Boundary-Aware Extinction Mapping

Pascal Gautron[†], Cyril Delalandre, Jean-Eudes Marvie, Pascal Lecocq
Technicolor

Figure 1: We introduce a novel formulation for fast scattering simulation in arbitrary participating media. This animated scene features a 1024^3 volcano smoke and 75 other media such as smoke and clouds (512^3 each), rendered interactively at 1.5 fps.

Abstract

We introduce *Boundary-Aware Extinction Maps* for interactive rendering of massive heterogeneous volumetric datasets. Our approach is based on the projection of the extinction along light rays into a boundary-aware function space, focusing on the most relevant sections of the light paths. This technique also provides an alternative representation of the set of participating media, allowing scattering simulation methods to be applied on arbitrary volume representations. Combined with a simple out-of-core rendering framework, *Boundary-Aware Extinction Maps* are valuable tools for interactive applications as well as production previsualization and rendering.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three Dimensional Graphics and Realism—Color, shading, shadowing, and texture; I.3.6 [Computer Graphics]: Methodology and Techniques—Graphics data structures; G.1.2 [Mathematics of Computing]: Numerical Analysis—Approximation

1. Introduction

The complex interactions of participating media with light contribute to visually compelling effects. However, their representation requires vast amounts of memory and expensive rendering algorithms. Further performance can be achieved using Deep Shadow Maps [LV00], which extend Shadow Maps [Wil78] by recording light attenuation samples along rays. Once combined with compression and filtering algorithms, Deep Shadow Maps are widely used in production rendering. However, the costs of such algorithms forbid a generic implementation on graphics hardware. We address this issue by introducing *Boundary-Aware Extinction Maps*, in which we accumulate arbitrary sample sets into a boundary-aware function space. The resulting extinction functions are then analytically integrated into transmittance functions. These maps can also be ray-marched [PH89] for

fast estimation of image-based lighting and multiple scattering across multiple media.

Our method particularly applies to real-time rendering of massive volumetric environments for fast previsualization of production assets. Scaling is achieved by integrating Boundary-Aware Extinction Maps into an out-of-core framework for adaptive loading of both static and animated production-quality participating media (Figure 1).

2. Related Work

Research towards efficient rendering of participating media produced numerous contributions [KVH84, CFP*05, EHK*06, WZH*11, JSYR12] aiming at solving the radiative transfer equation [Cha50]. Complex light behaviors can be simulated using variations of photon mapping [JC98, JZJ08, BPPP05, JNSJ11] and instant radiosity [Kel97, NNDJ12b, NNDJ12a, DKH*13]. However, such methods usually belong to the world of offline rendering [NJS*11]. Real-time

[†] Now affiliated to NVIDIA ARC

Figure 2: Overall shape of the transmittance function T through participating media. While the representation of T over the entire path is mandatory when rendering a medium filling the entire range (a), the function features constant sections between multiple media (b). Using an appropriate pseudometric g , T can be represented without gaps (c).

approaches tend to involve limitations such as homogeneous media [Mit05, WR08], specific representation [BNM*08, DGMF11] or heavy precomputations [ZRL*08].

Deep Shadow Maps (DSM) [LV00] represent transmittance functions by storing a linked list of all the volume samples along light rays. The sample lists can be efficiently quantized and filtered for higher performance. This approach inherently requires the storage of an unpredictable amount of samples per pixel depending on the depth complexity of the scene. Graphics hardware implementations of DSM [KN01, KPH*03, JHH*09, YHGT10] then typically introduce restrictions such as a fixed list size, potentially yielding visually disturbing artifacts. Massive volumetric datasets also introduce challenges in terms of real-time storage, sorting, quantization and filtering of samples.

Adaptive Volumetric Shadow Maps (AVSM) [SVLL10] simplify the DSM generation using a fixed number of transmittance values per texel. The contribution of each transmittance sample to the overall transmittance curve is estimated and compared to the contributions of the already stored transmittance values. This algorithm then eliminates weakly contributing samples on the fly. This technique proves effective in many cases, at the cost of searches and insertions for each transmittance sample. AVSM are then particularly suitable in scenes containing a small number of particles. Deep Opacity Maps [YK08] focus on hair rendering: the light entry points are first obtained by rendering the outer shell of the hair. Several opacity layers are then collected by moving the entry points by user-defined offsets. While effective for hair, other participating media require a fine tuning of the number and spacing of the layers. Furthermore, large volumetric datasets require a prohibitive number of layers.

Fourier Opacity Maps (FOM) [JB10] and Transmittance Function Maps (TFM) [DGMF11] avoid an explicit storage of volume samples by projecting the opacity of volume elements into Fourier series. Those techniques are typically suitable for media generating low frequency lighting variations such as clouds and smoke. However, they either impose voxel-based media [DGMF11] or estimate scattering very coarsely [JB10]. Also, the projection into Fourier series is most efficient with few moderately dense media and known boundaries of the integration domain. Ringing artifacts then tend to appear with sparse or multiple participating media. While using a FOM or TFM per medium could be achieved in simple cases, the linear scaling in memory consumption precludes the extension to massive scenes. Voxelized media are required by TFM as this technique relies on a sequential ray marching to evaluate transmittance along a ray. We avoid these constraints by evaluating transmittance using analytical integration. Also, ringing artifacts in complex scenes are drastically reduced using a nonlinear metric to avoid gaps between media.

