Progressive Photon Relaxation

BEN SPENCER
and
MARK W. JONES
Visual and Interactive Computing Group, Swansea University

We introduce a novel algorithm for progressively removing noise from view-independent photon maps while simultaneously minimizing residual bias. Our method refines a primal set of photons using data from multiple successive passes to estimate the incident flux local to each photon. We show how this information can be used to guide a relaxation step with the goal of enforcing a constant, per-photon flux. Using a reformulation of the radiance estimate, we demonstrate how the resulting blue noise photon distribution yields a radiance reconstruction in which error is significantly reduced. Our approach has an open-ended runtime of the same order as unbiased and asymptotically consistent rendering methods, converging over time to a stable result. We demonstrate its effectiveness at storing caustic illumination within a view-independent framework and at a fidelity visually comparable to reference images rendered using progressive photon mapping.

Categories and Subject Descriptors: I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Animation*; I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling—*Physically based modeling*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Photon Mapping, Blue Noise, Photon Relaxation, Global Illumination

1. INTRODUCTION

A caustic is the term given to the component of global illumination in which part of the light transport path contains a S^+D sub-path before reaching the eye. A commonly cited example is light that has been refracted through turbulent water forming characteristic rippled patterns on the bottom of a lake or pool. Other, less distinct examples include the illumination in a room lit by a glass light bulb or sunlight shining through a window onto a wall.

The space of all caustic paths exhibits a very high energy variance since only a fraction conveys any significant contribution from the light to the eye. As a consequence, locating these paths with a minimum of both time and error has presented a significant technical challenge. Bi-directional algorithms based on ray tracing have proven to be the most effective, and both biased and unbiased solutions are now widely adopted. Among the most popular is the photon mapping method [Jensen 1996]: a two-pass, consistent, particle-based algorithm which has become a mainstay in many production environments owing to its simplicity, versatility and speed.

Although photon mapping remains one of the fastest and most computationally efficient methods of synthesizing caustics, the relatively low resolution of the photon dataset means the resulting rendered images often lack the high visual fidelity of those generated by algorithms based on unbiased Monte Carlo integration. This is unsurprising given that such methods solve the rendering equation independently at each pixel by progressively sampling the global path space. Though inevitably costly, subtle details are correctly resolved and noise is guaranteed to vanish. Furthermore, unbiased Monte Carlo methods run within a near-constant memory bound, irrespective of the number of samples. By comparison, photon mapping exploits local integral correlation using a static, sparse point distribution in order to reconstruct exitant radiance. This approach is very fast since the relatively low number of photons means the dataset can be compiled quickly. However, discrepancy and noise in the point set become problematic when rendering S^+D caustic illumination since the photon map needs to be visualized directly.

More recent work has sought to address these limitations through the development of iterative density estimation with a progressively shrinking kernel bandwidth [Hachisuka et al. 2008]. This progressive solution to photon mapping allows exitant radiance to be reconstructed from very large numbers of incident photons without the need to store the entire dataset all at once. For scenes dominated by caustic illumination, this approach is particularly powerful. However, since the location of radiance estimates must be known *a priori* to photon propagation, it is an intrinsically view-dependent technique.

In this paper, we seek to combine the favorable qualities of view-independent photon mapping with those of unbiased or asymptotically consistent Monte Carlo methods. We accomplish this by using multiple passes of incident photons to progressively refine an initially noisy photon distribution. We employ a novel redistribution strategy based on point relaxation to equalize the incident radiant flux arriving over the Voronoi cell corresponding to each photon in the map. Our approach has an open-ended runtime similar to progressive Monte Carlo methods, but results in a view-independent photon map that can be re-rendered from arbitrary viewpoints. This yields the following principal advantages:

—The view-independent photon map generated by our method exhibits a blue noise spectral signature [Ulichney 1988] that is ideally suited to low-bandwidth kernels. We can achieve high-quality results with between 10- and 20-nearest neighbors re-

DOI 10.1145/1559755.1559763
http://doi.acm.org/10.1145/1559755.1559763

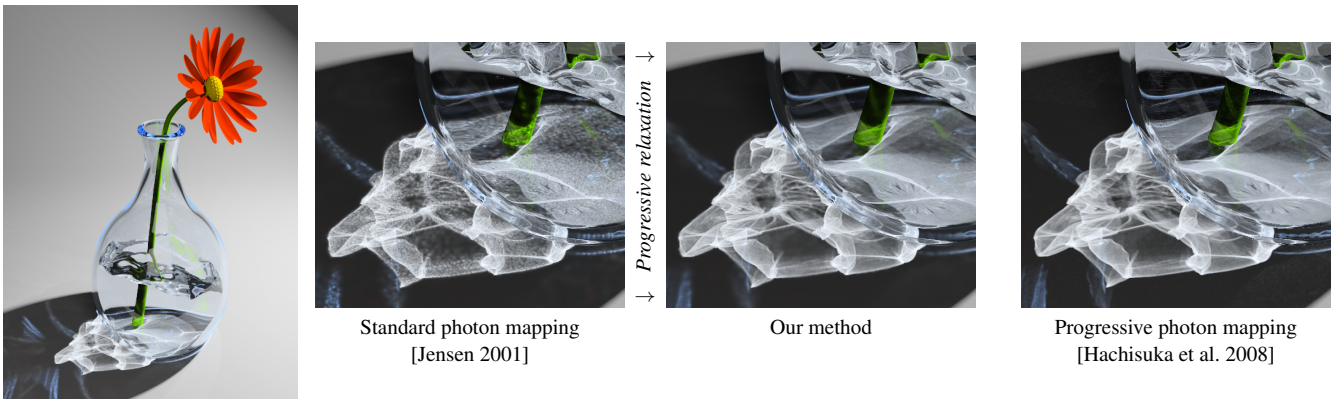

Fig. 1. A flower in a glass vase backlit by a strong point light source. Second from left: an unmodified photon map with 500k photons, rendered using 20-nearest neighbors in the radiance estimate. Second from right: using the map from the previous frame, our method progressively refines the distribution with data from 10M additional photons. These ancillary photons are cast individually and destroyed once their contributions have been accumulated, making the process memory efficient. Far right: the same scene rendered using progressive photon mapping and halted after 10M photons (a runtime equivalent to our method).

sulting in minimal levels of bias (Figure 1), very fast photon map queries and a corresponding reduction in rendering time.

—Unlike the method of photon redistribution described in Spencer and Jones [2009], our progressive method converges over time to a stable result. Furthermore, no explicit feature detection needs to be implemented.

2. RELATED WORK

Physically based methods of synthesizing caustics can be divided between biased and unbiased methods. Kajiya's [1986] pioneering work on the rendering equation led to the development of a generalized formulation to the global illumination problem. This application of Monte Carlo methods to solve the recursive integral of reflected radiance led to the introduction of path tracing as the first unbiased global illumination algorithm. While conceptually simple, Kajiya's solution is slow to converge when a significant proportion of the total energy arrives over a restricted set of paths. Lafortune and Willems [1993] and Veach and Guibas [1994] independently addressed this issue with the introduction of bi-directional path tracing (BDPT), thereby greatly lowering estimator variance in many scenes.

Despite improvements in importance sampling over multiple sub-paths [Veach and Guibas 1995], bi-directional path tracing still performs poorly under certain conditions. Veach and Guibas [1997] responded to these problems with the introduction of Metropolis light transport (MLT), an algorithm based on the Metropolis-Hastings approach to solving complex, multidimensional integrals through selective mutation and rejection. Later, Cline et al. [2005] proposed a variation based on a selective redistribution of energy between correlated integrals (in this case, nearby pixels). Lai et al. [2007] expanded on this concept by creating an initial pool of paths across the image domain and using them to guide successive iterations.

A second class of global illumination algorithms exploits integral correlation in a biased way such that the expected value of the estimator does not equal the function being estimated. These methods reduce variance and increase speed but often at the expense of visual accuracy. Photon mapping [Jensen 2001] uses a flux propagation pass to cache incident illumination which is later referenced by the rendering pass. The exitant radiance at a point is reconstructed from the local photon distribution using k-nearest neighbor (k-NN) density estimation. Although photon mapping is a general global illumination solution, its use of a sparse point dataset leaves it vulnerable to visible error due to point discrepancy as well as bias from the kernel.

