



HAL
open science

Optimised Light Rendering through Old Glass

Quentin Huan, François Rousselle, Christophe Renaud

► **To cite this version:**

Quentin Huan, François Rousselle, Christophe Renaud. Optimised Light Rendering through Old Glass. WSCG 2023 – 31. International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision, May 2023, Pilsen, CZ, Czech Republic. pp.258-267, 10.24132/CSRN.3301.30 . hal-04264368

HAL Id: hal-04264368

<https://hal.science/hal-04264368v1>

Submitted on 30 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimised Light Rendering through Old Glass

Quentin Huan
Université du Littoral
Côte d'Opale
LISIC
BP 719
France, 62228, Calais
cedex
quentin.huan@univ-
littoral.fr

Francois Rousselle
Université du Littoral
Côte d'Opale
LISIC
BP 719
France, 62228, Calais
cedex
francois.rousselle@univ-
littoral.fr

Christophe Renaud
Université du Littoral
Côte d'Opale
LISIC
BP 719
France, 62228, Calais
cedex
christophe.renaud@univ-
littoral.fr

ABSTRACT

We propose a rendering method for efficiently computing the transmitted caustics produced by a glass panel with arbitrary surface deformations, characteristic of old glass used in 3D reconstructions in virtual heritage. Using Fermat's principle of least time, we generalize the concept of Next Event Estimation to allow light sampling through two displaced refractive interfaces, which amount to numerically finding all stationary points of an objective function. Our work allows for an efficient estimation of the caustic while staying inside a standard Monte Carlo pathtracing framework. Our specific geometrical context allows our solver to converge significantly faster than the more general method Specular Manifold Sampling, while scaling well with the number of panels present in the scene.

Keywords

Ancient glass, Virtual heritage, Caustics, Pathtracing, Sampling, Newton's method

1 INTRODUCTION

Recent progress in lighting simulation offers a true opportunity for recreating the luminous atmosphere that existed in ancient architectures. The virtual restitution of architectural heritage gives to historians a new set of tools allowing a better analysis and understanding of life, work and worship conditions through comparison with the existing written archives.

For centuries, glass has been used in the windows of most buildings. However, the production of high quality glass is relatively recent and old glass, as used in windows and stained glass, has specific visual characteristics that greatly impact the lighting of indoor scenes. These include surface irregularities, or the presence of various bubbles and debris due to the different manufacturing processes (crown or cylinder blown sheet glass), which did not allow for perfectly flat or homogeneous surfaces. These irregularities, however small, produce distortions in the perception of the external environment but also complex lighting patterns

(caustics) on objects illuminated directly by a source (sun, flames) through these types of glass (see Fig. 1). The composition of the glass paste could also lead to slightly colored glass, influencing the color of perceived light.

These effects are still challenging to compute with modern rendering techniques, often requiring prohibitive rendering time to produce noise-free images. On the other hand, they are of great interest to the historian, in order to understand the differences in light atmosphere that may have existed in the past compared to the atmosphere produced by our modern glazing.




Figure 1: Photography of a caustic produced by a slightly irregular panel of glass.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

In this article, we are only considering the lighting effects produced by the irregularities of glass surfaces in windows. By neglecting volume irregularities (bubbles, impurities inclusions e.g.) and the possible colouring of the glass, we can model an ancient glass panel by two parallel and slightly displaced planes. By taking advantage of this particular geometrical context, we can significantly accelerate the rendering time of scenes including these complex objects.

The contributions of this paper are:

- a geometrical framework for modeling the surface irregularities of ancient glass panels ;
- a solver based on Fermat’s principle allowing the connection of two points through two refractive interfaces ;
- a robust initialisation strategy for path sampling that scales well with the number of glass panels present in the scene.

In the next paragraph, we review previous works relevant to our problem from the point of view of antique glass and caustics rendering. In paragraph 3, we detail the different steps of our approach. We then present the results obtained in terms of convergence speed of our solver, but also from the point of view of the interest of the initialization strategy that we propose when there are many windows. We discuss in section 5 some residual difficulties of our approach which are also present on the approaches of the state of the art, then conclude this article by proposing some avenues for future research.

2 PREVIOUS WORK

To our knowledge, no previous work specifically tackle the rendering of the caustics produced by ancient glass panels. Kider Jr et al. [Jr+09] recreate the characteristic caustic lighting pattern produced by early Islamic light sources using a specialised technique. This method doesn’t take in account the surface imperfections of the glass fixture, resulting in a simplified simulation of the effects. Grobe et al. [GNL20] measure the BSDF of various flat Roman window glass samples. They then proceed to simulate daylight lighting of an interior scene using a data-driven transmission model. While this work captures the small scale in-homogeneity of the material, it doesn’t take in considerations the larger surface variations that play an essential role for transparent and homogeneous glass samples that only refract incident light without significantly scattering it.

