pi^3} e^{i\vec{\mathbf{k}}_0 \cdot (\vec{\mathbf{r}} - \vec{\mathbf{r}}_0)} \\ & \times e^{-\frac{1}{2}(\vec{\mathbf{r}} - \vec{\mathbf{r}}_0)^\top (\Sigma + i\rho \mathbf{I})^{-1} (\vec{\mathbf{r}} - \vec{\mathbf{r}}_0)}, \end{aligned} \quad (12)$$

where $\mathbf{I}$ is the identity matrix. We use $\vec{\mathbf{r}}$ for spatial coordinates, and $\vec{\mathbf{k}}$ for wavevectors—directions of propagation scaled by the wavenumber $k = \frac{2\pi}{\lambda}$. In the definitions above, $\vec{\mathbf{r}}_0$ and $\vec{\mathbf{k}}_0$ are the *mean spatial position* and *mean wavevector*, respectively.

The above parameterization of Gaussian beams is chosen in a way that is most compatible with how we use them in light transport simulations: Σ is the initial spatial variance (symmetric positive-definite) matrix of the wave function, and $\rho \geq 0$ is a correlation factor that depends on propagation distance (initially $\rho = 0$ and, as will be shown, it grows linearly with the distance travelled by the wave function). Σ does not change on free-space propagation, but it may mutate upon light-matter interactions.

Under the PLT framework, the emission and sensing distributions of a rather general class of emitters and sensors can be written as a (countable) incoherent superposition of Gaussian beams. For example, the space-frequency (Wigner) distribution of an emitter is decomposed as

$$L(\vec{\mathbf{r}}, \vec{\mathbf{k}}) = \sum_j I_j \phi_j(\vec{\mathbf{r}}, \vec{\mathbf{k}}), \quad (13)$$

where ϕ_j are the Gaussian beams, and I_j is the power contained in each beam. A similar decomposition can be written for a sensor's sensitivity distribution. By design, Gaussian beams under PLT admit properties that are very useful for light transport simulations: they are *weakly-local*, superpose *linearly* in terms of intensity observed by a sensor, and form a *complete* basis. See STEINBERG et al. [SRB*24b] for more information.

Path tracing and rendering with PLT involves: sourcing a sample, i.e. a beam ϕ_j , from an emitter (or a sensor); propagating it across the scene and simulating its interaction with matter; and finally integrating over a sensor (or an emitter).

4.2. Tracing Beams

By discarding the tails (of both the spatial position and wavevector support) of the Gaussian beam in Eq. (11), the beam becomes weakly local. In our implementation we limit the beam's extent to three standard deviations. The weak-locality of Gaussian beams means that we may define a geometric *envelope*—outside of which the beam's power is negligible—and use that envelope for propagation, as well as intersection tests with geometric meshes, in our light

transport simulation. We show next that an *elliptical cone* serves as a tight geometric envelope for Gaussian beams.

Elliptical cones as geometric envelopes Trivially, the first-order moments of a Gaussian beam are

$$\int d\vec{r} d\vec{k} \vec{r} \phi(\vec{r}, \vec{k}) = \vec{r}_0 \quad (14)$$

$$\int d\vec{r} d\vec{k} \vec{k} \phi(\vec{r}, \vec{k}) = \vec{k}_0, \quad (15)$$

i.e. the mean spatial position and wavevector, as desired. The second-order moments (position and wavevector variances) are

$$\text{Var}_{\vec{x}} = \int d\vec{r} d\vec{k} \vec{r} \vec{r}^T \phi(\vec{r}, \vec{k}) - \vec{r}_0 \vec{r}_0^T = \frac{\rho^2}{2} \Sigma^{-1} + \frac{1}{2} \Sigma \quad (16)$$

$$\text{Var}_{\vec{k}} = \int d\vec{r} d\vec{k} \vec{k} \vec{k}^T \phi(\vec{r}, \vec{k}) - \vec{k}_0 \vec{k}_0^T = \frac{1}{2} \Sigma^{-1}. \quad (17)$$

On free-space propagation, a Gaussian beam transforms as:

$$\phi'(\vec{r}, \vec{k}) = \int d\vec{r}_i d\vec{k}_i K \times \phi(\vec{r}_i, \vec{k}_i) = \phi\left(\vec{r} - \frac{z}{k} \vec{k}, \vec{k}\right), \quad (18)$$

$$\text{where } K(\vec{r}, \vec{r}_i, \vec{k}, \vec{k}_i) = \delta\left(\vec{k} - \vec{k}_i\right) \delta\left(\vec{r} - \frac{z}{k} \vec{k}_i - \vec{r}_i\right).$$

Primed variables denote quantities after free-space propagation. K is the free-space propagation kernel (see TESTORF et al. [THO10]), z is the distance of propagation, and $k = |\vec{k}|$. It is easy to verify that after propagation the spatial mean is translated by the propagation distance and the wavevector mean remains unchanged:

$$\vec{r}'_0 = \vec{r}_0 + \frac{z}{k_0} \vec{k}_0 \quad \text{and} \quad \vec{k}'_0 = \vec{k}_0, \quad (19)$$

as expected. The variances and covariances after propagation are

$$\text{Var}'_{\vec{x}} = \frac{(z/k_0 + \rho)^2}{2} \Sigma^{-1} + \frac{1}{2} \Sigma \quad \text{and} \quad \text{Var}'_{\vec{k}} = \frac{1}{2} \Sigma^{-1} = \text{Var}_{\vec{k}}. \quad (20)$$

Therefore, we may conclude that on free-space propagation of distance z : (i) the spatial position mean $\vec{r}_0$ is shifted by $\frac{z}{k_0} \vec{k}_0$; and (ii) the correlation factor ρ is increased by $\frac{z}{k_0}$.