A wide variety of algorithms have been proposed to deal with the problem of noise in the radiance estimate. Jensen and Christensen [1995] suggested selecting an optimal kernel bandwidth through the use of differential checking. Schregle [2003] later proposed a more sophisticated bandwidth selector guided by the observation that the variation in noise over a random photon map conforms to a normal distribution. Further research focuses on complementing intelligent bandwidth selection by dynamically adjusting the filter support of the kernel, thereby minimizing proximity and boundary bias. Schjøth et al. [2005] proposed deriving a structure tensor from the local distribution in order to constrain smoothing perpendicular to the illumination gradient. Later, Schjøth et al. [2007] adapted Igehy's [1999] ray differential framework to store additional information about photon footprints, thereby greatly increasing the fidelity of highly focused caustics.

The photon relaxation method [Spencer and Jones 2009] addresses the problem of bias-variance trade-off by directly removing noise from the underlying photon distribution. The algorithm runs as two additional passes between photon propagation and rendering: feature detection and point relaxation respectively. Feature detection analyzes the photon distribution and isolates those points which lie near edges, boundaries and other important visual cues. During the relaxation step, each photon is repelled from its neighbors with a force proportional to their relative proximity. After a sufficient number of iterations, the spatially adjacent neighbors for any given photon become equidistant and, as a consequence, noise is diffused away. In regions of high-frequency detail, migration is constrained by feature bases defined in the previous pass. The resulting distribution has a blue noise spectral signature, one which has been shown to be an optimal sample pattern in many areas of computer graphics [Ulichney 1988], including photon mapping.

One of the advantages of photon relaxation is that noise is removed once as a pre-pass rather than at each radiance estimate.

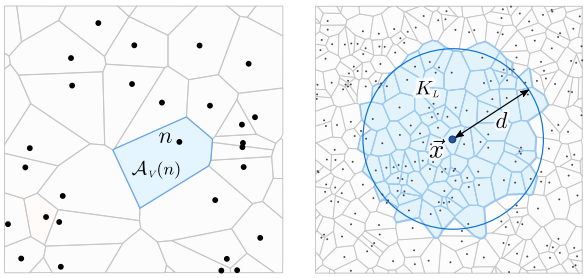

(a) (b)

Fig. 2. (a) The Voronoi tessellation links flux density with the discrete set of photon sites. (b) A k-NN estimator reconstructs exitant radiance area using the disk of radius d spanning the subset of K_L-nearest photons local to the query point $\vec{x}$.

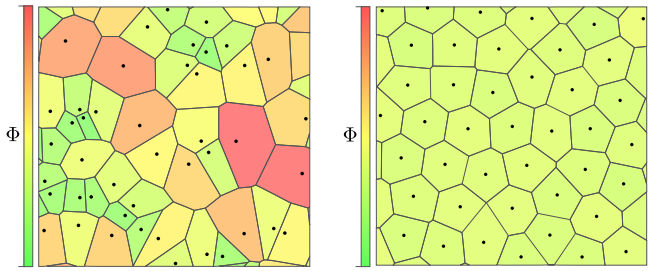

(a) Heterogeneous flux (b) Homogeneous flux

Fig. 3. A constant PDF encoded using both a stochastic and a blue noise photon distribution. (a) Stochastic sampling results in high variance in cell area which must be counterbalanced by high-variance photon flux. (b) A blue noise distribution results in low variance in cell area and thereby exhibits near-homogeneous photon flux.

This allows for the use of very low-bandwidth kernels and corresponds to a significant reduction in photon map rendering cost.

More recently, Hachisuka et al. [2008] introduced progressive photon mapping: a hybrid global illumination solution that retains the robustness of Jensen's original two-pass method while asymptotically converging to the correct result. Motivation for this algorithm stems from the observation that scenes dominated by caustic illumination are difficult to sample using Monte Carlo integration, and prohibitively expensive to store as a photon map due to the high memory overhead. Progressive photon mapping stores an array of primary and specular ray hit points which are used to gather flux from small sets of propagated photons. Photons landing within the radius of a hit point are captured and accumulated and the hit point's radius decreased accordingly. After a sufficiently large number of photons have been absorbed, noise and bias vanish. Although progressive photon mapping is not unbiased, the algorithm is consistent, converging asymptotically to the correct solution in the limit. The authors later expanded their method to support efficient sampling across multiple domains [Hachisuka and Jensen 2009], thus allowing rapid resolution of distributed ray tracing effects such as glossy reflection and depth-of-field blur. Knaus and Zwicker [2011] also proposed a formulation of progressive photon mapping that does not require the storage of local statistics and generalizes the technique to arbitrary kernels. Jarosz et al. [2011] proposed a progressive adaptation of photon beams to render scattering in participating media. More recently, Doidge et al. [2012] suggested mixing progressive photon mapping with Monte Carlo path tracing to address shortcomings inherent to both techniques.

3. RECONSTRUCTING THE ILLUMINATION FUNCTION USING THE PHOTON MAP

The photon map as defined by Jensen [2001] encodes the function of reflected radiance, L, over the combined space over all illuminated surfaces, σ, using propagated packets of energy which are cached as a point cloud. An estimate of L is typically reconstructed from a subset of photons local to a query point, $\vec{x}$. Since photons are effectively Dirac delta functions of radiance, the estimator must smooth the sparse point data over some interval, thereby allowing the function to be sampled continuously.

3.1 The Radiance Estimator

The *de facto* method of estimating L is through k-nearest neighbor density estimation. A disk of radius d is expanded around $\vec{x}$ until it contains K_L photons. The flux of each photon in this subset is multiplied by the surface BRDF, f_r, and the reciprocal of the area of the disk. Summing over all K_L yields the estimate, $\hat{L}$:

$$\hat{L}(x \to \vec{\omega}) = \frac{1}{\pi d^2} \sum_{i=1}^{K_L} f_r(x, \vec{\omega} \leftrightarrow \vec{\omega}_p(i)) \Delta\Phi_p(i)\,, \qquad (1)$$

where $\vec{\omega}$ is the exitant ray direction and $\vec{\omega}_p$ is the trajectory of the incident photon.

The bandwidth of a k-NN estimator can be adjusted by varying the size of K_L and should ideally be as small as possible so as to minimize bias from the span of the kernel. However, as the number of photons in K_L decreases, variance and discrepancy between each element translates into variance in the radiance estimate.

We explain the concept of variance between photons in Figure 2a where a uniform function of flux density, B, over the domain σ is encoded by a set of N photons seeded stochastically according to a continuous uniform distribution. By visualizing the Voronoi tessellation, V, of the underlying point cloud, we can link the discrete, finite set of photon sites with the continuous flux density function from which they are derived. Intuitively, the region of space associated with a single photon, n, is the area of its Voronoi cell, $\mathcal{A}_V(n)$.

It is apparent from this example that the difference between cell areas over V varies considerably. Since the flux density function we are trying to reconstruct is constant and each photon is initialized with a constant flux, $\Delta\Phi_p$, the area, $\mathcal{A}_V$, should in turn be uniform. In practice, stochastic seeding translates into variable cell areas and hence variance in the k-NN estimator.

A good strategy to reduce discrepancy between photons is to sample the space of the emitter using a quasi-random sequence [Keller 1995]. This approach is effective when photons are reflected and transmitted specularly, however, scattering BRDFs and complex area emitters scramble coherent photon paths resulting in randomization in the projected distribution.