The rendering of caustics produced by refractive or reflective objects has been a long lasting problem in computer graphics. While standard pathtracing is able to compute caustics generated by a specular object lighted

by an area light, its slow convergence rate makes it difficult to use in practice. Many different rendering techniques have been designed to accelerate the convergence of the standard algorithm.

Pathtracing

Pathtracing (PT) is a popular unbiased estimation method of the rendering equation originally described by Kajiya [Kaj86]. While being able to accurately render scenes with complex light transport (including caustics), the standard algorithm may suffer from impractically slow convergence rates even with GPU acceleration. Although slow, pathtracing can be used to produce accurate ground truth reference images.

The more sophisticated Bidirectional Path Tracing technique [LW93] speeds up convergence for Diffuse-Specular paths (see the A area in Fig. 2b), but remains unable to efficiently deal with Specular-Diffuse-Specular (SDS) paths that are commonly encountered with reflective or refractive objects (see Fig. 2b, where the B area is slower to converge).

Photon Mapping

Techniques based on Photon Mapping [Jen96] are generally well suited for caustics rendering. Photons coming from the light sources are traced around the scene and stored in an auxiliary data structure called a photon map. Since the amount of photons traced is finite, a photon density estimation step is needed in order to compute the rendering integral, which requires a voluminous amount of storage in certain cases.

In related previous work [Shi90], the contribution of illumination rays coming from the light sources were stored in a texture (an illumination map) associated to each diffuse surfaces. This however requires a sufficiently high resolution for the illumination map to avoid artifacts.

The stochastic progressive photon mapping technique (SPPM) proposed by Hachisuka and Jensen [HJ09] allows the progressive construction of the photon map during render time, alleviating the storage issue. While this technique handles SDS paths, it also introduces bias in the final image causing the caustic to appear blurry (see Fig. 2c).

Metropolis Light Transport

Metropolis Light Transport (MLT) of Veach and Guibas [VG97] uses Metropolis-Hasting integration to evaluate the rendering integral. A path space formulation of the rendering equation allows the construction and mutation of a group of bootstrap paths, allowing the




Figure 2: Comparison of different rendering techniques for a simple scene including a glass panel whose surfaces are not smooth (identical number of 200 samples per pixel (spp)). Images were rendered with PBRT-V4 using (a) Pathtracing, (b) Bidirectional Path Tracing, (c) Stochastic Progressive Photon Mapping and (d) Metropolis Light Transport. Image (e) Specular Manifold Sampling was rendered with the Path-SMS-MS integrator (using the author implementation in Mitsuba2 [Nim+19]) with 2 specular bounces (thus lacking the reflected caustic).

local exploration of path space. This approach is particularly efficient for scenes in which the light transport mainly occurs throughout few highly contributing paths that are difficult to sample using traditional techniques. Intrinsically, the convergence process is highly dependent of the random seed used and non-uniform by nature: some paths that are difficult to sample may be explored lately during render time, causing some brightness inconsistencies until the process completely converges (see Fig. 2d where the SDS paths can be alternatively brighter or darker than expected depending on the seed). This may lead to unpleasant flickering artifacts when rendering animations.

Manifold Exploration

Early work from Mitchell and Hanrahan [MH92] makes use of Fermat’s principle of least time combined with interval arithmetic for computing the illumination produced by reflective, implicitly defined surfaces. This method is deterministic and is able to compute the whole set of paths connecting two points through a one-bounce perfectly specular reflection.

Similarly with Manifold Exploration [JM12], Jakob and Marschner present a new set of mutations for MLT that are more suited to paths involving purely specular events. This set of paths is a manifold described by a set of specular reflection and refraction constraints C that each path satisfies. The use of the implicit function theorem allows to walk over this manifold by using the gradient of the specular constraint ∇C and a Newton type solver. This action is referred as Manifold Walk and has been applied in a standard pathtracing context by Hanika et al. in Manifold Next Event Estimation (MNEE) [HDF15] in order to construct paths connecting an observed point O to a point S sampled on a light source only visible through specular interactions. This technique is only able to find one solution and is thus limited to specular objects that are regular enough.

The Specular Manifold Sampling technique (SMS) of Zeltner et al. [ZGJ20] generalizes MNEE to the cases

where the geometry of the specular objects become more complex and more than one path between O and S may exist. This stochastic method allows unbiased rendering of caustics while staying in a standard Monte Carlo pathtracing context. Note that this method won’t produce more physically accurate caustics than standard pathtracing, but it will converge to the result faster.