From Eq. (20), we may note that the wavevector variance acts as the (anisotropic) solid angle into which the beam propagates, and the ellipse $\{\vec{r} \mid [(\vec{r} - \vec{r}_0)^T \text{Var}_{\vec{x}} (\vec{r} - \vec{r}_0)]^{1/2} < 3\}$ quantifies the spatial extent around the mean $\vec{r}_0$ where the Gaussian beam remains non-negligible. As the correlation factor ρ and variance $\text{Var}_{\vec{x}}$ increase linearly with propagation distance, that ellipse traces an *elliptical cone* and forms a tight envelope around the Gaussian beam (every point on the elliptical cone's shell is always within 3 standard deviations from the beam's mean).

As an aside, we may recall the fundamental n -dimensional *uncertainty relation* [MW95, Chapter 21]: $|\text{Var}_{\vec{x}} \text{Var}_{\vec{k}}|^{1/2} \geq \frac{1}{2^n}$, and note that (i) the variances of a Gaussian beam, Eqs. (16) and (17), indeed fulfil this relation (as $\rho \geq 0$); and, (ii) when $\rho = 0$ (i.e., initially at sourcing, before the Gaussian beam has propagated) the equality in the uncertainty relation is fulfilled, meaning that Gaussian beams are *as spatially compact as physically permissible*.

We may also note that the correlation between the spatial position and wavevector dimensions is

$$\text{Corr}_{\vec{x}\vec{k}} = \int d\vec{r} d\vec{k} \vec{r} \vec{k}^T \phi(\vec{r}, \vec{k}) - \vec{r}_0 \vec{k}_0^T = \frac{\rho}{2} \Sigma^{-1}. \quad (21)$$

As also observed by MAN'KO et al. [MW08], this implies that

propagation induces correlation between position and wavevector. This correlation is responsible to the increase in the spatial extent occupied by the beam, and means that we may not “break” a Gaussian beam into smaller beams after propagation, as that would violate the uncertainty relation scaled by the induced correlation factor.

Beam parameterization In our implementation, we parameterize each beam by:

1. mean wavelength λ_0 ;
2. Stokes parameters quantifying its polarization and power;
3. its elliptical conic envelope, parametrized by (i) mean spatial position $\vec{r}_0$ at its origin, (ii) mean direction of propagation $\hat{\vec{d}}$, (iii) major and minor axes $\vec{a}, \vec{b}$, of lengths $a = |\vec{a}|, b = |\vec{b}|$, and (vi) major and minor opening (half) angles $\alpha_a, \alpha_b \geq 0$.

The above fully describes the beam's initial spatial footprint as well as the solid angle it propagates into (and hence its spatial position and wavevector variances). As the Gaussian beam is fully described by its first two moments, the geometric envelope defined above fully quantifies the Gaussian beam ϕ , and vice versa.

Traversal Using our propagation formulae, Eqs. (19) and (20), the transformation of the elliptical conic envelope on free-space propagation of distance $z \geq 0$, is as follows:

$$\vec{r}_0 \rightarrow \vec{r}_0 + z \hat{\vec{d}}, \quad a \rightarrow a + z \tan \alpha_a \quad \text{and} \quad b \rightarrow b + z \tan \alpha_b,$$

with the other parameters unchanged. We use an accelerating data structure and geometry intersection test designed for elliptical cones [EKS25]. Traversal is similar to a ray-based traversal, though a fundamental difference is that a cone may intersect multiple triangles. As we traverse the accelerating data structure we collect intersected triangles, and keep track of the distance to the closest triangle $d_{\min}$. The traversal routines returns all triangles that intersect the elliptical cone within a distance range $[d_{\min}, d_{\min} + \delta]$, where δ controls the depth of the intersection range. The “intersection distance” is defined as the propagation distance along the mean $\hat{\vec{d}}$ (and not the radial distance), therefore the computed intersection region is an *elliptical conic frustum* (illustrated in orange in Fig. 2b). We set ad hoc $\delta = 2a$, i.e., twice the length of the major axis, which is sufficient for most applications and a good compromise between accuracy and performance for others.

Once the elliptical conic frustum that contains all the intersected triangles has been computed, it becomes the next interaction region $\mathbf{R}_{j+1}$. The wave tracing process then repeats recursively in order to sample the complete path $\vec{\mathbf{R}}$: the interaction of the beam with the geometry at $\mathbf{R}_{j+1}$ is simulated and a scattered beam—and its elliptical conic envelope—is sampled (discussed in Subsection 4.3), which is then beam traced to compute the next interaction region.