3.2 Reducing Kernel Variance

Increasing the bandwidth of the k-NN radiance estimator effectively smooths out noise due to discrepancy, as illustrated in Figure 2b. As K_L grows larger, the disk of the kernel more closely approximates the sum of $\mathcal{A}_V(n)$ in K_L. The drawback of large values of K_L is loss of detail due to increased kernel bias. If this is to be avoided, the resolution of the photon map must also be increased to compensate. This strategy is demonstrated in Figures 4a and 4b where the number of stored photons and the size of K_L is increased 10-fold

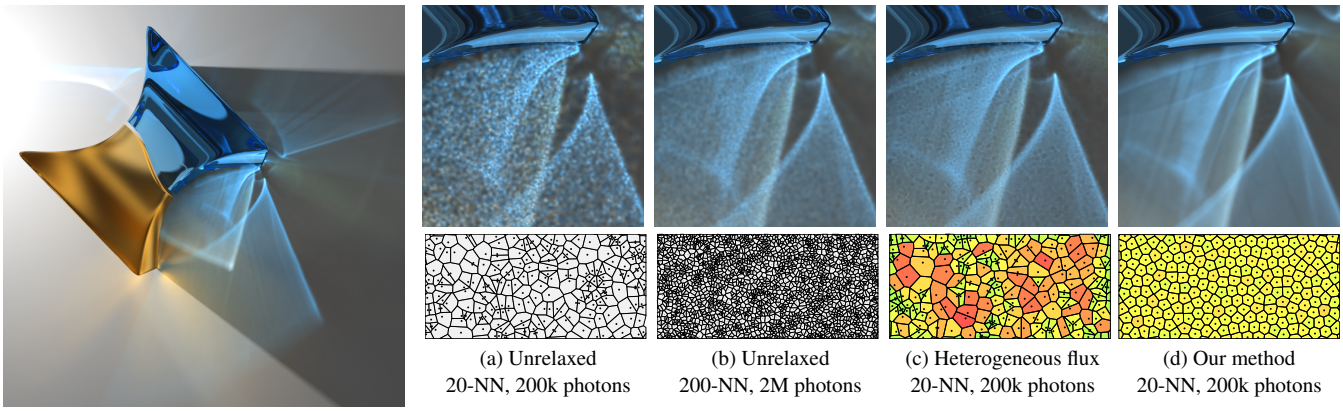

(a) Unrelaxed 20-NN, 200k photons | (b) Unrelaxed 200-NN, 2M photons | (c) Heterogeneous flux 20-NN, 200k photons | (d) Our method 20-NN, 200k photons

Fig. 4. A complex dielectric object generating both specular and glossy caustics. (a) An unrelaxed photon map exhibits high levels of noise. (b) Noise reduction by using an order of magnitude more photons in both the map and k-NN kernel. (c) Integrating flux density over each photon counterbalances variance but leaves residual error. (d) Our method relaxes the photon map thereby minimizing error due to discrepancy.

resulting in a reduction in noise while preventing any increase in smoothing due to bias.

Whereas the bias of an estimator increases in proportion to the square of the bandwidth, k^2, the variance diminishes according to the inverse relationship $\frac{1}{k}$ [Silverman 1986]. Thus, reducing the variance by half incurs a four-fold increase in bias which must be compensated by a similar rise in photon storage. This geometric increase in both storage and query time makes eliminating noise by increasing photon density impractical as a general solution. Hachisuka et al. [2008] show that, even when tens of millions of photons are stored, bias and noise often preclude the effective use of the photon map.

Given that increasing the number of stored photons is untenable, we instead consider how a stochastically seeded photon distribution can instead be "optimized", thereby reducing noise without the geometric storage increase. The source of noise, as discussed in Section 3.1, is shown to be variance in Voronoi cell area coupled with a homogeneous value of $\Delta\Phi_p$ across all photons. The first optimization we propose is to counterbalance variance in $\mathcal{A}_V$ by accurately estimating the integral of flux density over the domain of the cell of each photon. This approach shares conceptual similarities with Suykens' and Willems' [2000] solution to reducing photon storage in regions that reach a specified density.

Figure 3a illustrates the concept diagrammatically by encoding a uniform flux density function using a randomly seeded distribution of photons. In this example, the color of each cell associated with photon n corresponds to the relative magnitude of the estimated flux, $\Delta\Phi_p(n)$. Photons shaded in red encode more energy than those shaded in green and yellow, in this case because the larger areas accumulate more of the energy from the flux density function.

We demonstrate the effect of heterogeneous flux on the radiance estimate in Figure 4c. Here, the integral of B is solved progressively using the Monte Carlo method discussed in Section 4.1. Crucially, the number of stored photons is the same as Figure 4a. Despite no increase in photon density, compensating for local Voronoi cell variance by accurately computing $\Delta\Phi_p$ significantly reduces noise.

The limitation of this approach is visible as a distinctive mottling effect which remains even when the estimates of $\Delta\Phi_p$ are allowed to converge. These artifacts appear due to the circular approximation of the k-nearest Voronoi cell areas not precisely counterbalancing the Monte Carlo integral of flux (Figure 2b). Since the cell associated with each photon is not an explicitly defined primitive, compensating for this error by precisely computing its physical area is expensive. More significantly, irregular cells adapt poorly to changes in the flux density function and converge at variable rates. Hence, to minimize kernel artifacts we consider a new approach that is better suited to the circular kernel used by the k-NN radiance estimate.

3.3 Photon Redistribution

In order to minimize the mottling artifacts described above, we define a means of physically manipulating individual photon Voronoi cells so that their occupied area can be more accurately computed by the density estimation kernel. A blue noise distribution is optimal for this purpose since the variance in cell area is locally minimized. Opting for such a strategy introduces the scope for very low-bandwidth kernels, a desirable feature since both bias and the cost of k-nearest neighbor density estimation at render time remain low. Consequently, our goal is to relax the photon map by enforcing a capacity constraint, aiming to make the integral of flux over each photon homogeneous while allowing the area to change. A diagrammatic example is depicted in Figure 3b. Notice how each photon is at near-constant flux since the area of its Voronoi cell has been adjusted to match those of its neighbors.

In order to achieve this property we require a robust method for precisely controlling the area of each photon's Voronoi cell. While this is not a trivial task, the problem of producing good blue noise distributions has received considerable attention from the research community. Lloyd's method [1983] is one popular algorithm based on iterative relaxation. Given an initial point set, every member is repeatedly moved to the centroid of its corresponding Voronoi cell, the average migration distance attenuating as the number of iterations increases. Later, Secord [2002] successfully adapted the algorithm for non-uniform density functions by weighting each cell so that its area is proportional to the value of the function at that point. Balzer et al. [2009] formalized this concept with the definition of capacity constraint as a means to avoid the isotropic patterns that appear in sets converged using Lloyd's method. Succinctly, the optimal capacity, c_V, of a Voronoi cell associated with point i should be:

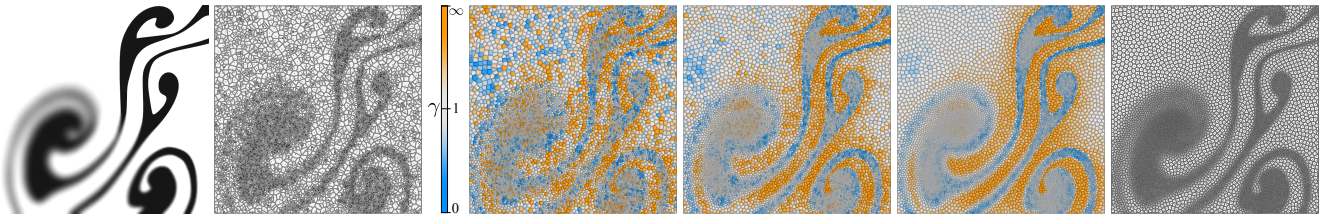

Fig. 5. Progressive photon relaxation. (a) The probability density function of B (inverted for clarity). (b) Stochastic sampling results in a high-discrepancy, persistent photon dataset. (c) Progressive refinement of the flux capacity of each cell. Here, orange represents values of γ that are greater than 1 and blue, values that are less than 1. Residual flux imbalance is due to diffusion bias injected by the relaxation kernel, however, this is corrected by the density estimation kernel. (d) After relaxation, the distribution exhibits a characteristic blue noise signature.

$$c_V(i) = \frac{\int_\sigma \rho(x)da}{g}, \tag{2}$$

where ρ is the bounded density function over domain σ, and g is the number of cells. By enforcing this rule for every point during relaxation, the distribution converges to a blue noise distribution of extremely high quality.

Directly applying similar techniques to redistribute the photon map is impractical. Balzer et al.'s method implements a discrete space adaptation of Lloyd's method which uses a high-density point superset to enforce the capacity of each Voronoi cell. While this could theoretically be extended to 2.5-dimensional photon space, the memory overhead required to store a fine-grained sampling of ρ at the photon level would be prohibitive. Furthermore, the running time for the capacity-constrained method given in Balzer et al. [2009] is $\mathcal{O}(g^2 + gh\log\frac{h}{g})$ for g cells and h points. Given that caustic photon maps containing in excess of 1,000,000 photons are not uncommon, this geometric running time becomes problematic. Despite these drawbacks we can see that a capacity-constrained photon distribution over incident flux would be highly desirable. In the following sections we propose a method to achieve flux capacity constraint progressively, thereby avoiding excessive storage overheads.