3. Problem Statement

Rendering a participating medium intensively evaluates the transmittance function at points x along rays:

$$T(x) = \exp\left(-\int_0^x \sigma_t(y) dy\right) \quad (1)$$

where $\sigma_t(y)$ is the medium's extinction at a point y . The integral can be evaluated numerically by a discretization into ray marching steps:

$$T(x_j) = T(x_{j-1}) \exp(-\sigma_t(x_j) d_j) \quad (2)$$

where d_j is the length of step j . The TFM [DGMF11] are generated using this technique to evaluate $T(x)$ along light rays. The transmittance samples are projected into an orthonormal basis $\{b_i(x)\}_{i \in \mathbb{N}}$, yielding coefficients T_i :

$$T(x) = \sum_{i=0}^{\infty} T_i b_i(x), \quad T_i = \int_0^D T(x) b_i(x) dx \quad (3)$$

where $D > 0$ is the length of the ray. In the case of multiple media separated by gaps all the regions are then treated equally, hence degrading the quality of the regions where media are effectively present.

We introduce a *Boundary-Aware Pseudometric* (Section 4) to reduce such gaps and increase the quality of the reconstructed signal. Then, we devise *Boundary-Aware Extinction Maps* (BAEM) for efficient transmittance evaluation in arbitrary participating media (Section 5). Our solution also provides a unified solution for fast ray marching within complex environments (Section 6). We apply this capability to interactive image-based lighting and multiple scattering. Finally, we combine BAEM with out-of-core management for efficient rendering of massive scenes comprising many detailed participating media (Section 7).

4. Boundary-Aware Pseudometric

As shown in Figure 2 the transmittance function is constant between participating media. Those constant sections are embedded as well in FOM and TFM [JB10, DGMF11] due to the use of non-localized function bases, resulting in increased ringing artifacts when the media are separated by large distances (Figures 7). The piecewise-linear transmittance of DSM [LV00] avoids this problem by explicitly storing the most characteristic points of the transmittance curve, at the cost of additional memory and per-textel management of linked lists.

The principles of DSM and TFM could be merged by explicitly storing the intervals where $T(x)$ is constant, and projecting the transmittance function only within the participating media (Figure 2). To this end, we can replace the Euclidean distance by a pseudometric [SS70] accounting for the ray sections traversing the participating media. A uniform sampling of the pseudometric would result in a similarly uniform sampling within the media and large steps across empty zones (Figure 2c). Such *Boundary-Aware Pseudometric* g between two points x_1 and x_2 is:

$$g(x_1, x_2) = \int_{x_1}^{x_2} t(x) dx, \quad t(x) = \begin{cases} 1 & \text{if } \sigma_t(x) > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

As the number of traversed media depends on the environment, we propose an approximate formulation for g using a fixed number of coefficients.

The binary function t can be evaluated at any point x along a ray. The resulting values can then be projected into an arbitrary orthonormal basis $\{c_i(x)\}_{i \in \mathbb{N}}$, yielding coefficients t_i :

$$t_i = \int_{x_{\min}}^{x_{\max}} t(x) c_i(x) dx = \sum_j C_i(x_{j+}) - C_i(x_{j-}) \quad (5)$$

where x_{j+} and x_{j-} are the boundaries of the j^{th} sample and C_i is indefinite integral of c_i . By convention $t(x) = 1$ within the volume sample, and hence vanishes in the derivation.

A noticeable effect of our series representation is the independence of the samples, whose contributions can be combined in any order. This typically makes our technique usable with renderers generating arbitrary volume samples such as Pixar's RenderMan, or when rendering multiple independent media on graphics hardware. The function t can then be reconstructed:

$$t(x) = \sum_{i=0}^{\infty} t_i c_i(x) \quad (6)$$

Plugging Equation 6 into Equation 4 we obtain a coefficient-based representation of our pseudometric:

$$g(x_1, x_2) = \int_{x_1}^{x_2} \sum_{i=0}^{\infty} t_i c_i(x) dx \quad (7)$$

$$= \sum_{i=0}^{\infty} t_i (C_i(x_2) - C_i(x_1)) \quad (8)$$

Total coeff.	$t(x)$	$g(x)$	$T(x)$	$T(g(x))$
4	81.63%	14.21%	10.79%	9.08%
8	60.52%	7.08%	4.92%	1.28%
12	46.50%	3.56%	2.75%	0.59%

Table 1: RMS error of the reconstruction of the binary function t , the pseudometric g and the transmittance T of Figure 3 without and with our pseudometric.