Ballistic propagation In order to allow a scattered beam to propagate away from an interaction region (without extensive self-intersections), as well as to allow it to explore tight regions (e.g., a waveguide that might be narrower than the geometric envelope), we perform *ballistic propagations* [LVD98]. Immediately after sourcing or interaction, we assume the beam takes a *ballistic path*—traces a ray with a vanishing cross section—for a short distance, as follows:




Figure 3: Comparison with a wave solver. a collimated beam impinges upon a screen with two slits cut in it, and a couple of conductive reflectors are placed after the screen (setup illustrated on the left). Diffracted light reaches the far wall and the resulting diffraction pattern is rendered. We use a BEM wave solver [BS21] to solve for the electric field over a 2D vertical slice, and compare with results computed using our renderer (UTD free-space diffractions), see plots. Our results are shown below the plots. Multi-edge diffractions happen over two regions—slits and then reflectors—before reaching the wall. We also perform the experiment without the reflectors (slits only) and plot analytic solutions to Fraunhofer and UTD diffractions: as expected our results closely match UTD and are more accurate than Fraunhofer diffraction. **Simulation costs.** Ours, full 3D: 4 min and 195 MB of memory; BEM, 2D slice: 2 hours, 27 min and 60 GB.

1. We propagate the beam as a ballistic particle—via ray tracing—for a segment of up to B wavelengths in distance.
2. If the ray intersects geometry, a normal interaction is performed; after which we return to step 1.
3. Otherwise, we attempt to resume cone tracing from the end of the ballistic segment:
 - a. if the elliptical cone immediately (without any further propagation) intersects geometry:
 - i. set $B \rightarrow s \times B$, up to some predefined maximum $B_{\max}$;
 - ii. return to step 1.
 - b. otherwise, elliptical cone tracing is continued.

The scaling of the ballistic segment length B by $s > 1$ is designed to avoid excessive attempts to terminate ballistic propagation. Ballistic propagation also alleviates some of the chronic pains of cone tracing, like a beam propagating parallel to a wall and repeatedly intersecting it. A ballistic path is illustrated via a thick line in Fig. 2b.

In addition to enabling a beam to explore tight confines in a physically-sensible manner, ballistic propagation also serves as a major performance optimization: the most expensive part of cone tracing is when the cone passes close to geometry and many expensive cone-primitive intersection tests are required. Given a wavelength λ , our renderer sets ad hoc the initial ballistic segment length $B = 2\lambda$, the max length $B_{\max} = 2^{16}\lambda$, and the scaling factor $s = 4$.

4.3. Interactions and Diffractions

Let $\mathbf{R}_j$ be an interaction region: that is, the intersection of an elliptical conic frustum with geometry. The incident beam ϕ_j interacts with the matter and materials, and the scattered distribution from an interfering superposition of $N \geq 1$ triangles (and their materials) is a *bilinear* combination that accounts for interference:

$$\Phi(\vec{r}, \vec{k}) = \sum_{l=1}^N \mathcal{T}_l \phi_j + 2 \sum_{m>l}^N \mathcal{T}_{lm} \phi_j. \quad (22)$$

The terms $\mathcal{T}_l$ are operators which compute the scattered light produced by each triangle (or material) l . The double-subscripted scattering operators $\mathcal{T}_{lm}$ quantify the bilinear interference terms between all triangle pairs. This is a very general expression—any interfering wave-optical interaction needs to account for such interference terms. For example, if a triangle’s material models a rough

surface, then $\mathcal{T}_l$ computes the scattering of a Gaussian beam by that surface over the triangle’s extent; often, an approximate surface scattering model will be used. See our supplemental material for more details.

Our goal is to sample a new beam and its envelope from Eq. (22), enabling us to beam trace it and compute the next interaction region $\mathbf{R}_{j+1}$. Formally, the power contained in some scattered Gaussian beam ϕ' is the inner functional product:

$$I' = \langle \phi', \Phi \rangle = \sum_{l=1}^N \langle \phi', \mathcal{T}_l \phi_j \rangle + 2 \sum_{m>l}^N \langle \phi', \mathcal{T}_{lm} \phi_j \rangle, \quad (23)$$

with the inner product taken over the space-frequency (Wigner) space: $\langle \phi', \Phi \rangle = \int \phi' \Phi^* d\vec{r} d\vec{k}$. The (unknown) beam ϕ' defines its geometric envelope, which in turn defines the next interaction region. Therefore, if we could derive from Eq. (23) and expression for I' as a function of the envelope’s geometry, we could importance sample a scattered beam from the interaction.

PLT invariants help us narrow down what kind of ϕ' we should consider: The space-frequency bandwidth of the Gaussian beams ϕ' that decompose the scattered light Φ cannot decrease (we may not “break a beam into smaller beams”). Because the Gaussian beam is fully described by its elliptical conic envelope, this constraint implies a geometric invariant: the envelope of the new ϕ' contains the current interaction region $\mathbf{R}_j$ and propagates into an identical solid angle as the incident beam.