Our present work will make use of the same context described in [ZGJ20], with a problem specific solver based on Fermat’s principle instead of manifold walk (Fermat Next Event Estimation, FNEE). As illustrated by Fig.2, we chose to compare our method only to Specular Manifold Sampling since this technique seemed to achieve the best balance between performance and fidelity for our use case.

In the following, we will focus on the transmitted caustic, but our work can be adapted to generate the reflected caustics as well.

3 METHODOLOGY

A glass windowpane displaying surface deformations is represented by two parametric surfaces $\{\Gamma_1, \Gamma_2\}$ (Fig. 3), each described by an elevation function $h_i(u, v)$ $i \in \{1, 2\}$ such that $\Gamma_i = \{X \in \mathbb{R}^3 / X = (u, v, h_i(u, v))\}$ with $(u, v) \in [0, 1]^2$. In the following, the h_i functions are considered twice differentiable. The ray - surface intersections are resolved using sphere tracing [Har96] in order to avoid making further assumption on how the h_i are defined (bi-cubic interpolation of a heightmap e.g.).

In the most general case where the two interfaces are not flat (i.e. the $h_i(u, v)$ functions are not constant), several paths connecting O and S may exist. This results in the formation of a caustic that is strongly dependent on the geometry of the two interfaces (Fig. 1).

Our objective is to compute the direct lighting that comes through the panel. This is done by extending the well known Next Event Estimation technique [SWZ96] to the case where two refractive interfaces occlude the

observed point O from the point S sampled on a light source.




Figure 3: Transmitted ray parametrization. There is in general multiple paths connecting an observed point O to a source S .

3.1 Finding an admissible path

In the following, we will describe a stochastic procedure to construct one valid path and compute its contribution. The procedure can then be re-conducted for each path tracing samples, effectively computing the contribution of all feasible paths when the number of samples is large enough.

Fermat's principle

Fermat's principle (or *principle of least time*) states that for any given two points of space O and S , the optical length of the paths taken by a light ray between O and S is stationary.

Let $L = \int_{\mathbf{P}} \eta(s) ds$ be the optical length of a given path \mathbf{P} , with $\eta(s)$ the index of refraction of the medium at position s . According to Fermat's principle, the set of admissible paths connecting O to S satisfy $dL = 0$.

- In the case of two points in an homogeneous medium of refraction index η , the solution is uniquely given by the straight line connecting the two points.
- With a straight windowpane in between O and S , the solution is still unique, but is now composed of three line segments OX_1 , X_1X_2 and X_2S .
- If the windowpane has irregular surfaces, the solution is no longer unique (see Fig. 3). With the added hypothesis of the h_i being small and regular enough, we can still consider that the solutions are composed of three line segments.

Hence,

$$L(u_1, v_1, u_2, v_2) = \eta_e \|OX_1\| + \eta_i \|X_1X_2\| + \eta_e \|X_2S\|$$

with η_e , η_i the refractive indexes of respectively the exterior or interior medium, and (u_i, v_i) the parametric coordinates of the point X_i . The differential dL is given by:

$$dL = \eta_e \left(\frac{OX_1}{\|OX_1\|} \cdot dX_1 + \frac{X_2S}{\|X_2S\|} \cdot dX_2 \right) + \eta_i \left(\frac{X_1X_2}{\|X_1X_2\|} \cdot (dX_1 + dX_2) \right) \quad (1)$$

Since $dL = 0 \Leftrightarrow \nabla L = 0$, we have to solve a system of 4 non-linear equations of 4 unknowns (u_1, v_1, u_2, v_2) . The solutions of this system will give us the parametric coordinates of X_1 and X_2 , allowing us to construct admissible paths.

Other configurations

Note that we can derive in the same way a system of 2 equations of 2 unknowns (u_2, v_2) for solving caustics generated by light being reflected on one face of the panel (see figure 2a for example). In that case,

$$L(X_2) = \eta_e \left\| \overrightarrow{OX_2} \right\| + \eta_e \left\| \overrightarrow{X_2S} \right\|$$

$$dL = \eta_e \left(\frac{OX_2}{\|OX_2\|} + \frac{X_2S}{\|X_2S\|} \right) \cdot dX_2$$

We could do the same for an arbitrary number of specular reflections or transmissions (similarly to Manifold walk, the Hessian matrix would have a diagonal block structure). The method used to solve the system of equation and to compute paths contributions are then similar for all configurations.