For conciseness we use $g(x) = g(0, x)$ from this point on. Using this pseudometric, any function $f(x)$ can be projected into another function basis $\{b_i(x)\}_{i \in \mathbb{N}}$ as follows:

$$f_i = \int_{g(0)}^{g(D)} f(u) b_i(u) du, \quad f(x) = \sum_{i=0}^{\infty} f_i b_i(g(x)) \quad (9)$$

We implemented function expansion into Fourier cosine series, whose basis functions are defined as:

$$b_i(x) = \cos\left(\frac{i\pi x}{D}\right), \quad B_i(x) = \frac{D}{i\pi} \sin\left(\frac{i\pi x}{D}\right) \quad (10)$$

where $x \in [0, D]$, $D > 0$ and B_i is the indefinite integral of b_i . For conciseness we choose $c_i(x) = b_i(x)$, and use only $b_i(x)$ to identify basis functions in the remainder of the document. As the practicality of our method comes from the use of non-localized basis functions with a closed form integral, any other basis fulfilling these requirements could also be used with similar results.

The exact representation of the binary function t within a function basis would require an infinite number of projection coefficients. The use of a band-limited projection results in inevitable oscillations in the reconstructed pseudometric (Figure 3, top). However, the purpose of the *boundary-aware pseudometric* is only the reduction of the empty spaces between participating media. As shown in Equation 9 the same pseudometric is used for both projecting and reconstructing the transmittance function. This ensures consistency between the original and projected signals regardless of the potential ringing present in the representation of g . As shown in Table 1, even few coefficients significantly improve the quality of the representation of the transmittance function for a given total number of coefficients.

5. Boundary-Aware Extinction Maps

Following Jansen and Bavoil [JB10], we reformulate the transmittance (Equation 1) by considering the integrand $\sigma_t(x)$ in Euclidean space. Using an arbitrary set of extinction values obtained by sampling the participating media, we project the extinction function into a function basis $\{b_i(x)\}_{i \in \mathbb{N}}$:

$$\sigma_t(x) = \sum_{i=0}^{\infty} s_i b_i(x), \quad s_i = \sum_{j=0}^{\infty} \sigma_t(x_j) \int_{x_{j-}}^{x_{j+}} b_i(x) dx \quad (11)$$

Figure 3: Plots of the projected binary function t , analytically reconstructed pseudometric g , transmittance T without using the pseudometric and transmittance $T(g)$ using the pseudometric. The theoretical values are shown in red, while blue curves indicate the approximations of the functions using $n = 4, 8$ and 12 coefficients from left to right. The bottom row is computed using $n/2$ coefficients for the pseudometric map and $n/2$ coefficients for the extinction map. The use of the pseudometric reduces the oscillations in the reconstructed transmittance function, hence improving the scattering estimation (Figure 7).

FOM are based on infinitely thin particles yielding Dirac impulses on the extinction signal. Our approach tackles the more general problem of volumes composed of arbitrary sets of thick samples through which light gets scattered. In the spirit of the previous section, we use analytical integration to express the transmittance in terms of the projected extinction function:

$$T(x) = \exp \left(- \int_0^x \sum_{i=0}^{\infty} s_i b_i(y) dy \right) \quad (12)$$

$$= \exp \left(- \sum_{i=0}^{\infty} s_i (B_i(x) - B_i(0)) \right) \quad (13)$$

The transmittance (or optical depth) expressing the effect of each volume sample on the entire light ray is then recovered analytically from the projection coefficients.

Our approach also underlines a major efficiency difference with TFM. In Equation 1 the exponential of the transmit-

tance function is embedded within the TFM coefficients. Conversely, the exponential remains analytical in Equation 12. The coefficients then only encode the lower energy features representing the variations of the extinction along light rays, hence surpassing the quality of TFM while offering the genericity of FOM (Figures 5, 6).

Even though Fourier-based representations can only encode the lower frequency variations of the signal, the original high frequencies of the extinction are summed to obtain the transmittance (Equation 1). Therefore, the integration averages out the high frequency variations, preserving only the main, lower frequency components. In our approach the high frequencies are already averaged by the Fourier-based representation of the extinction function, while the integration is performed analytically. Combined with the analytical exponential, this representation effectively improves the overall quality of the reconstructed transmittance. Compared to Adaptive Volumetric Shadow Maps [SVLL10] this representation provides a closer fit to the original transmittance function using a similar storage size (Figure 4).

Further quality can be obtained by following Equation 11 to perform the projection within our pseudometric space instead of the classical Euclidean space:

$$s_i = \sum_{j=0}^{\infty} \sigma_i(x_j) \int_{x_{j-}}^{x_{j+}} b_i(g(x)) dg(x_{j-}, x) \quad (14)$$

Assuming $B_i(0) = 0$, we rewrite the transmittance function using this representation, yielding the formula for Boundary-Aware Extinction Maps:

$$T(x) = \exp \left(- \sum_{i=0}^{\infty} s_i B_i(g(x)) \right) \quad (15)$$

The compression of the integration domain provided by our pseudometric relaxes the relationship between the size of the domain and the representation quality. In particular, unlike FOM and TFM, participating media can be spaced arbitrarily with significantly lower impact on image quality (Figure 7). Further insights on the projection quality of the pseudometric are discussed in Section 10.

Figure 4: Our extinction projection method closely matches the original transmittance curve using 12 coefficients, yielding a RMS error of 3.9%. TFM and AVSM represent the same transmittance with respective error rates of 76% and 13.6% using 12 coefficients/nodes.

Figure 5: Compared to the reference values (red), our extinction-based method reconstructs transmittance with less oscillations than TFM (Figure 7).