To satisfy this geometric invariant, when constructing a new elliptical conic envelope for an arbitrary new tracing direction $\hat{\mathbf{d}}'$ we proceed as follows: we solve for the elliptical cone that contains $\mathbf{R}_j$ and propagates into the same solid angle as the incident beam, i.e. $\tan \alpha'_a \tan \alpha'_b = \tan \alpha_a \tan \alpha_b$. The ratio α'_a / α'_b is chosen such that the beam’s eccentricity remains constant (this is valid and simplifies some intersection tests). This fixes the envelope, which in turn defines the scattered Gaussian beam ϕ' .

In general, computing closed forms expressions for the inner products in Eq. (23) can be difficult. In our supplemental material, Section S6, we study what assumptions (or approximations) need to be imposed upon the scattering operators $\mathcal{T}_l$, in order to be able to write these inner products as functions of only incident and scattering directions and interaction footprint. We call such materials *Fraunhofer materials*, and show that for these materials the inner



Figure 5: Wave optics rendering. A metal screen is placed between a powerful collimated light source and a wall. A geometric aperture is cut in the middle of the screen: (left) star shaped or (right) double horizontal slits. Our simulation reproduces the expected diffraction patterns. To make the light transport more challenging we place a (left) mirror or (right) dielectric bunny, making direct sampling techniques difficult. This shows that we are able to importance sample the entire multi-edge diffraction well. This scene is miniaturized: the length of the far wall is 8 cm. The cross-sectional areas of the elliptical conic envelopes on interactions are $(0.2 \pm 1.0) \text{ mm}^2$ (mean and standard deviation), and intersect with 14 ± 261 triangles. Fraunhofer edge-based free-space diffractions are used.




Figure 6: Indoor 5GHz simulations. We simulate signal coverage with a 5GHz wireless carrier in two indoor scenes with more challenging light transport. Visualized is the received signal strength (RSS) in decibels that reaches the floor. No geometric simplifications are done; the geometry used for the wireless simulation is as visualized in the optical rendering insets. The surface materials are different from the optical materials: at 5GHz frequency, surfaces are effectively perfectly specular, and the material (refractive index) is selected from the appropriate construction material [ITU23]: wood, plywood, chipboard, plasterboard, concrete, or metal. (Office) This scene contains complex geometry and fine geometric details. The transmitter is placed at the top right corner; note the increased RSS around the lobby on the left and the meeting room on the right. (Building) RSS is computed on the 3rd floor, while the transmitter is placed at 2nd floor level in the stairwell (top right). Due to the transmitter's placement, most emitted energy does not propagate to the 3rd floor, or is diffracted by the staircase or railing. Paths that diffract several times are needed to reach the far areas of the large room on the left. Both scenes contain many glass windows, through which significant energy radiates away. The elliptical conic envelopes' cross-section areas and primitive counts on interactions (means and standard deviations): (Office) $(1.100 \pm 1.364) \text{ m}^2$, 34 ± 470 ; (Building) $(2 \pm 40) \text{ m}^2$, 29 ± 905 . UTD-based free-space diffractions are used (see Subsection 4.3.2).

edge diffraction operator $\mathcal{T}_{e_i}$, but not for the more complicated interference operators $\mathcal{T}_{e_{i,m}}$. Our importance sampling strategy from Subsection 4.3.1 improves upon theirs by importance sampling the entire interfering sum of edge diffractions, enabling high quality sampling. We show in Fig. 5 that we are able to render complex diffraction patterns under complex light transport.

UTD-based diffraction Fraunhofer-based diffractions are less accurate for long-wavelength radiation, assume perfect conductors, ignore wedge geometry, and it can be difficult to correctly classify edges which are partially occluded by triangles within the interaction region. As an alternative, we propose a method that relies on the uniform theory of diffraction (UTD) [MPM90] to simulate free-

space diffractions. This method can be understood as tracing a ray bundle confined within the beam's envelope. At each interaction region $\mathbf{R}_j$, we perform UTD-based diffraction by connecting rays from the centre of the previous region $\mathbf{R}_{j-1}$ to the centre of the next one $\mathbf{R}_{j+1}$ over the edges in $\mathbf{R}_j$, by checking if the diffracted Keller cone for each edge in $\mathbf{R}_j$ falls upon the destination. Using the region's centres is a simplifying approximation; proper bundle-to-bundle connections are left for future work.

Similar to before, the rays in the bundle give rise to diffracted terms $\mathcal{T}_{e_n}$, however these are not Fraunhofer materials, and importance sampling is more difficult. Our sampling approach is very basic: a conservative frequency-dependent Gaussian is fitted to the UTD edge diffraction function, and at each interaction region an

edge (that is involved in the diffraction) is selected at random and the next direction is sampled from that Gaussian. We do not consider the interference terms when sampling, nor perform multiple importance sampling (MIS) over the different edges. This leads to rather poor sampling; a sophisticated importance sampling strategy for UTD-based ray bundle diffractions is left for future work.

5. Discussion

Subsections 4.2 and 4.3 describe the primary components of wave tracing as done in our system. We construct region-to-region paths by tracing the elliptical conic envelopes of Gaussian beam samples. Cone tracing allows us to sample a complete interaction region, with all the contained geometry that the Gaussian beam interacts with. In contrast to ray-based frameworks like Sionna [HCA*22], we do not need to perform a shooting-bouncing rays exploration pass in order to detect edges (which might fail to detect relevant edges that may lie in shadowed, unexplored regions, or due to high tessellation of the geometry). Cones are inherently able to sample zero-measure features, like the edges needed for free-space diffraction, and could also be used to sample caustics in a unidirectional path tracer; to be explored in future work.