Solving the system

The system is solved using Newton's method for optimisation. Newton's method is known for converging to stationary points of the objective function regardless of their nature (minimal, maximal or saddle). While that can be an issue in a general optimization context where we want to either maximize or minimize an objective function, it is a useful property in our case. Since h_i are twice differentiable, we can derive an exact analytic expression for the 4×4 Hessian matrix H from the expression of ∇L .

The method is iterative and consists (at step k) in finding a descent direction v_k using curvature information at the point $\theta_k = (u_1^k, v_1^k, u_2^k, v_2^k)$. We can then compute

the next point θ_{k+1} by moving along v_k by a step of size t_k (Eq. 2).

$$\begin{aligned} v_k &= -[H|\theta_k]^{-1} \cdot \nabla L|_{\theta_k} \\ \theta_{k+1} &= \theta_k + t_k \cdot v_k \end{aligned} \quad (2)$$

The step size t_k is found by the Armijo rule [Arm66] using a backtracking strategy (see Algorithm 1) with constants $\alpha = 0.45$ and $\beta = 0.5$.

Algorithm 1 Line search with backtracking

```

 $t_k \leftarrow 1$ 
while  $(L(\theta_k + t_k \cdot v_k) > L(\theta_k) + \alpha \cdot t_k \cdot \nabla L|_{\theta_k} \cdot v_k)$  do
     $t_k \leftarrow \beta \cdot t_k$ 
end while
 $\theta_k \leftarrow \theta_k + t_k \cdot v_k$ 

```

In practice, we limit the number of iterations to 20 and define a convergence threshold to $\|\nabla L|_{\theta_k}\| < \nabla L_\epsilon = 10^{-4}$. If the threshold is not reached after 20 iterations, we consider that the process has failed to converge.

Newton’s method convergence behavior is well known to be complex even with relatively simple functions [HSS01]. Since the h_i functions are potentially complex, finding a good initial guess is in general challenging. The initial value θ_0 is thus chosen randomly inside $[0, 1]^4$.

We also found that imposing a minimal value for t_k is generally beneficial to prevent the solver from getting stuck ($t_{min} = 0.004$).

3.2 Path contribution

After finding a potential path, we need to compute its contribution to the lighting of the observed point O .

Let $\bar{X} = (O, X_1, X_2, S)$ be the path found by the solver. The contribution of \bar{X} is given by

$$f(\bar{X}) = L_e \cdot \frac{G(O \leftrightarrow S)}{PDF(\bar{X})} \cdot BSDF_{X_1} \cdot BSDF_{X_2} \cdot V(\bar{X}) \quad (3)$$

with L_e the radiance emitted from the light source and $V(\bar{X})$ being 0 if the path \bar{X} is occluded, 1 otherwise. $BSDF_{X_i}$ is the value of the bidirectional scattering distribution function at vertex X_i , with incoming light direction $\overrightarrow{X_1 X_2}$ and outgoing light direction $\overrightarrow{X_1 O}$ (respectively, we have for $BSDF_{X_2}$ incoming direction $\overrightarrow{X_2 S}$ and $\overrightarrow{X_2 X_1}$ for outgoing direction).

The ratio $\frac{G(O \leftrightarrow S)}{PDF(\bar{X})}$ is composed of the generalized geometry factor [JM12], and the probability density of finding the solution path \bar{X} using our solver.

Generalized geometry factor

The generalized geometry factor is defined as:

$$G(X_0 \leftrightarrow X_n) = \frac{d\omega^\perp}{dA_n} = \frac{d\omega^\perp}{dA_1} \cdot \frac{dA_1}{dA_n}$$

Intuitively, it represents the tendency of a ray bundle to spread out or focus when subject to a sequence of refraction or reflection.




Figure 4: Generalized geometry factor parametrization in the case of light transmission through a glass panel.

The term $\frac{d\omega^\perp}{dA_1}$ is simply the usual Geometry factor given by:

$$G(X_0 \leftrightarrow X_1) = \frac{d\omega^\perp}{dA_1} = \frac{|N(X_0) \cdot \overrightarrow{X_0 X_1}| \cdot |N(X_1) \cdot \overrightarrow{X_1 X_0}|}{\|X_0 - X_1\|^2}$$


with $N(X_i)$ the normal vector at vertex X_i and ω^\perp the projected solid angle.

The other factor $\frac{dA_1}{dA_n}$ is then computed from ∇C , the derivative of the specular constraint of \bar{X} [JM12].

Alternatively, pencil tracing can also be used to compute G by the mean of the transfer matrices of the optical system [STN87] [KHD14].

Inverse probability estimation

As explained in [ZGJ20], the probability of sampling the solution \bar{X} corresponds to the volume of the convergence basin of this solution. For a given solution \bar{X} , its convergence basin is defined as the set of all the initial values θ_0 that converges to this solution (Fig. 5).