Our representation of extinction along a set of rays using a BAEM provides a simple means of evaluating the transmittance function along any light ray, without prior knowledge of the structure, number and arrangement of participating media. For any light source compliant with the shadow mapping approach, Extinction-Transmittance Maps can be used to generate compelling lighting and shadowing effects. Compared to the computation of a TFM for each medium, our approach provides a simple and unified solution inherently accounting for both overlapping media and vacuum sections in between.

6. Ray Marching Arbitrary Representations

Extinction representation techniques such as DSM and BAEM provide an approximate representation of volumetric extinction at any point visible from the light source. This representation can be used to perform further light scattering computations independently from the original volume representation. While on-the-fly scene voxelization could also be

Figure 6: Plumes of smoke rendered using Boundary-Aware Extinction Maps without pseudometric and TFM, both with 4 coefficients per texel. The right part of the images show the difference with a reference solution, with pixel intensities multiplied by 40. TFM reconstruction errors typically appear as horizontal stripes in the denser smoke (bottom).

performed, massive scenes would require vast amounts of memory and processing power while potentially introducing aliasing. In this section we consider the use of BAEM for fast estimation of light scattering using ray marching.

The band-limited representation and projection direction used by BAEM tends to result in non-negligible smoothing of the input media (Figure 8). However, light scattering can be seen as a low-pass filter, blurring out the high frequency details of the media [dLE07, d11]. Also, the convolution of attenuated radiance with the medium phase function along light paths can be connected to the computation of indirect illumination, in which the incoming lighting is convolved with the reflectance function. In this context, indirect lighting computations can be performed on a coarse model without a significant loss of precision [TL04].

Therefore, BAEM can be seen as a viable alternate representation enabling ray marching within otherwise non-marchable representations such as particle clouds. We il-

Figure 7: The projection of the extinction function into a non-localized function basis tends to introduce oscillations along long paths, resulting in undesired dimming of the volumes farthest from the light source (a). Our pseudometric cancels out those empty spaces, avoiding such artifacts (b).

Figure 8: Opacity visualization of a medium using an explicit voxel-based description (a) and BAEM (b, c). While (b,c) lack the high frequency features of the original medium, our representation provides a simple approximate solution for tracing light paths through arbitrary media.

illustrate this principle with the computation of image-based lighting and multiple scattering.

6.1. Image-Based Lighting

Fast rendering of heterogeneous participating media under image-based lighting is a challenging task. Light potentially comes from any direction, generating the complex lighting and shadowing effects visible in the real world.

The radiance scattered from any point $\mathbf{p}$ in direction ω is:

$$L_i(\mathbf{p}, \omega) = \sigma_s(\mathbf{p}) \int_{\Omega} L(\omega_i) T(\mathbf{p}, \omega_i) p(\mathbf{p}, \omega_i, \omega) d\omega_i \quad (16)$$

$$= \sigma_s(\mathbf{p}) \int_{\Omega} L_{ri}(\omega_i) p(\mathbf{p}, \omega_i, \omega) d\omega_i \quad (17)$$

where Ω is the sphere of directions surrounding $\mathbf{p}$ and p is the phase function of the medium at $\mathbf{p}$. This formulation highlights the need for sampling the entirety of the media to determine the scattered radiance at any point. While samples and light transfer could be precomputed for real-time rendering, this would preclude rendering dynamic media.

We extend the principle of Delalandre *et al.* [DGM13] by evaluating the transmittance around any point by marching through the BAEM in a set of directions covering the entire sphere. This approach reduces the costs of transmittance evaluation using a simple caching strategy. For a number of points in the scene the transmittance is evaluated in a number of directions and projected into the spherical harmonics basis [Mül66]. The overall incoming radiance at any point is then deduced by interpolating the closest spherical harmonics records and using the orthonormality of the basis to efficiently compute the dot product of the attenuated radiance L_{ri} with the phase function p of the medium.

The fast evaluation of the extinction function allows our method to regenerate the spherical harmonics records for each frame. Depending on the application the records can be evenly spaced in the scene, or organized in a raster-oriented grid for more efficiency. Figure 9 compares our approach with a reference solution obtained by classical ray marching through a voxelized medium. We compared those images us-

Figure 9: For each frame, the transmittance is projected into 9 spherical harmonics coefficients at 50^3 uniformly distributed points, and interpolated for fast image-based lighting (a). We render an animated smoke (512^3) using ray marching through a voxel grid (b) and the BAEM (c).

ing the Structured SIMilarity method [WBSS04], yielding an accuracy of 99.62% with a BAEM resolution of 1024^2 .

6.2. Multiple Scattering

A classical method for estimating multiple scattering is Monte Carlo path tracing, in which the choices of the random direction and distance to the next scattering event are driven by the optical properties of the medium [HK93, Sta95]. We then carry out the integration along arbitrary rays by marching through the BAEM and evaluating the medium extinction at any point directly from the projection coefficients (Equation 14). Figure 10 illustrates the use of BAEM for providing a reliable estimate of the extinction for the purpose of dual bounce multiple scattering computation. Note that an arbitrary number of light bounces can be added by further marching through the BAEM. Besides this simple approach, our representation can also benefit more advanced solutions [Fat09, PAT*04].