The beams and their envelopes often exhibit significant anisotropy, for example, when a beam is scattered into a grazing angle on surface interaction. Proper wave tracing then requires support for full elliptical cone tracing, with accelerating data structure traversal and primitive (edges and triangles) intersection tests for arbitrary elliptical cones. We provide implementation details in our supplemental material.

5.1. Results

We have validated our approach by comparing results to a BEM wave solver; see Fig. 3. We find that our renderer overall matches the reference well, yet computes solutions in 3D (rather than a 2D slice) while requiring a fraction in computation and memory resources. We have also evaluated our algorithms for optical rendering and modeling radio wave propagation.

In these results, we also report the means and standard deviations of the areas of the beams' cross-sectional areas and primitive count at interaction regions. Note the large difference in areas between different wavelengths. For convenience, these statistics are also summarised in Table 2 for all of our results.

Optical rendering Our optical results, Fig. 5, demonstrate the ability of our implementation to simulate free-space diffractions in scenes with extensive geometric detail and complex light transport. All renders are spectral and polarimetric.

These results were rendered with a bidirectional path tracing (BDPT) extension of our weakly-local path integral. When forming connections between paths under BDPT, we wave trace each beam from the two path vertices involved in the connection strategy halfway towards each other. If both beams are unoccluded, the beams meet and we integrate the beams over their respective space-frequency footprint [SRB*24b]. Compared with the classical approach, multiple importance sampling (MIS) for BDPT changes as




Figure 7: Received signal strength (RSS) differences between ray tracing and wave tracing. Blue indicates areas where a ray tracing simulation computes a higher RSS compared with wave tracing, and red indicates regions where the situation is reversed.

some connection strategies may be occluded while others are not. For our implementation, we use the classical strategy as an approximation, and leave correcting MIS for BDPT for future work. See Section S4 of our supplemental material for more information on our BDPT extension and explicit formulae.

Signal coverage We perform signal coverage simulations with long-wavelength radiation in two large-scale city scenes (Fig. 4) and two more complex indoor scenes (Fig. 6). At these wavelengths, all the surfaces become perfectly specular, and light transport is simulated unidirectionally from the emitter only. The wavelength-dependent complex refractive indices for all materials are given by ITU [ITU23] recommendations for construction materials, and include concrete, glass, brick, wood, and metal surfaces. For comparison, same scenes are rendered with pure ray tracing; note the significantly higher received signal strength that reaches the shadowed regions with wave tracing. See our supplemental for additional analysis.

Ballistic paths For our optical results (Fig. 5), less than 10% of the path segments are ballistic. These ballistic paths arise around corners (e.g., when light is reflected off the very bottom of the screen into the table), as well as in tight confines (around the flowers). Transport with long-wavelength simulations requires larger-aperture cones, and as expected a higher ratio of path segments are ballistic: about 17% and 25% for the Etoile and München scenes (Fig. 4), respectively.

Ballistic propagation is an effective, practical method to enable particle transport in the presence of complex geometry.

5.1.1. Additional Comparison with ray tracing

We compare the received signal strength (RSS) that reaches the floor in the Building and Office scenes between ray tracing and wave tracing, see Fig. 7. The color-coded blue and red areas indicate areas where ray tracing or wave tracing, respectively, produce a greater RSS. As expected, over most areas wave tracing gives rise to a higher RSS, as light is able to diffract and penetrate into harder-to-reach regions. Nevertheless, in some areas ray tracing dominates, at times substantially so. These blue areas are *errors induced by the ray tracing simulation*: light that should have diffracted continues to propagate unobstructed. For example, light

that is reflected by ceiling fixtures towards the lobby in the Office scene (Fig. 7a left), or light that passes through the railings around the stairwell in the Building scene (Fig. 7b top right).

5.1.2. Comparison with pure-ballistic propagation

STEINBERG et al. [SRB*24a] present a light transport approach where incoherent ray tracing is used for propagation and diffraction effects are simulated at ray intersection points by considering the geometry around the intersection. This can be understood as propagating beams using purely ballistic paths—by ray tracing—and doing full beam–matter interactions at the ray intersection points. We compare our beam tracing approach with such pure-ballistic propagation in Fig. 8. Observe that when light passes close to geometry it diffracts, resulting in some of its energy scattering elsewhere: shadow penumbras are enlarged under wave optics. Pure-ballistic propagation cannot reproduce that effect, as nearby geometry is not detected and rays continue to propagate undisturbed.

Pure-ballistic propagation methods are *inherently biased*, nor do they conserve energy: the energy for the diffracted lobes should come from light that passes through the screen openings and diffracts (as done by our method), however that is not possible with classical ray tracing. Therefore, in order to reproduce the diffraction lobes when a ballistically-propagated beam intersects geometry, STEINBERG et al. [SRB*24a] inject additional forward-scattered energy into the system. This is unavoidable: as discussed in Section 3, and in Section S1 in our supplemental material, incoherent ray tracing—formalized by the classical path integral—cannot faithfully reproduce diffraction phenomena, and a generalized formulation of light transport is required. Our method does not suffer from this bias because we only perform full beam–matter interactions when proper beam tracing is done, and not on ballistic intersections (instead, on ballistic intersections we perform classical ray–matter interactions).