(a) Convergence basins at a point O for a source S for a simple blown glass profile. (b) Zoom on the solution cluster.

Figure 5: The solver converges to 3 solutions colored in red, blue and green. Each point is colored depending on which solution the solver converges to, starting from this point. Their respective convergence basins are highly irregular and nearly cover the whole space. These 4D basins are projected onto the 2D plane for visualization purpose (i.e. we choose $(u_1, v_1) \in [0, 1]^2$ and impose $u_1 = u_2$ and $v_1 = v_2$).

These convergence basins are in general highly irregular and their volume cannot be computed exactly. The inverse of their volume can however be estimated by an iterative process that resorts to counting the number of trials necessary for the solver to converge to \bar{X} from a random initial starting point.

$$\frac{1}{PDF(\bar{X})} \simeq N_{trial}$$

In practice, it is necessary to clamp this estimator to N_{max} to avoid potentially infinite loops if the basins are close to being singular (typical value $N_{max} = 1000$).

3.3 Initialisation strategy

In this section, we discuss the initialisation strategy we use for selecting an appropriate $\{\Gamma_1, \Gamma_2\}$ surface pair given two points O and S that we try to connect. Suppose that multiple specular surfaces are present in the scene. In the general case where the interfaces can have any shape, we can't easily choose an appropriate couple $\{\Gamma_1, \Gamma_2\}$ for initializing the solver. We still can resort to selecting a random pair for each sample, which is the strategy used by Specular Manifold Sampling: select a *caustic caster* marked shape at random, then launch one estimate per *caustic bouncer* marked shape. Unfortunately, in practice, most pairs lead to no solutions and are thus bad initialisation choices (see Fig. 6).




Figure 6: Choice of an interface pair: most combinations lead to a poor initialisation (red paths) that prevents the solver from converging to a solution.

In our case, having to handle many windows in a scene, we can make a better choice.

Consider that we have N glass panes. For each glass pane g_i , the most natural pairing would be to use $\Gamma_1^{g_i}$ and $\Gamma_2^{g_i}$ the front and back faces of the glass pane, thus reducing the number of initialisation choices from $4N^2$ (since there is $2N$ faces) to N .

Since the caustic generated by the window is mostly contained inside the shadow cast by a flat glass pane (that we will refer to as the *approximate caustic area* in the following), we can further reduce the number of possibilities by selecting the glass pane that occluded the shadow ray cast from O in the direction of S (see point O in Fig. 7). That generally leaves us with a unique initialisation possibility. In the contrary case, we resort to standard pathtracing (see point O').




Figure 7: Window selection during light sampling. The *approximate caustic area* is colored in light yellow.

This strategy however produces unnaturally sharp shadows (see Fig. 8b) for occluders situated in front or behind the chosen window (see point O'' in Fig. 7). In practice, the caustic tends to bleed in a small region outside of the approximate caustic area (see Fig. 8 a). We

can easily account for this case by choosing to launch the solver with a random glass pane whenever an object occludes the shadow ray. This may come at the cost of a loss of performance (see Fig. 8 c). This solution isn't entirely satisfactory and more sophisticated approach [WHY20] may be considered in future work.

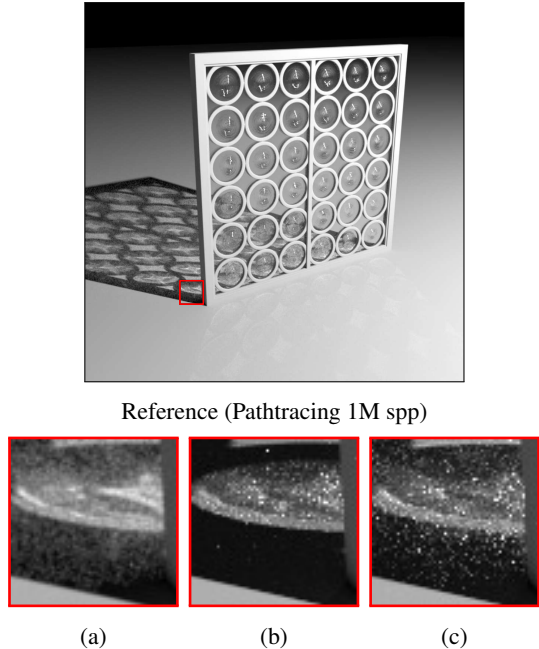


Figure 8: Illustration of the bias introduced by the two initialisation strategies: (a) Reference, (b) FNEE with sharp shadows (1000 sec), (c) FNEE with caustic bleed (1000 sec).