Figure 10: Dual bounce multiple scattering using ray marching in a voxelized medium (a) and BAEM (b,c). The right part of the images show image differences multiplied by 40. Images (b) and (c) are respectively rendered at 5 fps using 10 direction samples for each step of the marching (100 steps per pixel).

7. Applications

Our solution provides a generic framework for fast estimation of light scattering, suitable for a wide range of applications. In this section we apply our approach to the visualization of massive volumetric environments.

7.1. Cloudy Skies

A cloudy sky is composed of a large number of participating media of various sizes and shapes, potentially overlapping or

separated by large distances. Those media cast and receive soft shadows, yielding complex lighting effects (Figure 11).

Regardless of the representation of the volumes, the BAEM represent the extinction function and analytically integrate the function to obtain the transmittance at any point visible from the light source. The empty space between the clouds gets mostly canceled out by our boundary-aware pseudometric. The large volumetric dataset is adaptively streamed using a simple out-of-core technique.

7.2. Out-Of-Core Volumes

Scenes featuring many complex participating media are common in the world of production rendering, but the media are usually generated independently and arranged using crude representations. The lighting design then becomes a tedious task involving numerous overnight test renders to achieve the desired result. Furthermore, the complexity of even a single medium may exceed the memory available on current graphics hardware. Many approaches consider real-time visualization of large static volumetric datasets [CNLE09, GMI08], overlooking dynamic changes and scattering simulation. We describe a simple yet efficient out-of-core method for rendering massive volumetric scenes using Boundary-Aware Extinction Maps.

Volume Representation Our method is based on a decomposition of each participating medium into an octree representation [Mea82] in which nodes are bricks of fixed resolution. In practice, 64^3 bricks proved to be a satisfying compromise between granularity and octree depth. While leaf nodes represent the data at its highest resolution, the bricks of inner nodes are computed from bottom up by filtering children bricks. This technique provides a simple way of representing multiple levels of detail on a volume. Also, the quality of the representation can be adjusted by adaptively refining the structure at the desired locations.

Note that instead of traversing the octree on graphics hardware we consider each internal node of the octree as a *localized* mip level of the overall volume. This solution simultaneously addresses two issues: first, parts of the volumes

Figure 11: Sky comprising 75 media (512^3 each), rendered at 1.5fps with single scattering and image-based lighting using BAEM with adaptive out-of-core rendering.

© 2013 The Author(s)

Computer Graphics Forum © 2013 The Eurographics Association and John Wiley & Sons Ltd

can be loaded independently for fast rendering of arbitrarily complex media. Second, mip levels are intrinsically introduced by the octree structure, providing the high quality filtering capabilities required in production rendering.

The octree is then stored into a coherent hard disk representation, in which the contents of each node can be accessed linearly. Note that while other volume subdivision techniques [LW90, DH92, YIC*10, YIC*11, Mus13] could also be used, we chose octrees for their ease of implementation and filtering.

Rendering The representation described above can be seen as a subdivision of a main participating medium into a set of smaller, independent media casting and receiving shadows. The subdivision level of each branch of the octree is determined using the projected area of the node bounding boxes [ST99], and loaded adaptively from the hard disk.

Complex scenes may contain numerous, potentially overlapping participating media. Consequently the octree subdivisions result in a large number of independent nodes. Using our approach those nodes are seamlessly combined regardless of their relative positions or rendering order.

8. Implementation

The generation of the BAEM is split into two passes. The pseudometric map is first generated by rendering the bounding boxes of each medium of the scene as seen from the light source. For each fragment we intersect the corresponding ray with the box. Using Equation 5 the contribution of the medium to the pseudometric is computed and combined to the contributions of the other media using additive blending. The resulting coefficients for each ray are stored within a floating-point buffer in GPU memory.

The generation of the extinction map involves the evaluation of the extinction values. For each brick we perform an adaptive ray marching (1-50 steps depending on the traversal length through the medium) and evaluate the pseudometric at each step using the previously computed coefficients. The step contribution to the projection coefficient of the extinction is evaluated using Equation 14. The combined contributions are returned by the fragment shader before being additively blended with the contributions of the other media. In the case of arbitrary media only available through sampling (eg. particle clouds), the contributions of each sample are simply computed using Equation 14 and blended into a floating-point GPU buffer.

Final rendering requires back to front sorting to ensure a proper blending of the visible media, as is the case when rasterizing any set of transparent objects. For each brick or volume sample we evaluate single scattering by fetching the coefficients for both the pseudometric and the extinction maps and analytically integrating the BAEM (Equation 15). The blending factor is the transmittance of the brick or sample.

Scene Contents	Dynamic Lighting	Static Lighting
512 ³	17.24 fps	36.8 fps
1024 ³	12.00 fps	27.5 fps
3 × 512 ³	11.20 fps	16.9 fps
6 × 512 ³	6.21 fps	9.31 fps

Table 2: Rendering performance of BAEM using production-quality media with mutual shadowing.