The extent by which shadow penumbras are enlarged is roughly equal to the cross-sectional length of the traced beams. See Table 2 for a summary of these statistics for all our results. Note that for long-wavelength radiation, this extent can be significant (roughly several-to-many metres in length).




Figure 8: Comparison with pure ballistic propagation. We setup a similar experiment to Fig. 3: (a) Light impinges upon and is diffracted by a screen with several openings cut in it, thereafter the pattern it produces upon the wall is imaged. We compare several simulation approaches: our beam tracing method; pure-ballistic propagation, where beams are propagated via ray tracing (as done by STEINBERG et al. [SRB*24a]); and geometric optics. The simple planar setup also enables computing a numeric ground truth using UTD. We (b) plot the intensity (its square root for visualization) of light falling upon the wall, as well as (c) the rendered pattern (square root of intensity and Decibels for the top and bottom half, respectively, of each pattern). Both our method and pure-ballistic propagation reproduce the diffraction pattern with good accuracy; however, the latter overestimates the forward-scattered energy, which has to account for both the geometric optics direct lobes as well as the wave optics diffraction lobes. The screen elements are 2.5 mm in length, and spaced 3.5 mm apart.

5.1.3. Comparison with Sionna

Sionna [HCA*22] is a communications systems simulation framework that employs UTD to simulate diffractions. An initial ray tracing pass is used to find triangles and their edges that might participate in diffraction. All these edges, across all samples, are stored in memory, and a subsequent pass attempts to connect emitter–edge–receiver diffraction paths via every found edge and to every receiver (or pixel in a coverage map). No diffractions are mixed with reflections or scattering events, likewise no multiple diffractions, or multi-edge diffractions are simulated.

SCENE		ELLIPTICAL CONE STATISTICS		THROUGHPUT (samples/millisecond)	
		cross-sectional area (mean $\pm$ std. dev.)	triangles per cone (mean $\pm$ std. dev.)		
Etoile	(Fig. 4)	30 GHz	(56 $\pm$ 250) m ²	7 $\pm$ 21	5.1k
		1 GHz	(90 $\pm$ 290) m ²	9 $\pm$ 22	5.9k
München	(Fig. 4)	30 GHz	(49 $\pm$ 298) m ²	11 $\pm$ 129	900
		1 GHz	(84 $\pm$ 364) m ²	10 $\pm$ 107	1.1k
Office	(Fig. 6)	5 GHz	(1.100 $\pm$ 1.364) m ²	34 $\pm$ 470	150
Building	(Fig. 6)	5 GHz	(2 $\pm$ 40) m ²	29 $\pm$ 905	1.05k
Room	(Fig. 5)	optical	(0.2 $\pm$ 1.0) mm ²	14 $\pm$ 261	1.2k

Table 2: Performance statistics. A summary of the average cross-sectional area of a traced elliptical cone on an interaction, as well as average count of triangles per interaction; as well as performance (samples per millisecond) for all of the scenes in our results.




Figure 9: Comparison with Sionna. The scene is “simple street canyon” from HOYDIS et al. [HCA*22]. The position of emitter illustrated by the yellow circle. The received signal strength (RSS) in decibels that reaches the street level is visualized, with a rendering of the occluding and diffracting building geometry overlaid (buildings’ colors are for visualization only). Not only can our approach account for path that involve both diffractions and reflections, but ours uses significantly less memory as well.

We do a comparison with Sionna (version 0.19) in a very simple scene, where single-edge, direct (emitter–edge–receiver) diffractions dominate, see Fig. 9. There are small differences between the material and emitter models, nevertheless most differences are minor. Most differences arise in the bottom parts of the image, where the majority of the energy arrives by reflecting off the buildings first and then diffracting, which cannot be simulated with Sionna.

Note the memory requirements: Sionna’s memory requirements become prohibitive with more complex scenes that require higher sample counts.

6. Conclusion

Our discussion began with studying the power—in terms of reproducing wave-optical phenomena—of the classical path integral formulation of light transport (see our supplemental material, Section S1, for the complete discussion); then, we generalized it to two formulations that are used in wave simulations: (i) ray-based bilinear transport; and (ii) weakly-local (region-to-region) transport. Region-to-region path tracing, as formalized by our weakly-local path integral formulation, enables formulating—and in turn, importance sampling—scattering functions that account for the entire interfering interaction within a region. Contrast this to typical shooting-bouncing rays or UTD-based approaches, where interference is resolved later at the sensor (as formalized by our bilinear path integral), and importance sampling turns into a global problem.

We discussed how to design a system for wave simulations, based on our weakly-local path integral formulation, and have released a complete rendering system that implements the ideas in this paper. Our renderer is a fully polarimetric spectral renderer, designed from scratch to target elliptical cone tracing and wave simulations across the EM spectrum. Our system supports multi-

material and multi-edge interactions, as well as multiple such interactions over a path.