4. OUR METHOD - PROGRESSIVE PHOTON RELAXATION

The goal of our algorithm is to enforce a capacity constraint similar to Equation 2 over all photons in N. The density function, ρ, becomes the function of radiant flux density, B. Thus, for an optimized, converged distribution, every photon, $n \in N$, should conform to the following equality:

$$\int_{\mathcal{A}_V(n)} B(a)da = \frac{1}{|N|}\int_\sigma B(a)da, \tag{3}$$

where σ is the combined space of the photon distribution on all illuminated surfaces in the scene. This equation stipulates that the integral of the flux density over each Voronoi cell, $V(n)$, should be equal to $\frac{1}{|N|}$ of the integral of flux density over σ, or the mean flux incident to all photons in N.

We initialize our algorithm by propagating the primal set, N, of photons which are scattered and stored using the standard techniques described by Jensen [2001]. The initial power of each photon is set to $\frac{1}{|N|}\Phi_\sigma$, where Φ_σ is the combined radiant power of all light sources in the scene. We describe these photons as *persistent* since they remain in memory and are referenced at render time to estimate L. In the pseudocode example, Algorithm 1, this step is encapsulated by the function CASTPERSISTENTPHOTONS().

Algorithm 1 PROGRESSIVERELAX()

```
N ← CastPersistentPhotons()
T = 1
δt = 1/|N|
repeat
    T_{i+1} = αT        {Eq. 13}
    while T < T_{i+1} do
        m ← CastTransientPhoton()
        if m is stored then
            T = T + δt
            n ← FindIntersectedCell(m)
            Σ_Φ(n) = Σ_Φ(n) + ΔΦ_p(m)
        end if
        DeletePhoton(m)
    end while
    for all n ∈ N do
        r_V(n) ← EstimateLocalArea(n)
        r_γ(n) = r_V(n) sqrt(γ̂(n)^{W(n)})        {Eqs. 8 and 15}
    end for
    for all n ∈ N do
        x⃗(n) = x⃗(n) + W(n) f⃗(n)        {Eq. 15 }
    end for
until User stop
```

Initial discrepancy in the distribution means that the flux capacity of any given persistent photon in N is likely to be in an unbalanced state (Figure 2a). Consequently, we need to adjust the area, $\mathcal{A}_V$, of each Voronoi cell until an equilibrium is reached. In order to determine by how much a cell needs to be enlarged or reduced to meet Equation 3, we calculate the disparity, γ, between the two sides of the expression:

$$\gamma(n) = \frac{1}{|N|}\int_\sigma B(a)da \left[\int_{\mathcal{A}_V(n)} B(a)da\right]^{-1}. \tag{4}$$

The coefficient computed by this function implies that if a persistent photon cell's flux capacity is p-times the mean cell flux capac-

ity, its area should be scaled by $\frac{1}{p}$ in order to balance the constraint equation. This is based on the assumption that the energy across the cell and outlying space is homogeneous. In practice, variation in flux density necessitates iterative refinement, thereby allowing the estimator to converge toward the correct solution.

4.1 Monte Carlo Estimation of Flux

The first problem we encounter when enforcing the constraint equation is how to accurately evaluate the integral of flux density over the area spanned by each individual Voronoi cell. This is problematic because B is unknown *a priori* to rendering and hence cannot be explicitly sampled. In order to seek a solution, we re-express the function of flux density as its derivative, $B = \frac{d\Phi}{da}$. Substituting this into the integral in Equation 3 yields:

$$\int_{A_V(n)} B(a)da = \int_{A_V(n)} \frac{d\Phi(a)}{da} da\,. \tag{5}$$

We next consider that a photon map stores radiant energy as discrete packets of flux with power $\Delta\Phi_p$. It follows that the function of incident illumination can be precisely encoded using an ancillary photon map, M, of infinite size and infinitesimal per-photon flux. Such a discretization allows the integral of B over a Voronoi cell to be expressed as the limit of the following summation:

$$\int_{A_V(n)} \frac{d\Phi(a)}{da} da \equiv \lim_{\frac{1}{|M|}\to 0} \sum_{m\in V(n)} \Delta\Phi_p(m)\,. \tag{6}$$

Here, m is an element of the set of all photons in M that intersect with the Voronoi cell, V.

Assuming photon trajectories are functions of independent and identically distributed random variables, we substitute Equation 6 with a Monte Carlo estimator of flux capacity. The variance of such an estimator diminishes asymptotically as more samples are added. This allows us to progressively refine the solution to an arbitrary precision in an open-ended way. We classify photons used to refine the estimate as being *transient*. This is because they are destroyed once their contribution has been accumulated resulting in a constant memory bound. Throughout the following sections, transient photons are referred to explicitly. All other references to photons are implicitly persistent or else refer to those used by standard or progressive photon mapping.

To ensure correctness, the primal flux, $\Delta\Phi_p$, of all transient photons in M must equal those in N. The estimate must be also normalized as more samples of $\Delta\Phi$ are included. This yields the following Monte Carlo summation:

$$\int_{A_V(n)} \frac{d\Phi(a)}{da} da \approx \frac{1}{T} \sum_{m\in V(n)} \Delta\Phi_p(m)\,. \tag{7}$$

Here, T is the sample density, defined as $1 + \frac{|M|}{|N|}$ where $|M|$ is the total number of transient photons cast. $\Delta\Phi_p(m)$ is the flux carried by m. Note that the flux encoded by each persistent photon is included in its estimate, hence $T = 1$ when $|M| = 0$.

Our algorithm progresses by casting one transient photon at a time (CastTransientPhoton() in Algorithm 1), locating the Voronoi cell, n, in the persistent set with which it intersects, then adding its value of $\Delta\Phi_p$ to the accumulated estimate at n. The transient photon is then destroyed. Intersected cells can be found rapidly in $O(\log|N|)$ time using a balanced k-d tree. Christensen [1999] employs this approach when querying precomputed irradiance estimates.

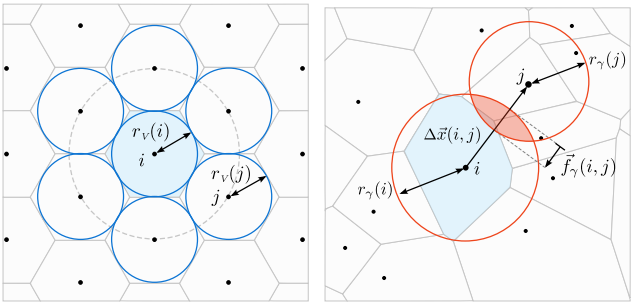

(a) (b)

Fig. 6. (a) A lattice of photons and associated bounding spheres arranged in an ideal blue noise configuration. The radius, r_V, is directly estimated from the underlying distribution. (b) r_V is adjusted by $\hat{\gamma}$, yielding r_γ. This value is used by the relaxation operator to repel or attract neighboring photons. The magnitude of the resultant force, f_γ, is calculated to move i so that its bounding sphere grazes that of j.

Using the discretization in Equation 7, we can re-express Equation 4 in terms of an estimate of γ:

$$\hat{\gamma}(\mathrm{n}) = \frac{1}{|N|} \sum_{m\in M} \Delta\Phi_p(m) \left[\sum_{m\in V(n)} \Delta\Phi_p(m) \right]^{-1}. \tag{8}$$

For brevity, Equation 8 simplifies the estimation of incident flux by reducing it to a scalar quantity. To allow for variations in chromaticity, we also store the accumulated estimation of the energy of the photon in RGB color space.

4.2 Cell Area Approximation

The value of $\hat{\gamma}$ computed by Equation 8 specifies the change in area of each persistent photon's Voronoi cell required to satisfy the flux capacity constraint. Unfortunately, accurately adjusting cell area is a non-trivial operation for several reasons. As stated in Section 3.2, each Voronoi cell is not an explicitly defined primitive and hence cannot be arbitrarily scaled. In addition, variation in the position of neighboring photons further influences its resultant size and shape making it hard to accurately constrain.

Our solution is to apply iterative point relaxation through controlled diffusion, encouraging neighboring persistent photons to adjust their proximity by attraction and repulsion. To achieve this, we approximate the Voronoi cell, $V(n)$, using a sphere centered around each photon's position. The area of each sphere is estimated from the local distribution and can be used to guide the movement of adjacent neighbors. The resultant migration of points then indirectly affects their corresponding cell areas. Since a sphere is scaled by a single parameter, r, the space to be explored is a monotonic, univariate function. Hence, increasing a photon's radius increases the integral of B over the cell area and vice versa.