9. Results

We implemented Boundary-Aware Extinction Maps on a Intel Xeon X5680 3.36GHz processor running a Nvidia GeForce GTX580 GPU. The timings are measured when generating images with resolution 1280 × 720. The BAEM has a resolution of 1024 × 1024 with a total of 12 half precision floating point values per texel, yielding a 96MB GRAM footprint.

We first consider multiple production-quality volumes with mutual shadowing (Figure 7). Our solution provides interactive to real-time performance in those cases (Table 2), even when refreshing the BAEM for each frame.

Our main test scene features a main smoke medium with resolution 1024³ surrounded by 74 cloud and smoke volumes (Figures 1 and 11), resulting in 10.5GB of volumetric data. The octree decomposition of the media took 8.8s per 512³ medium, and 50.3s for the 1024³ smoke. The set of nodes is then considered as a large set of independent participating media. Single scattering and image-based lighting were estimated at 1.5 fps using a budget of 1GB of graphics memory. As shown in Figure 12 the media are adaptively refined with respect to the viewpoint. Our boundary-aware pseudometric finds a particular use for rendering clouds under the overhanging smoke and lava bomb.

The occasional popping visible in the accompanying video has two origins, both being orthogonal to our method. First, the octrees representing the media are refined on the fly by replacing obsolete nodes depending on the viewpoint location and hardware performance. A smoother blending could reduce the visibility of these artifacts. Second, the nodes are rendered from back to front in the final image to ensure a consistent visual blending of the translucent objects, as is the case with many real-time renderers. The sorting instability then introduces temporally sharp changes in overlapping objects. This issue could be solved using an order-independent transparency technique [BM08].

The BAEM find a particular use in this scene due to the decomposition of the media into a large number of subsets. The pseudometric allows our method to efficiently avoid the premature transmittance drop plotted in Figure 3. This results in the avoidance of unnatural light dimming due to the oscillations of the representation (Figure 7).

Movie post-production makes intensive use of complex volumetric assets, requiring fast yet accurate previsualization tools for scene and lighting setup. We then refresh the

BAEM for each frame to support dynamic environments and ensure consistency between the BAEM and the final image. As shown in Figure 13, our solution provides images consistent with Pixar’s RenderMan using DSM in terms of tint and shadowing features.

10. Discussion

Pseudometric Properties A pseudometric g must satisfy the triangle inequality $g(x, y) + g(y, z) \geq g(x, z)$. As shown in Figure 3, the representation of the Boundary-Aware Pseudometric tends to feature oscillations, hence failing to verify this inequality. In practice those errors may result in nonzero extinction values “bleeding” backward along light paths. However, those oscillations tend to be mostly located on the constant sections of the pseudometric (Figure 3), making most reconstruction errors occur in the vacuum between participating media. Furthermore, due to the properties of the cosine function, the average of the oscillations is very small and does not introduce noticeable artifacts in our test scenes.

Pseudometric Ringing In TFM and FOM visible oscillations appear in the image as the reconstructed transmittance is inconsistent with the original signal. In the case of BAEM the pseudometric is only a space warping function reducing the gaps between media, with the goal of increasing the consistency of the projected and reconstructed extinction functions in the final image. The projected pseudometric is then used consistently between projection and reconstruction, making the potential oscillations of this function negligible.

Filtering The extinction functions stored in neighboring texels of the BAEM are defined in spaces involving different pseudometrics, and require a full reprojection of the extinction for precise filtering. A less accurate yet faster solution is a direct interpolation of both the pseudometric and the extinction function coefficients (Figure 14). As g is coherent in projective space, this approximation did not lead to visible artifacts in our images. A more precise derivation of filtering for BAEM is left for future work.

Figure 12: The participating media are adaptively loaded and refined according to their projected area. Each node of the octrees is rendered independently into the BAEM.

(a) BAEM, 12 coefficients (b) Pixar's RenderMan
Figure 13: BAEM (a) provide an interactive estimate of production-quality images obtained using DSM (b).

Figure 14: The BAEM texels express the transmittance using distinct pseudometrics (a,b). An accurate filtering requires a costly reprojection (c, red), while a direct combination of both coefficients sets (c, blue) may introduce errors (top). However, neighboring texels involve similar pseudometrics (bottom), making such errors negligible in practical cases.

Overlapping Media The media present in complex environments usually overlap (Figure 1). While TFM require a simultaneous access to all overlapping media to evaluate the transmittance function, BAEM implicitly handle such overlaps. The strength of our method is that extinction samples are simply combined into the projected extinction function coefficients using Equation 14. Therefore, BAEM do not introduce any overhead in the presence of overlapping media.

Limitations and Future Work BAEM also involve a number of limitations: In particular the generation of the pseudometric introduces an additional pass compared to FOM and TFM. However, this pass only requires rendering the bounding boxes of the media using a simple shader. In our tests this additional cost appeared to be very small compared to the rendering of the volumetric data in the subsequent passes.

Band-limited representations of high frequency signals usually introduce ringing and undesired smoothing. As BAEM may excessively smooth fine shadows in very dense media, adapting the pseudometric to the medium densities could provide enhanced results. Further removal of empty path sections could be achieved using a tree-based space subdivision to estimate the distance to the first intersected medium.