4 RESULTS

Solver comparison

In this section, we compare our solver against the predictor-corrector scheme used by Specular Manifold Sampling. For comparison fairness, the two solvers are implemented inside the same integrator in Mitsuba2 [Nim+19] and thus use the initialisation scheme described in the previous section. Both solvers use double precision arithmetic which is necessary to deal with thin windows (thickness around 1mm).

The images produced by FNEE and SMS were compared against a ground truth image generated by pathtracing. The Mean Squared Error (MSE) metric was used:

$$MSE = \frac{1}{n} \sum_{i=0}^n (x_i - \hat{x}_i)^2$$

with x_i the pixels values of the ground truth image (generated by pathtracing) and \hat{x}_i the pixels values of the image being compared.

Our benchmark scene (Fig. 9) is a window composed of 4 different glass panes with various elevation profiles

on the face facing the camera. These profiles are modeled from typical windows of the XIV to XIX centuries (crown and cylinder blown sheet glass).




Figure 9: Visual comparison of our method (FNEE) with Specular Manifold Sampling (SMS) on various perturbation profiles (equal time $t = 100s$). Reference computed with Pathtracing with 10M samples per pixels.




Figure 10: Convergence rate comparison of the two solvers (ours in blue, SMS in green). Using the scene from Fig.9, we compare images computed using each solver to the reference image at various rendering times and report the corresponding MSE error measure.

Our Fermat-based solver achieves noticeably faster convergence on all deformation profiles of the benchmark scene (Fig. 9 and Fig. 10). Both solvers are prone to producing outliers that would need to be eliminated in post treatment, or with a robust Monte Carlo estimator [ZHD18] [BDR21].

Solver	h_{eval}	<i>Solutions</i>	<i>Success</i>
FNEE	956M	76832	3.76%
SMS	8457M	48754	1.87%

Table 1: Comparison of the two solvers for various performance metrics (same scene as Fig. 9 with image size of 256×256 px and 100 spp): number of elevation function evaluations (in millions), number of unique solutions discovered, solver success probability (ratio of the number of solver calls that converges to a solution over the total number of calls).

More in depth comparisons (Tab. 1) reveals that the rendering of caustics produced by thin glass panels is particularly challenging for the manifold walk solver used by SMS. Our formulation displays an around 8 time

smaller number of elevation function calls and a higher success probability (we compute the success probability for each solver as the ratio of the number of solver calls that converges to a solution over the total number of calls registered). This allows FNEE to explore the solution space more efficiently than SMS (larger number of solutions discovered).

Initialisation method comparison

Here, we compare our initialisation strategy to the random specular shape sampling used by SMS on a restitution scene (work in progress) from the Digital Field of Cloth of Gold project. This indoor scene (Fig. 11) is lighted by a distant spherical light source simulating sunlight. The sun shines through many old glass windows, producing characteristic caustics.




Figure 11: Visual comparison of our method with SMS (equal time $t = 100s$). Reference uses pathtracing with 30M samples per pixel.

The large number of specular interfaces makes the random pairing initialisation particularly inefficient. In contrast, the convergence speed of the proposed method remains practically constant with the number of windows used. This particularly highlights the need for a robust initialisation strategy for these types of scenes.

5 DISCUSSION

Both *FNEE* and *SMS* share the same difficulties dealing with glass panels with intricate, high frequency details, where the invPDF estimation process becomes performance intensive. Using the above parameters ($N_{newton} = 20$, $\nabla L_{\epsilon} = 10^{-4}$, $t_{min} = 0.004$) with finely detailed profiles leads to unnaturally dark caustics (Fig. 12).





Figure 12: invPDF under-estimation in the case of a finely detailed glass profile (using our method *FNEE*).

Understandably, as the surface gets more detailed, the generalized geometric term $G(O \leftrightarrow S)$ tends to become increasingly small in certain area (being directly dependent to the partial derivatives of the normal vector $\frac{\partial n}{\partial u}$ and $\frac{\partial n}{\partial v}$). As the ray contribution is proportional to the product $G(O \leftrightarrow S) \cdot invPDF(\bar{X})$, if the allowed number of iterations N_{max} during the invPDF estimation process is too low, the invPDF term generally won't be large enough to compensate for the luminosity loss caused by $G(O \leftrightarrow S)$. This results in an underestimation of the illumination.

Since the convergence basins are generally large for smooth perturbation profiles, the estimation can be done in a few iterations. It is however not the case anymore when the perturbations are finer since the basins become increasingly small (Fig. 13) as the number of solutions dramatically increases. In this case, taking a sufficiently large N_{max} may lead to an impractically slow convergence rate for both *SMS* and *FNEE*.