As stated previously, our goal is to move persistent photons so they become locally equidistributed with a characteristic blue noise spectral signature. Thus, we consider a hexagonal lattice of perfectly arranged points which represent an ideal configuration, then determine what criteria are met for said points to be considered relaxed.

This concept is illustrated in Figure 6a where a group of photons are shown with their Voronoi cells and representative bounding spheres. The radius of each sphere has been chosen to most closely match the area of its associated cell. We compute this radius, r_V,

as being half of the sum of the mean distance and standard deviation of the distance to its 6-nearest neighbors. In this configuration, sphere boundaries precisely graze those adjacent to them. Thus, we propose that an optimally relaxed distribution will minimize the distance between the perimeter of its spherical boundary and those of its neighbors. In the event of a neighboring sphere boundary being either too close or too far, the photon must subsequently be repelled from or attracted to that neighbor respectively.

4.3 Relaxation Operator

We calculate a force, $\vec{f}$, that is used to modify the position of each persistent photon based on the compound influence of its spatially adjacent neighbors. The force between a photon, i, and its neighbor, j, is a function of the finite difference in position, $\Delta\vec{x}$. This is defined as $\Delta\vec{x}(i,j) = \vec{x}(j) - \vec{x}(i)$ (illustrated in Figure 6b). The goal of iterative relaxation is to move each point until it comes to rest in a state of equilibrium.

The magnitude of the force acting along each $\Delta\vec{x}(i,j)$ is the sum of two components. The first applies a positive or negative pressure proportional to the estimated flux disparity, $\hat{\gamma}$. We use Equation 8 to modify the radius of each photon's cell area approximation, r_v, yielding a new term, r_γ:

$$r_\gamma(n) = r_v(n)\sqrt{\hat{\gamma}(n)}\,. \tag{9}$$

Photons are encouraged to move apart or cluster together depending on the specified change in cell area. We illustrate this force in Figure 6b where the change in measured radius due to $\hat{\gamma}$ causes photon radii to overlap. The force between i and j is the motion required to push them apart (or pull them together) so that their bounding spheres come as close as possible to grazing. We define this scalar component, f_γ, as:

$$f_\gamma(i,j) = \|\Delta\vec{x}(i,j)\| - (r_\gamma(j) + r_\gamma(i))\,. \tag{10}$$

To assist in equidistribution we also apply a diffusive force that acts to move photons into the ideal configuration in Figure 6a. Here, i is away from or toward j by the difference in r_v with the aim of equalizing the areas of adjacent cells. We define this scalar component, f_v, as:

$$f_v(i,j) = r_v(j) - r_v(i)\,. \tag{11}$$

Equation 11 is moderated by Equation 10 so diffusion is inhibited when it violates the flux capacity constraint.

Combining these two components, we express the resultant force on photon i as the mean force exerted by all spatially adjacent photons, K_f.

$$\vec{f}(i) = \frac{1}{|K_f|}\sum_{j\in K_f}\frac{\Delta\vec{x}(i,j)}{\|\Delta\vec{x}(i,j)\|}\left[\, f_\gamma(i,j) + f_v(i,j)\,\right]. \tag{12}$$

We determine which photons are spatially adjacent and therefore belong in K_f if *either* they are a member of the 6-nearest neighbors, *or* they satisfy the condition $\|\Delta\vec{x}(i,j)\| \le 2r_\gamma(i)$.

Finally, to allow for correct redistribution on curved surfaces, relaxation must be performed in the tangent space of i. Hence, $\vec{f}$ is projected into the plane derived from the photon's surface normal before being added to $\vec{x}$.

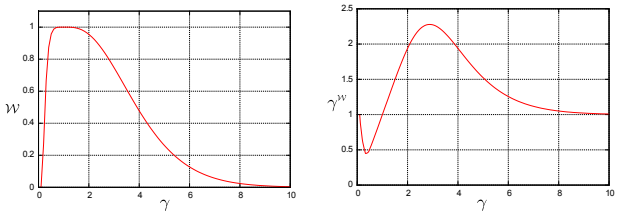

(a) (b)

Fig. 7. (a) A plot of γ against $\mathcal{W}$. Note the flattened peak of the curve resulting in minimal attenuation when the value of γ is close to 1. (b) A plot of γ against $\gamma^{\mathcal{W}}$. If the value of γ deviates significantly from unity, it is damped back to 1.

4.4 Progressive Refinement

As transient photons are cast and absorbed throughout the scene, the estimate of flux at each persistent photon and, in turn, the estimates of γ begin to converge. After casting a predetermined number of transient photons, we calculate each photon's flux-adjusted bounding sphere, r_γ (Equation 9), before relaxing the distribution using Equation 12.

As with all Monte Carlo integrals, we observe that the error, ϵ, in the estimate of flux over each cell diminishes in proportion to the inverse square root of the mean number of transient photons per cell. Since $\hat{\gamma}$ is a function of this estimate, the change in cell area is a function of ϵ. It is therefore computationally efficient to choose the interval between relaxation iterations based upon a similar relation so as to reflect the continually decreasing gradient of the error. For a given iteration, i, at time T_i, the expected value of the error obeys the following correspondence: $E[\epsilon] \propto \frac{1}{\sqrt{T_i}}$. Hence, the next iteration at time T_{i+1} should occur once the mean residual error has attenuated by a constant factor since T_i. This can be expressed as a recurrence relation:

$$\begin{aligned} T_0 &= 1 \\ T_{i+1} &= \alpha T_i\,, \end{aligned} \tag{13}$$

where α is the scaling parameter. We use a value of $\alpha = 500^{\frac{1}{60}} \approx 1.1091$. This yields approximately 60 iterations by $T = 500$ and allows the distribution to adequately relax without incurring a prohibitive computational overhead.

4.5 Detecting and Constraining Edge Cases

One drawback of using a sparse point cloud to encode B is that the boundary of the set is poorly defined. In particular, the transition between zero and non-zero irradiance results in unbounded Voronoi cells with infinite area. Since no real centroid of such a cell exists, the relaxation operator experiences unchecked repulsion. We illustrate this problem in Figure 8 where a group of persistent photons lies on the edge of a flux discontinuity. The cells shaded in red in Figure 8b are infinitely large and this translates into poorly estimated repulsion radii, as depicted in Figure 8c.

To avoid an unchecked increase in r_v we capitalize on the property that escaped photons' $\hat{\gamma}$ values increase monotonically once they migrate into an area with a significantly low flux density. We therefore specify a weighting function, $\mathcal{W}$, that tends to zero as $\hat{\gamma}$ grows large:

$$\mathcal{W}(n) = \exp(-\tfrac{1}{5}\ln(\gamma(n))^4)\,. \tag{14}$$

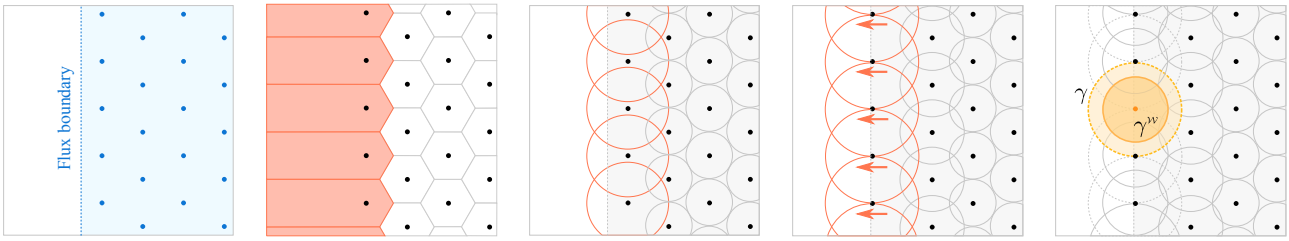

Fig. 8. The result of unchecked persistent photon migration. (a) Photons arrive with equal probability to the right of the flux boundary and zero probability to the left. (b) The Voronoi cells of the photons shaded in red have infinite area because of their position near the boundary. (c) This translates to an overestimation of r_V from the k-nearest neighbors. (d) Relaxation repels the edge photons from their neighbors across the flux boundary. The error is compounded at the next iteration with more photons becoming affected. (e) The repulsion radius computed from γ (dotted outer circle) is attenuated using the weighting function to arrive at the reduced radius $\gamma^{\mathcal{W}}$ (solid inner circle), thereby limiting unchecked migration.