11. Conclusion

We introduced Boundary-Aware Extinction Maps for efficient estimation of light scattering within arbitrary participating media. Based on a projection of the extinction function using a Boundary-Aware Pseudometric, our approach represents both local extinction and global transmittance using a single set of projection coefficients. We then compute single scattering efficiently by analytical integration on graphics hardware for interactive high quality rendering.

Boundary-Aware Extinction Maps also provide a unified representation of the volumetric datasets of the scene, compatible with many other rendering methods. We apply this feature to interactive image-based lighting and multiple scattering estimation, opening the way towards interactive implementations of more advanced rendering techniques.

Our technique finds a particular use for real-time editing and navigation in production environments featuring massive volumetric datasets. As the Boundary-Aware Extinction Maps do not introduce any precomputation, the components of the scene can be modified in real-time.

References

- [BM08] BAVOIL L., MYERS K.: *Order independent transparency with dual depth peeling*. Tech. rep., NVIDIA, 2008. 8
- [BNM*08] BOUTHORS A., NEYRET F., MAX N., BRUNETON E., CRASSIN C.: Interactive multiple anisotropic scattering in clouds. In *Proceedings of I3D* (2008), pp. 173–182. 2
- [BPPP05] BOUDET A., PITOT P., PRATMARTY D., PAULIN M.: Photon splatting for participating media. In *Proceedings of GRAPHITE* (2005), pp. 197–204. 1
- [CFP*05] CEREZO E., FRANCISCO S., PEREZ F., SILLION F., PUEYO X.: A survey on participating media rendering techniques. *Visual Computer* 21, 5 (2005), 303–328. 1
- [Cha50] CHANDRASEKHAR S.: *Radiative transfer*. Clarendon Press. 1
- [CNLE09] CRASSIN C., NEYRET F., LEFEBVRE S., EISEMANN E.: Gigavoxels : Ray-guided streaming for efficient and detailed voxel rendering. In *Proceedings of I3D* (2009). 7
- [DGM13] DELALANDRE C., GAUTRON P., MARVIE J.-E.: Real-Time Image-Based Volume Lighting. In *Proceedings of Eurographics* (2013), pp. 57–60. 6
- [DGMF11] DELALANDRE C., GAUTRON P., MARVIE J.-E., FRANÇOIS G.: Transmittance function mapping. In *Proceedings of I3D* (2011), pp. 31–38. 2, 3
- [DH92] DANSKIN J., HANRAHAN P.: Fast algorithms for volume ray tracing. In *Proceedings of the Workshop on Volume visualization* (1992), pp. 91–98. 7
- [dI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. In *Proceedings of SIGGRAPH* (2011), pp. 56:1–56:14. 5
- [DKH*13] DACHSBACHER C., KŘIVÁNEK J., HAŠAN M., ARBREE A., WALTER B., NOVÁK J.: Scalable realistic rendering with many-light methods. In *Proceedings of Eurographics - STAR program* (2013). 1