Both our path integral generalizations serve to relax some of the limitations of the classical path integral. The far-field and high-frequency assumptions are often relaxed (the degree of which depends on the optical formalism), and the restriction to statistically stationary materials is eliminated entirely. Future work may focus on eliminating the time-averaging assumption, enabling interference across spectral samples for acoustics simulations, for example. Future work may also target the differentiability of our rendering system for wave-optics modeling; differentiable rendering is likely to be more efficient in a non-bilinear framework like ours.

References

- [Arv95] ARVO, JAMES. “Analytic Methods for Simulated Light Transport”. PhD thesis. Yale University, 1995 2.
- [BB17] BELCOUR, LAURENT et al. “A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence”. *ACM Trans. Graph.* 36.4 (July 2017), 65:1–65:14. ISSN: 0730-0301. DOI: [10 . 1145 / 3072959.3073620](https://doi.org/10.1145/3072959.3073620) 2.
- [BS21] BETCKE, TIMO et al. “Bempp-cl: A fast Python based just-in-time compiling boundary element library”. *J. Open Source Softw.* 6.59 (Mar. 2021), 2879 7.
- [BVD20] BILIBASHI, D et al. “Dynamic ray tracing: Introduction and concept”. *2020 14th European Conference on Antennas and Propagation (EuCAP)*. IEEE, 2020, 1–5 2.
- [CHB*12] CUYPERS, TOM et al. “Reflectance model for diffraction”. *ACM Transactions on Graphics* 31.5 (Aug. 2012), 1–11. DOI: [10 . 1145/2231816.2231820](https://doi.org/10.1145/2231816.2231820) 2.
- [CRW04] CASELLA, GEORGE et al. “Generalized Accept-Reject sampling schemes”. *A Festschrift for Herman Rubin*. Institute of Mathematical Statistics, 2004, 342–347. DOI: [10.1214/lnms/1196285403](https://doi.org/10.1214/lnms/1196285403) 8.
- [EKS25] EMRE, UMUT et al. “High-Performance Elliptical Cone Tracing”. *Computer Graphics Forum* 44.7 (Oct. 2025). DOI: [10 . 1111 / cgf.70230](https://doi.org/10.1111/cgf.70230) 3, 6.
- [FJF20] FALSTER, V. et al. “Computing the Bidirectional Scattering of a Microstructure Using Scalar Diffraction Theory and Path Tracing”. *Computer Graphics Forum* 39.7 (Oct. 2020), 231–242. DOI: [10.1111 / cgf.14140](https://doi.org/10.1111/cgf.14140) 2.
- [GMG*20] GUILLÉN, IBÓN et al. “A General Framework for Pearlescent Materials”. *ACM Transactions on Graphics* 39.6 (2020). DOI: [10 . 1145/3414685.3417782](https://doi.org/10.1145/3414685.3417782) 2.
- [HCA*22] HOYDIS, JAKOB et al. *Sionna*. Version 1.0.2. <https://nvlabs.github.io/sionna/>. 2022 8, 10–12.
- [HIK*20] HUANG, WEIZHEN et al. “Chemomechanical simulation of soap film flow on spherical bubbles”. *ACM Transactions on Graphics* 39.4 (July 2020). DOI: [10.1145/3386569.3392094](https://doi.org/10.1145/3386569.3392094) 2.
- [HP17] HOLZSCHUCH, NICOLAS et al. “A Two-scale Microfacet Reflectance Model Combining Reflection and Diffraction”. *ACM Trans. Graph.* 36.4 (July 2017), 66:1–66:12. ISSN: 0730-0301. DOI: [10 . 1145/3072959.3073621](https://doi.org/10.1145/3072959.3073621) 2.
- [ITU23] ITU. *Recommendation ITU-R P.2040-3*. 2023. URL: https://www.itu.int/dms_pubrec/itu-r/rec/p/R-REC-P.2040-3-202308-I!!PDF-E.pdf 8–10.
- [JB91] JENSEN, KL et al. “The methodology of simulating particle trajectories through tunneling structures using a Wigner distribution approach”. *IEEE Transactions on Electron Devices* 38.10 (1991), 2337–2347 2.