A plot of this function is illustrated in Figure 7a. We raise the exponent to the 4^{th} power in order to lower the second derivative $\frac{d^2\mathcal{W}}{d\gamma^2}$ around $\gamma = 1$, flattening out the curve at its apex and preventing unnecessary constraint to photons merely affected by noise. The computed value of $\hat{\gamma}$ is constrained by raising it to the power of $\mathcal{W}$, preventing escaped photons' repulsion radii from growing exponentially large. This function is plotted in Figure 7b and its effect illustrated diagrammatically in Figure 8e. We also attenuate the computed relaxation force $\vec{f}(n)$ (Equation 12) by its weight $\mathcal{W}$. This yields the following modified values for γ and $\vec{f}$:

$$\begin{aligned} \gamma'(n) &= \gamma(n)^{\mathcal{W}(n)}\,, \\ \vec{f}'(n) &= \mathcal{W}(n)\vec{f}(n)\,. \end{aligned} \tag{15}$$

4.6 Modifying the Radiance Estimator

The relaxation operator used by our approach approximates Voronoi cell area, r_V, using a density estimation kernel and is thereby prone to a form of proximity bias. This error translates to a mis-estimation of cell area and prevents Voronoi cells from precisely adhering to the flux capacity constraint in Equation 3. The effect is highlighted in Figure 5 where a sample PDF is used to stochastically seed a persistent photon dataset. As the distribution is relaxed, the color encoding reveals that values of $\hat{\gamma}$ in regions where the illumination gradient is non-zero do not converge to 1. This is due to each Voronoi cell's area being affected by the residual diffusion of the distribution in its local neighborhood.

By using the accumulated approximation of flux across each photon's Voronoi cell, the effects of proximity bias are effectively counterbalanced. The original radiance estimate as defined by Jensen must be modified accordingly:

$$L(\vec{x} \to \vec{\omega}) \approx \frac{1}{T\pi d^2} \sum_{i=1}^{K_L} \left[f_r(x, \vec{\omega} \leftrightarrow \vec{\omega}_p(i)) \sum_{m \in V(i)} \Delta\Phi_p(m) \right], \tag{16}$$

where all notation corresponds to Equations 1 and 7.

5. RESULTS

All renders were carried out on an Intel Core i7 with 8GB of RAM running Windows 7. We found that nearly all stages of our algorithm were well-suited to vectorization and so used a total of 7 threads over 4 cores with 1 additional manager thread. Ray casting, transient photon merging and persistent photon relaxation were all successfully load balanced and we found that increasing the number of threads from 1 upward scaled well and left comparatively few processing bottlenecks.

5.1 Comparative Algorithms

Progressive photon relaxation (PPR) is a hybrid approach that uses a persistent point dataset refined through progressive integration. As such, it exhibits a memory bound proportional to Jensen's original photon mapping method and a runtime in the order of unbiased Monte Carlo methods. In order to draw objective comparisons, we compare our approach against both classic photon mapping (PM) [Jensen 2001] and progressive photon mapping (PPM) [Hachisuka et al. 2008].

Classic photon mapping forms the framework upon which our approach is based and is homologous to an unrefined and unrelaxed, persistent photon map. Hence, it represents an equal-memory test case. In all examples, images are rendered using 20-nearest neighboring photons to ensure comparable levels of kernel bias. An isotropic Epanechnikov filter [Silverman 1986] is used to minimize kernel bias. Where appropriate, we also test against photon relaxation [Spencer and Jones 2009]. Like photon mapping, this method uses approximately the same amount of memory but requires a short preprocessing step proportional to the number of persistent photons.

Progressive photon mapping has been implemented according to the default parameters supplied by Hachisuka et al. [2008]. The initial radius of each radiance estimate is based on the pixel footprint derived from the distance to its neighboring hit points. The runtimes of both PPM and PPR are of the order $O(n)$ for n transient photons, so we draw an equal-time comparison by referring to the total number of photons cast. We chose progressive photon mapping over unbiased Monte Carlo methods due to its robustness at handling reflections and transmissions of caustics. Our implementation does not use the stochastic sampling algorithm described in Hachisuka and Jensen [2009], so images rendered using PPM are not anti-aliased.

In Table I we supply timings for all test scenes. These comprise photon propagation (applicable to both PPR and PPM) and image render time (applicable only to our approach). We also list the number of persistent photons, $|N|$, and the sample density, T. Note here that the render time of Figure 4 is significantly lower because no direct lighting computations were necessary in this scene.

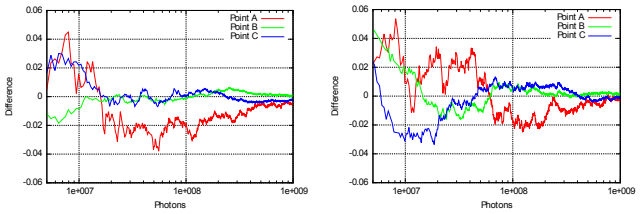

(a) Our method (b) Progressive photon mapping

Fig. 9. Measured difference in pixel intensity over time for the Torus Cube scene in Figure 12.

5.2 Convergence Tests

To visualize the convergence properties of our approach we plot the change in pixel intensity over a selection of 3 sample points as progressively more transient photons are cast. For this test, we use the Torus scene from Figure 12. As a reference, we use a converged image rendered using progressive photon mapping. The results are illustrated in Figure 9.

We observe that, unlike progressive photon mapping, the asymptote around which each pixel converges in Figure 9a may not be that of the reference. This is because the measured error of $\hat{L}$ is the sum of both the sampling noise in the Monte Carlo estimates of $\Delta\Phi_p$, and the bias from the kernel itself. Since the number of photons used in the reconstruction of L remains constant, estimator bias does not diminish asymptotically to zero. However, as demonstrated by the rendered examples in Section 5.3, visual noise is still successfully removed while the residual effects of bias are minimized through the use of low-bandwidth kernels.

5.3 Test Scenes

In Figure 1, a bright light is shone through a glass vase containing recently disturbed water. This scene demonstrates the ability of our algorithm to optimize high-frequency caustics without diffusing away detail. Our progressive approach preserves intricate details caused by the turbulent nature of the water's surface.

Figure 4 is designed to demonstrate caustics generated by materials with different scattering properties. A glass prism is split into two pieces, one rendered with a glossy isotropic Ward material, the other a perfect specular BRDF. Approximately 220,000 persistent photons are stored and refined with an additional 88,000,000 transient photons. Note how the fine striations caused by imperfections in the glass surface are visible after progressive refinement. In the unrelaxed comparison (Figure 4a), noise results in these details becoming lost.

In Figure 10, the WinOSI global illumination test scene [Granz 2012] with modifications inspired by Hachisuka et al. [2008] offers a further example of complex caustic illumination. Here, two small area emitters have been placed inside specular reflective tubes forcing photons to bounce multiple times before exiting through the aperture in the base. Both a faceted and smooth, specular reflective sphere, and a dielectric ball further reflect and refract the photons. 700,000 persistent photons were relaxed using approximately 230,000,000 transient photons in the progressive image. In this scene, the relaxed caustic photon map is also used to compute diffuse interrefection, both for our method and PPM, via a final gather pass. This scene also reveals that residual bias cannot be altogether eliminated by our approach, despite using a low-bandwidth radiance estimate. Examples of this error can be observed at the base of the walls and around other illumination discontinuities.

Figure 11 depicts a triangular glass prism mounted below a curved, reflective strip. Coherent white light is dispersed by the prism before being repeatedly reflected and finally absorbed. 500,000 persistent photons are stored and 20-nearest neighbors used in the radiance estimate. Our algorithm was run to 1,000 passes making a total of 500,000,000 accumulated photons. We also compare the performance of our method against three other algorithms. *Standard photon mapping* (Figure 11a) is tested using a map containing 5,000,000 photons with 200-nearest neighbors in the radiance estimate. This example emphasizes the argument made in Section 3.2 that increased physical storage of photons (in this case, one order of magnitude more) is not a practical means of removing noise. *Photon relaxation* (Figure 11b) reduces noise but severely degrades fine detail due to the failure of the feature detection algorithm to capture and preserve the fine caustic filaments. We also test against a map that has been allowed to converge without photon redistribution (Section 3.2). The mottling effect demonstrated in Figure 4c is also clearly visible.