- [dLE07] D'EON E., LUEBKE D., ENDERTON E.: Efficient rendering of human skin. In *Proceedings of Eurographics Symposium on Rendering* (2007), pp. 147–157. 5
- [EHK*06] ENGEL K., HADWIGER M., KNISS J., REZK-SALAMA C., WEISKOPF D.: *Real-time Volume Graphics*. AK Peters, 2006. 1
- [Fat09] FATTAL R.: Participating media illumination using light propagation maps. *ACM Trans. Graph.* 28, 1 (2009), 1–11. 6
- [GMI08] GOBBETTI E., MARTON F., IGLESIAS GUITIÁN J.: A single-pass GPU ray casting framework for interactive out-of-core rendering of massive volumetric datasets. *The Visual Computer* 24, 7-9 (2008), 797–806. 7
- [HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH* (1993), pp. 165–174. 6
- [JB10] JANSEN J., BAVOIL L.: Fourier opacity mapping. In *Proceedings of I3D* (2010), pp. 165–172. 2, 3
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of SIGGRAPH* (1998), pp. 311–320. 1
- [JHH*09] JOHNSON G. S., HUNT W. A., HUX A., MARK W. R., BURNS C. A., JUNKINS S.: Soft irregular shadow mapping: fast, high-quality, and robust soft shadows. In *Proceedings of I3D* (2009), pp. 57–66. 2
- [JNSJ11] JAROSZ W., NOWROUZEZAHRAI D., SADEGHI I., JENSEN H. W.: A comprehensive theory of volumetric radiance estimation using photon points and beams. In *Proceedings of SIGGRAPH* (2011), vol. 30, pp. 5:1–5:19. 1
- [JSYR12] JÖNSSON D., SUNDÉN E., YNNERMAN A., ROPINSKI T.: Interactive Volume Rendering with Volumetric Illumination. In *Eurographics STAR program* (2012). 1
- [JZJ08] JAROSZ W., ZWICKER M., JENSEN H. W.: The beam radiance estimate for volumetric photon mapping. In *Proceedings of Eurographics* (2008), vol. 27, pp. 557–566. 1
- [Kel97] KELLER A.: Instant radiosity. In *Proceedings of SIGGRAPH* (1997), pp. 49–56. 1
- [KN01] KIM T.-Y., NEUMANN U.: Opacity shadow maps. In *Proceedings of Eurographics Workshop on Rendering* (2001), pp. 177–182. 2
- [KPH*03] KNISS J., PREMOZE S., HANSEN C., SHIRLEY P., MCPHERSON A.: A model for volume lighting and modeling. *IEEE Trans. on Vis. and Comp. Graph.* 9, 2 (2003), 150–162. 2
- [KVH84] KAJIYA J. T., VON HERZEN B. P.: Ray tracing volume densities. *Proceedings of SIGGRAPH* 18, 3 (1984), 165–174. 1
- [LV00] LOKOVIC T., VEACH E.: Deep shadow maps. In *Proceedings of SIGGRAPH* (2000), pp. 385–392. 1, 2, 3
- [LW90] LEVOY M., WHITAKER R.: Gaze-directed volume rendering. *Proceedings of SIGGRAPH* 24, 2 (1990), 217–223. 7
- [Mea82] MEAGHER D.: Geometric modeling using octree encoding. *Computer Graphics and Image Processing* 19, 2 (1982), 129–147. 7
- [Mit05] MITCHELL J. L.: *ShaderX3: Light Shaft Rendering*. Charles River Media, 2005, pp. 573–588. 2
- [Mül66] MÜLLER C.: *Spherical Harmonics*. Springer, 1966. 6
- [Mus13] MUSETH K.: VDB: High-resolution sparse volumes with dynamic topology. *ACM Transaction on Graphics* 32, 3 (2013), 27:1–27:22. 7
- [NJS*11] NOWROUZEZAHRAI D., JOHNSON J., SELLE A., LACEWELL D., KASCHALK M., JAROSZ W.: A programmable system for artistic volumetric lighting. In *Proceedings of SIGGRAPH* (Aug. 2011), vol. 30, pp. 29:1–29:8. 1
- [NNDJ12a] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Progressive virtual beam lights. In *Proceedings of EGSR* (2012), pp. 1407–1413. 1
- [NNDJ12b] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Virtual ray lights for rendering scenes with participating media. In *Proceedings of SIGGRAPH* (2012), pp. 60:1–60:11. 1
- [PAT*04] PREMOZE S., ASHIKHMIN M., TESSENDORF J., RAMAMOORTHY R., NAYAR S.: Practical rendering of multiple scattering effects in participating media. In *Proceedings of Eurographics Symposium on Rendering* (2004), pp. 363–374. 6
- [PH89] PERLIN K., HOFFERT E. M.: Hypertexture. In *Proceedings of SIGGRAPH* (1989), pp. 253–262. 1
- [SS70] STEEN L. A., SEEBACH J. A. J.: *Counterexamples in Topology*. Holt, Rinehart and Winston, Inc., 1970. 3
- [ST99] SCHMALSTIEG D., TOBLER R. F.: Fast projected area computation for three-dimensional bounding boxes. *Journal of graphics, gpu, and game tools* 4, 2 (1999), 37–43. 7
- [Sta95] STAM J.: Multiple scattering as a diffusion process. In *Proceedings of Eurographics Workshop on Rendering* (1995), pp. 41–50. 6
- [SVLL10] SALVI M., VIDIMČE K., LAURITZEN A., LEFOHN A.: Adaptive volumetric shadow maps. In *Eurographics Symposium on Rendering* (2010), pp. 1289–1296. 2, 4
- [TL04] TABELLION E., LAMORLETTE A.: An approximate global illumination system for computer generated films. In *Proceedings of SIGGRAPH* (2004), pp. 469–476. 5
- [WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. 6
- [Wil78] WILLIAMS L.: Casting curved shadows on curved surfaces. In *Proceedings of SIGGRAPH* 12, 3 (1978), 270–274. 1
- [WR08] WYMAN C., RAMSEY S.: Interactive volumetric shadows in participating media with single-scattering. In *Proceedings of Symposium on Interactive Ray Tracing* (2008), pp. 87–92. 2
- [WZH*11] WRENNINGE M., ZAFAR N. B., HARDING O., GRAHAM G., TESSENDORF J., GRANT V., CLINTON A., BOUTHORS A.: Production volume rendering. *SIGGRAPH Courses* (2011). 1
- [YHGT10] YANG J. C., HENSLEY J., GRÜN H., THIBIEROZ N.: Real-time concurrent linked list construction on the GPU. *Computer Graphics Forum* 29, 4 (2010), 1297–1304. 2
- [YIC*10] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. In *Proceedings of SIGGRAPH Asia* (2010), pp. 177:1–177:8. 7
- [YIC*11] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Toward optimal space partitioning for unbiased, adaptive free path sampling of inhomogeneous participating media. *Computer Graphics Forum* 30, 7 (2011), 1911–1919. 7
- [YK08] YUKSEL C., KEYSER J.: Deep opacity maps. *Computer Graphics Forum* 27, 2 (2008). 2
- [ZRL*08] ZHOU K., REN Z., LIN S., BAO H., GUO B., SHUM H.-Y.: Real-time smoke rendering using compensated ray marching. In *Proceedings of SIGGRAPH* (2008), pp. 1–12. 2