6 CONCLUSION

We have presented a method to compute the transmitted caustic produced by a displaced glass panel lighted by a light source. Our method based on the Fermat's principle displays faster convergence than *SMS* and allows to handle efficiently scenes including many glass panes. These scenes are commons in architecture and represent a situation for which *SMS* initialisation scheme is inefficient. The two methods being closely related, they share the same difficulty dealing with finely detailed surfaces: since many solutions exists, the estimation of

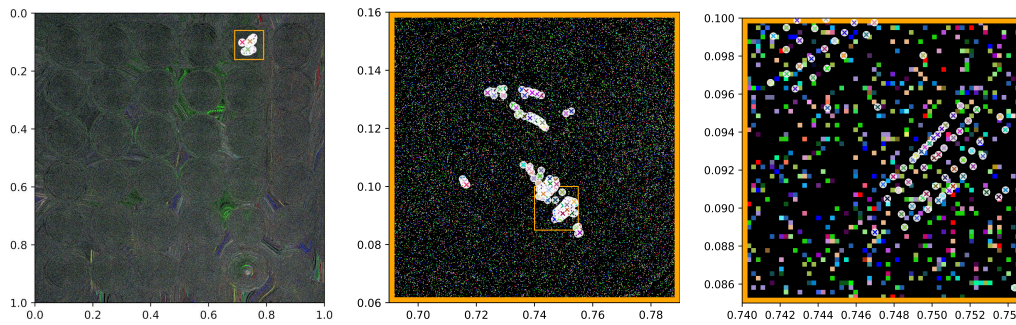


Figure 13: Convergence basins for a crown glass window. The finely detailed surface of the panel gives rise to several clusters of solutions. Each individual solution has a small and irregular convergence basin, making the invPDF estimation process inefficient.

the invPDF term requires many iterations to converge. Finding a better way to estimate this term is an avenue for future work and would greatly benefit the two techniques. Our present work is still a simplification of the real world problem since we neglected the volume irregularities that exist in real ancient glass panels. Finding ways to take these irregularities into account would be a natural direction for future research.

7 ACKNOWLEDGMENTS

We would like to thank the Hauts-de-France region, which is helping to finance this work through Quentin Huan’s PhD grant. The Digital Field of Cloth of Gold is a project of the Université Lille Nord-Europe granted by the I-SITE ULNE foundation.

REFERENCES

- [Arm66] Larry Armijo. “Minimization of functions having Lipschitz continuous first partial derivatives”. In: *Pacific Journal of Mathematics* 16.1 (Jan. 1966), pp. 1–3. ISSN: 0030-8730, 0030-8730. DOI: 10.2140/pjm.1966.16.1.
- [BDR21] Jérôme Buisine, Samuel Delepouille, and Christophe Renaud. “Firefly Removal in Monte Carlo Rendering with Adaptive Median of means”. In: *Eurographics Symposium on Rendering - DL-only Track* (2021). Artwork Size: 12 pages ISBN: 9783038681571 Publisher: The Eurographics Association Version Number: 121-132, 12 pages. ISSN: 1727-3463. DOI: 10.2312/SR.20211296.
- [GNL20] Lars Oliver Grobe, Andreas Noback, and Franziska Lang. “Data-Driven Modelling of Daylight Scattering by Roman Window Glass”. In: *Journal on Computing and Cultural Heritage* 13.1 (Feb. 2020), pp. 1–20. ISSN: 1556-4673, 1556-4711. DOI: 10.1145/3350428.
- [Har96] John C. Hart. “Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces”. In: *The Visual Computer* 12.10 (Dec. 1996), pp. 527–545. ISSN: 01782789. DOI: 10.1007/s003710050084.
- [HDF15] Johannes Hanika, Marc Droske, and Luca Fascione. “Manifold Next Event Estimation”. In: *Computer Graphics Forum* 34.4 (July 2015), pp. 87–97. ISSN: 0167-7055, 1467-8659. DOI: 10.1111/cgf.12681.
- [HJ09] Toshiya Hachisuka and Henrik Wann Jensen. “Stochastic progressive photon mapping”. In: *ACM SIGGRAPH Asia 2009 papers on - SIGGRAPH Asia '09*. Yokohama, Japan: ACM Press, 2009, p. 1. ISBN: 978-1-60558-858-2. DOI: 10.1145/1661412.1618487.
- [HSS01] John Hubbard, Dierk Schleicher, and Scott Sutherland. “How to find all roots of complex polynomials by Newton’s method”. In: *Inventiones mathematicae* 146.1 (Oct. 2001), pp. 1–33. ISSN: 0020-9910, 1432-1297. DOI: 10.1007/s002220100149.
- [Jen96] Henrik Wann Jensen. “Global Illumination using Photon Maps”. In: *Rendering Techniques '96*. Ed. by Xavier Pueyo and Peter Schröder. Series Title: Eurographics. Vienna: Springer Vienna, 1996, pp. 21–30. ISBN: 978-3-211-82883-0 978-3-7091-7484-5. DOI: 10.1007/978-3-7091-7484-5_3.
- [JM12] Wenzel Jakob and Steve Marschner. “Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport”. In: *ACM Trans. Graph.* 31.4 (July 2012). Place: New York, NY, USA