Figure 12 is inspired by the caustic test scene devised by Cline et al. [2005]. A diffuse torus is suspended in a transparent dielectric cube which, in turn, is suspended in mid-air. The persistent photon map contains approximately 500,000 photons with our progressive relaxation method integrating a further 500,000,000 transient photons to remove noise. We also compare our approach with progressive photon mapping using the same number of transient photons for both tests. An analysis of the convergence rates at selected sample points in this scene is given in Figure 9.

	Fig. 1	Fig. 4	Fig. 11	Fig. 10	Fig. 12
Resolution	548 x 470	640 x 640	640 x 750	640 x 480	800 x 800
$\|N\|$	500k	220k	500k	700k	500k
T	20	500	1000	350	1000
PPR/PPM t	6m	25m	27m	45m	1h 35m
PPR render t	1m 58s	1m 21s	0m 8s	1m 23s	1m 18s

Table I. Render statistics for included figures. *PPR/PPM t* indicates the time taken to cast transient photons and construct both the progressive photon map and relaxed photon map scene in each example. *PPR render t* is the approximate time required to re-render the scene using progressive photon relaxation from an arbitrary viewpoint. For progressive photon mapping, this time would be approximately equal to *PPR/PPM t* at every frame.

6. DISCUSSION

The original photon mapping method is known to be susceptible both to bias from the density estimation kernel and to noise due to discrepancy in the underlying distribution. These artifacts are particularly detrimental when caustics form the dominating component of global illumination. Noise removal either incurs a geometric increase in storage overhead, or results in loss of detail and blurring due to kernel bias. Our technique avoids these penalties by optimizing the underlying particle distribution. This strategy improves rendering quality while avoiding an increase in the stored number of photons. As shown by the results in Section 5, it is possible to obtain images of similar fidelity to unbiased or asymptotically consistent methods while retaining the flexibility to reconstruct an image from an arbitrary viewpoint.

While evaluating the applications of our progressive method, we found it useful to consider potential application areas. Architectural scenes and product designs would see a significant benefit since in these scenarios the photon map is often static. With thousands

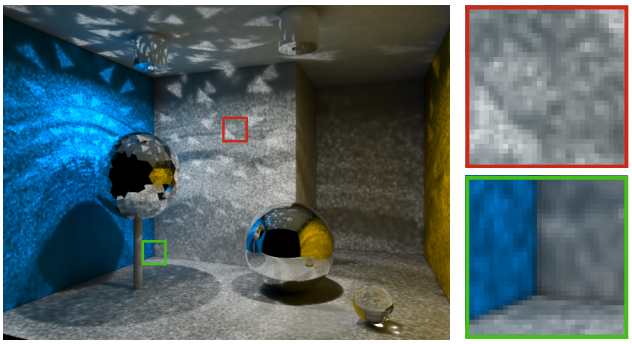

Standard photon mapping [Jensen 2001].
700k photons.

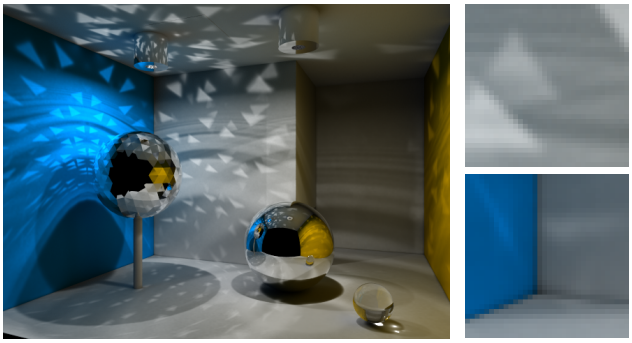

Progressive photon relaxation.
700k persistent photons. 245M transient photons.

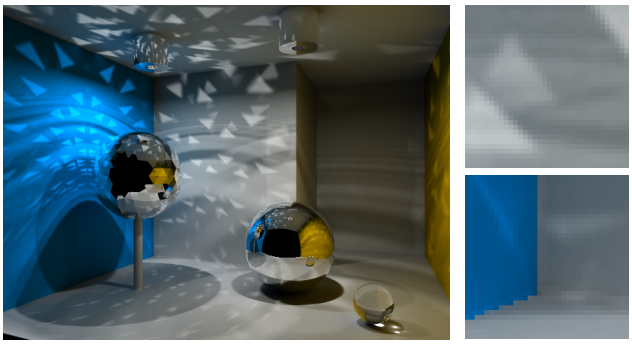

Progressive photon mapping [Hachisuka et al. 2008].
245M transient photons.

Fig. 10. The WinOSI global illumination test scene [Granz 2012] with modifications similar to those made by Hachisuka et al. [2008].

of rendered frames, the one-off preprocessing cost of progressive photon relaxation is low compared to the potential overhead of re-generating a new photon map at every frame. An additional benefit of our method is its construction on top of the basic photon mapping framework established by Jensen [2001]. If the infrastructure for photon tracing already exists (either using the CPU or GPU), our algorithm can bolt on to existing software.

7. FUTURE WORK

We have identified several promising avenues of further research into improving both the efficiency and versatility of progressive photon relaxation.

For materials with non-Lambertian BRDFs, variance in the sparse encoding of incident illumination by the photon map introduces noise which is not completely eliminated by our technique. Increasing the bandwidth of $\hat{L}$ to compensate is problematic for the reasons discussed in Section 3.2. A possible solution would be to construct an estimator that takes advantage of the de-noised distribution to differentiate between bias in radiance and irradiance. This would make PPR more flexible over a broader range of materials.

Further analysis could also be made into how the change in Voronoi cell area affects the estimate, $\hat{\gamma}$. As demonstrated empirically in Figure 9a, variance introduced by cell fluctuation diminishes as the distribution converges. However, quantification of this error would allow for its compensation and could improve the speed of convergence.

Another area of interest is the initial density of the persistent photon map. Currently, this is based upon the same criteria as conventional photon maps. A heuristic that selects photon density automatically, perhaps based upon a statistical analysis of the scene, would be advantageous. Comprehensive analysis of the convergence rate of our approach would also aid in this goal.

Investigation into Schjøth et al's [2007] adaptation of the ray differentials framework could improve the performance of progressive photon relaxation when rendering complex specular caustics. In particular, instances where ray footprints become highly anisotropic would likely benefit from this technique.

8. CONCLUSION

In this paper we have presented a novel approach to photon mapping using point relaxation and progressive refinement to optimize the persistent photon distribution. We have demonstrated the advantages of preserving flux homogeneity by generating noise-free renders of intricate caustics using very low-bandwidth kernels. We have also shown that complex illumination effects such as spectral dispersion can be effectively integrated. We highlight the key benefits of progressive photon relaxation as view-independence, robust noise removal, low memory overhead, suitability to vectorization, compatibility with existing photon mapping implementations and broad scope for expansion.

ACKNOWLEDGMENTS

The work presented in this paper was funded by the Wales Research Institute of Visual Computing and EPSRC grant number EP/I031243/1. The Prism image was inspired by the work of Zipper at the deviantART community (zipper.deviantart.com). The model in Figure 12 was kindly made available by Toshiya Hachisuka (cs.au.dk/~toshiya). Finally, we would like to thank the anonymous reviewers for their valuable feedback and recommendations.

REFERENCES

BALZER, M., SCHLÖMER, T., AND DEUSSEN, O. 2009. Capacity-constrained point distributions: a variant of Lloyd's method. *ACM Trans. Graph. (Proceedings of SIGGRAPH) 28,* 3 (July), 86:1–86:8.

CHRISTENSEN, P. H. 1999. Faster photon map global illumination. *J. Graph. Tools 4,* 3, 1–10.

CLINE, D., TALBOT, J., AND EGBERT, P. 2005. Energy redistribution path tracing. *ACM Trans. Graph. (Proceedings of SIGGRAPH) 24,* 3 (July), 1186–1195.

DOIDGE, I., JONES, M. W., AND MORA, B. 2012. Mixing monte carlo and progressive rendering for improved global illumination. *The Visual Computer 28,* 6–8 (June), 603–612.