- [Kaj86] KAJIYA, JAMES T. “The rendering equation”. *Proceedings of the 13th annual conference on Computer graphics and interactive techniques - SIGGRAPH '86*. ACM Press, 1986. DOI: [10.1145/15922.159024](https://doi.org/10.1145/15922.159024).
- [KGK19] KNEIPHOF, TOM et al. “Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence”. *Computer Graphics Forum* 38.4 (2019), 77–85. DOI: [10.1111/cgf.13772](https://doi.org/10.1111/cgf.13772).
- [Kry06] KRYWONOS, ANDREY. “Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction”. PhD thesis. University of Central Florida, Jan. 2006 [2](#).
- [LVD98] LEMIEUX, P-A et al. “Diffusing-light spectroscopies beyond the diffusion limit: The role of ballistic transport and anisotropic scattering”. en. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* 57.4 (Apr. 1998), 4498–4515 [6](#).
- [MJ21] MACKAY, JACOB et al. “Millimetre wave ray tracing simulator with phase and beam effects using the Wigner distribution function”. *Passive and Active Millimeter-Wave Imaging XXIV*. Vol. 11745. SPIE. 2021, 43–57 [2](#).
- [MPM90] MCNAMARA, D A et al. *Introduction to the uniform geometrical theory of diffraction*. en. Antennas & Propagation Library. Norwood, MA: Artech House, Jan. 1990 [9](#).
- [MW08] MAN'KO, VLADIMIR I et al. “The influence of spherical aberration on Gaussian beam propagation”. *Lie Methods in Optics*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, 207–225 [6](#).
- [MW95] MANDEL, LEONARD et al. *Optical coherence and quantum optics*. Cambridge: Cambridge University Press, 1995. ISBN: 978-0521417112 [6](#).
- [MWB*18] MOUT, MARCO et al. “Ray tracing the Wigner distribution function for optical simulations”. *Optical Engineering* 57.1 (2018), 014106–014106 [2](#).
- [SM99] SON, HAE-WON et al. “A deterministic ray tube method for microcellular wave propagation prediction model”. *IEEE Transactions on Antennas and Propagation* 47.8 (1999), 1344–1350 [2](#).
- [SRB*24a] STEINBERG, SHLOMI et al. “A Free-Space Diffraction BPDF”. *ACM Trans. Graph.* 43.4 (July 2024). ISSN: 0730-0301. DOI: [10.1145/3658166](https://doi.org/10.1145/3658166) [2](#), [8](#), [11](#).
- [SRB*24b] STEINBERG, SHLOMI et al. “A Generalized Ray Formulation For Wave-Optical Light Transport”. *ACM Trans. Graph.* 43.6 (Nov. 2024). ISSN: 0730-0301. DOI: [10.1145/3687902](https://doi.org/10.1145/3687902) [2](#), [5](#), [10](#).
- [SSY22] STEINBERG, SHLOMI et al. “Towards Practical Physical-Optics Rendering”. *ACM Transactions on Graphics* 41.4 (July 2022), 1–13. DOI: [10.1145/3528223.3530119](https://doi.org/10.1145/3528223.3530119) [2](#), [5](#).
- [Sta99] STAM, JOS. “Diffraction shaders”. *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99*. ACM Press, 1999. DOI: [10.1145/311535.311546](https://doi.org/10.1145/311535.311546) [2](#).
- [SY21] STEINBERG, SHLOMI et al. “A Generic Framework for Physical Light Transport”. *ACM Transactions on Graphics* 40.4 (Aug. 2021), 1–20. DOI: [10.1145/3450626.3459791](https://doi.org/10.1145/3450626.3459791) [2](#).
- [SY22] STEINBERG, SHLOMI et al. “Rendering of Subjective Speckle Formed by Rough Statistical Surfaces”. *ACM Trans. Graph.* 41.1 (Feb. 2022). ISSN: 0730-0301. DOI: [10.1145/3472293](https://doi.org/10.1145/3472293) [2](#).
- [TG17] TOISOUL, ANTOINE et al. “Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance”. *ACM Transactions on Graphics* 36.5 (July 2017), 1–16. DOI: [10.1145/3012001](https://doi.org/10.1145/3012001) [2](#).
- [THO10] TESTORF, MARKUS et al. *Phase-Space Optics: Fundamentals and Applications*. McGraw-Hill Education, 2010 [6](#).
- [Vea97] VEACH, ERIC. “Robust Monte Carlo Methods For Light Transport Simulation”. PhD thesis. Stanford University, 1997 [2–4](#).
- [VG97] VEACH, ERIC et al. “Metropolis light transport”. *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, 65–76 [2](#).
- [VWH18] VELINOV, Z. et al. “Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces”. *Computer Graphics Forum* 37.2 (2018), 123–134. DOI: [10.1111/cgf.13347](https://doi.org/10.1111/cgf.13347) [2](#).
- [Wei06] WEINMANN, FRANK. “Ray tracing with PO/PTD for RCS modeling of large complex objects”. *IEEE Transactions on Antennas and Propagation* 54.6 (2006), 1797–1806. DOI: [10.1109/TAP.2006.875910](https://doi.org/10.1109/TAP.2006.875910) [2](#).
- [WMB*25] WU, LIFAN et al. “Unbiased Differential Visibility Using Fixed-Step Walk-on-Spherical-Caps And Closest Silhouettes”. *ACM Trans. Graph.* 44.4 (July 2025). ISSN: 0730-0301. DOI: [10.1145/3731174](https://doi.org/10.1145/3731174) [3](#).
- [WVJH17] WERNER, SEBASTIAN et al. “Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure”. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36.6 (Nov. 2017). DOI: [10.1145/3130800.3130840](https://doi.org/10.1145/3130800.3130840) [2](#).
- [YHM*22] YI, HAOFAN et al. “Ray tracing meets terahertz: Challenges and opportunities”. *IEEE Communications Magazine* (2022) [2](#).
- [YXW*23] YU, YUNCHEN et al. “A Full-Wave Reference Simulator for Computing Surface Reflectance”. *ACM Trans. Graph.* 42.4 (July 2023). ISSN: 0730-0301. DOI: [10.1145/3592414](https://doi.org/10.1145/3592414) [2](#).