- Publisher: Association for Computing Machinery. ISSN: 0730-0301. DOI: 10.1145/2185520.2185554.
- [Jr+09] Joseph T. Kider Jr. et al. “Recreating Early Islamic Glass Lamp Lighting”. In: *VAST: International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. Ed. by Kurt Debattista et al. ISSN: 1811-864X. The Eurographics Association, 2009. ISBN: 978-3-905674-18-7. DOI: 10.2312/VAST/VAST09/033-040.
- [Kaj86] James T. Kajiya. “The rendering equation”. In: *ACM SIGGRAPH Computer Graphics* 20.4 (Aug. 1986), pp. 143–150. ISSN: 0097-8930. DOI: 10.1145/15886.15902.
- [KHD14] Anton S. Kaplanyan, Johannes Hanika, and Carsten Dachsbacher. “The natural-constraint representation of the path space for efficient light transport simulation”. In: *ACM Transactions on Graphics* 33.4 (July 2014), pp. 1–13. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/2601097.2601108.
- [LW93] Eric P. Lafortune and Yves D. Willems. “Bi-directional path tracing”. In: *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)*. Alvor, Portugal, Dec. 1993, pp. 145–153.
- [MH92] Don Mitchell and Pat Hanrahan. “Illumination from Curved Reflectors”. In: *SIGGRAPH Comput. Graph.* 26.2 (July 1992). Place: New York, NY, USA Publisher: Association for Computing Machinery, pp. 283–291. ISSN: 0097-8930. DOI: 10.1145/142920.134082.
- [Nim+19] Merlin Nimier-David et al. “Mitsuba 2: a retargetable forward and inverse renderer”. In: *ACM Transactions on Graphics* 38.6 (Dec. 2019), pp. 1–17. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/3355089.3356498.
- [Shi90] Peter Shirley. “A ray tracing method for illumination calculation in diffuse-specular scenes”. In: *Proceedings of Graphics Interface '90*. GI 1990. ISSN: 0713-5424 event-place: Halifax, Nova Scotia, Canada. Toronto, Ontario, Canada: Canadian Information Processing Society, 1990, pp. 205–212. DOI: 10.20380/GI1990.25.
- [STN87] Mikio Shinya, T. Takahashi, and Seiichiro Naito. “Principles and Applications of Pencil Tracing”. In: *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '87. New York, NY, USA: Association for Computing Machinery, 1987, pp. 45–54. ISBN: 0-89791-227-6. DOI: 10.1145/37401.37408.
- [SWZ96] Peter Shirley, Changyaw Wang, and Kurt Zimmerman. “Monte Carlo techniques for direct lighting calculations”. In: *ACM Transactions on Graphics* 15.1 (Jan. 1996), pp. 1–36. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/226150.226151.
- [VG97] Eric Veach and Leonidas J. Guibas. “Metropolis light transport”. In: *Proceedings of the 24th annual conference on Computer graphics and interactive techniques - SIGGRAPH '97*. Not Known: ACM Press, 1997, pp. 65–76. ISBN: 978-0-89791-896-1. DOI: 10.1145/258734.258775.
- [WHY20] Beibei Wang, Miloš Hašan, and Ling-Qi Yan. “Path cuts: efficient rendering of pure specular light transport”. In: *ACM Transactions on Graphics* 39.6 (Dec. 2020), pp. 1–12. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/3414685.3417792.
- [ZGJ20] Tizian Zeltner, Iliyan Georgiev, and Wenzel Jakob. “Specular manifold sampling for rendering high-frequency caustics and glints”. In: *ACM Transactions on Graphics* 39.4 (Aug. 2020). ISSN: 0730-0301, 1557-7368. DOI: 10.1145/3386569.3392408.
- [ZHD18] Tobias Zirr, Johannes Hanika, and Carsten Dachsbacher. “Re-Weighting Firefly Samples for Improved Finite-Sample Monte Carlo Estimates”. In: *Computer Graphics Forum* 37.6 (2018), pp. 410–421. DOI: <https://doi.org/10.1111/cgf.13335>.