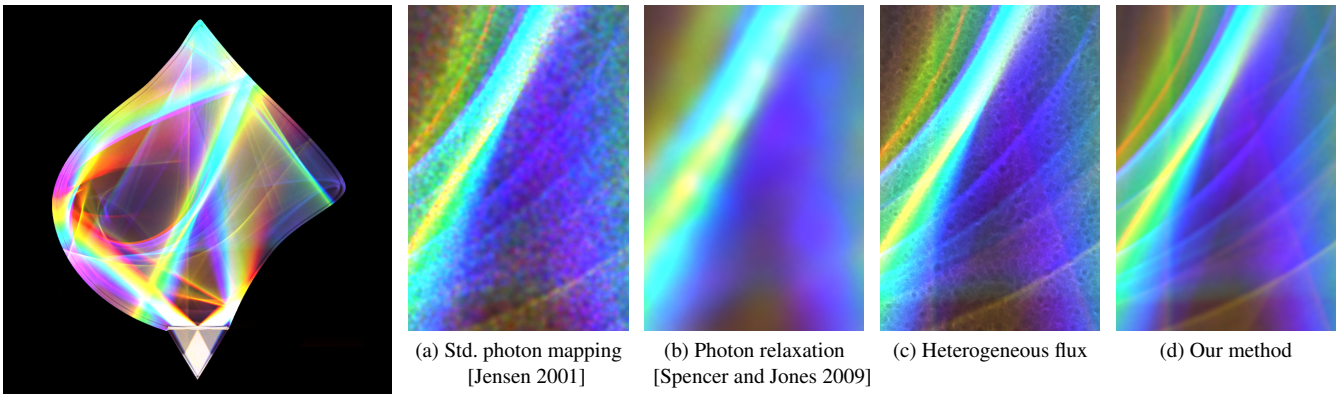

(a) Std. photon mapping [Jensen 2001] (b) Photon relaxation [Spencer and Jones 2009] (c) Heterogeneous flux (d) Our method

Fig. 11. Coherent white light refracted through a glass prism into a curved metallic strip. In each magnified frame we compare our technique against photon relaxation, standard photon mapping and the heterogeneous approach discussed in Section 3. Note that the map used by the standard method contains an order of magnitude more persistent photons than the remaining three.

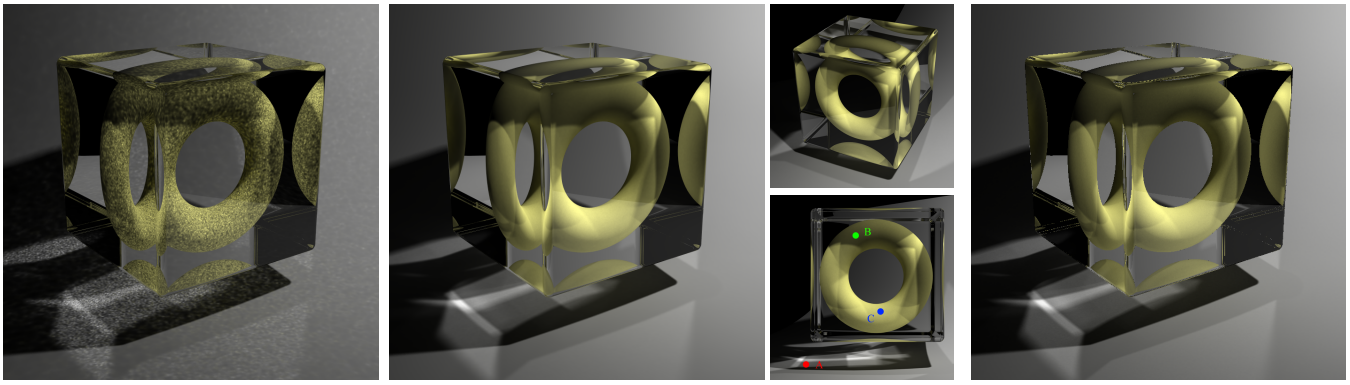

Standard photon mapping
500k persistent photons.

Progressive photon relaxation
500k persistent photons, 500M transient photons.

Progressive photon mapping
500M transient photons.

Fig. 12. Torus in a glass cube based on the caustic illumination test scene devised by Cline [2005]. Note that progressive photon relaxation generates a view-independent photon map which can be rapidly re-rendered from multiple angles. The sample points for the convergence tests in Figure 9 are also shown.

GRANZ, M. 2012. WinOSi homepage. `http://www.winosi.onlinehome.de`. Retrieved March 2012.

HACHISUKA, T. AND JENSEN, H. W. 2009. Stochastic progressive photon mapping. In *SIGGRAPH Asia '09: ACM SIGGRAPH Asia 2009 papers*. ACM, New York, NY, USA, Article No. 141.

HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive photon mapping. In *SIGGRAPH Asia '08: ACM SIGGRAPH Asia 2008 papers*. ACM, New York, NY, USA, Article No. 130.

IGEHY, H. 1999. Tracing ray differentials. In *SIGGRAPH '99: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 179–186.

JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. *ACM Trans. Graph. (Proceedings of SIGGRAPH Asia) 30,* 6, 181:1–181:12.

JENSEN, H. W. 1996. The Photon Map in Global Illumination. Ph.D. thesis, Technical University of Denmark, Lyngby.

JENSEN, H. W. 2001. *Realistic Image Synthesis using Photon Mapping*. A. K. Peters, Ltd.

JENSEN, H. W. AND CHRISTENSEN, N. J. 1995. Photon maps in bidirectional Monte Carlo ray tracing of complex objects. *Computers and Graphics 19,* 2, 215–224.

KAJIYA, J. T. 1986. The rendering equation. In *SIGGRAPH '86: Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press, New York, NY, USA, 143–150.

KELLER, A. 1995. A quasi-Monte Carlo algorithm for the global illumination in the radiosity setting. In *Monte-Carlo and Quasi-Monte Carlo Methods in Scientific Computing*. Springer, 239–251.

KNAUS, C. AND ZWICKER, M. 2011. Progressive photon mapping: A probabilistic approach. *ACM Trans. Graph. 30*, 25:1–25:13.

LAFORTUNE, E. P. AND WILLEMS, Y. D. 1993. Bi-directional path tracing. *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics 93)*, 145–153.

LAI, Y.-C., FAN, S., CHENNEY, S., AND DYER, C. 2007. Photorealistic image rendering with population Monte Carlo energy redistribution. In *Proc. Eurographics Symposium on Rendering*. 287–296.

LLOYD, S. A. 1983. An optimization approach to relaxation labeling algorithms. *Image and Vision Computing 1,* 2 (May), 85–91.

SCHJØTH, L., FRISVAD, J. R., ERLEBEN, K., AND SPORRING, J. 2007. Photon differentials. In *GRAPHITE '07*. ACM Press, New York, NY, USA, 179–186.

SCHJØTH, L., OLSEN, O. F., AND SPORRING, J. 2005. Diffusion based photon mapping. Ph.D. thesis, IT University Copenhagen.

SCHREGLE, R. 2003. Bias compensation for photon maps. *Computer Graphics Forum 22,* 4, 729–742.

SECORD, A. 2002. Weighted Voronoi stippling. In *NPAR '02: Proceedings of the 2nd International Symposium on Non-photorealistic Animation and Rendering*. ACM, 37–43.

SILVERMAN, B. W. 1986. *Density Estimation for Statistics and Data Analysis*. Springer.

SPENCER, B. AND JONES, M. W. 2009. Into the blue: Better caustics through photon relaxation. *Eurographics 2009, Computer Graphics Forum 28,* 2 (March), 319–328.

SUYKENS, F. AND WILLEMS, Y. D. 2000. Density control for photon maps. In *Rendering Techniques: Proceedings of the Eurographics Workshop on Rendering*. Springer-Verlag, London, UK, 23–34.

ULICHNEY, R. A. 1988. Dithering with blue noise. *Proceeding of the IEEE 76,* 1, 56–79.

VEACH, E. AND GUIBAS, L. J. 1994. Bidirectional estimators for light transport. *Proceedings of the 1994 Eurographics Rendering Workshop*, 147–162.

VEACH, E. AND GUIBAS, L. J. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *SIGGRAPH '95: Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. ACM, New York, NY, USA, 419–428.

VEACH, E. AND GUIBAS, L. J. 1997. Metropolis light transport. In *SIGGRAPH '97: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 65–76.

Received September 2008; accepted March 2009