

Adaptive Mesh Subdivision for Precomputed Radiance Transfer

Jaroslav Křivánek*
Univ. of Central Florida
Czech Technical University
INRIA/IRISA

Sumanta Pattanaik†
Univ. of Central Florida

Jiří Žára‡
Czech Technical University

Abstract

In precomputed radiance transfer (PRT), the radiance transfer operator is precomputed on vertices of a mesh. During rendering the transfer operator is combined with illumination to produce global illumination effects at real-time frame rates. However, visible error can be introduced by interpolating the transferred radiance between vertices. We propose to adaptively subdivide meshes with PRT in order to avoid the interpolation error. The mesh density is increased where the transferred radiance can change rapidly and introduce high error. We exploit the fact that the illumination frequency in PRT is bounded and thus we are able to perform the subdivision during the preprocessing phase. This obviates the need for dynamic mesh subdivision at render-time. The adaptive subdivision is guided by an error measure based on the transfer operator.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Shading, Shadowing

Keywords: Adaptive subdivision, Spherical Harmonic Lighting, Precomputed Radiance Transfer

1 Introduction

Precomputed radiance transfer (PRT) [Sloan et al. 2002; Sloan et al. 2003] is a technique to compute lighting on 3D models from arbitrary low-frequency directional light

sources in real-time. PRT supports advanced global illumination effects such as soft-shadows, interreflections and subsurface scattering. Spherical harmonics (SH) are used in PRT to represent illumination and bidirectional reflectance distribution functions (BRDFs) [Westin et al. 1992; Kautz et al. 2002; Ramamoorthi 2002].

There is an extensive preprocess step in PRT which simulates how light is shadowed and/or reflected before it reaches a point on an object’s surface and how BRDF turns it into reflected light. Results of the preprocess are stored in the form of one transfer matrix \mathbf{M}_p per “illumination point” \mathbf{p} . The illumination points can be either object vertices or texels of a texture mapped onto the object.

At render-time, the transfer matrix \mathbf{M}_p is multiplied by the SH coefficient vector \mathbf{l} representing the illumination. The resulting vector $\mathbf{l}_{tr} = \mathbf{M}_p \mathbf{l}$ is then dotted with the vector of SH bases evaluated for the viewing direction. This is how the reflected luminance (color) is computed for the illumination points (vertices, texel centers). The luminance of image pixels to which no illumination point projects is computed by linear interpolation.

The illumination points (points at which the transfer matrix was computed) form point sampling of the transfer operator defined over whole object surface. If the sampling density is not high enough, artifacts occur in the reconstructed shading (Figure 1(a)). A possible solution to the transfer operator undersampling is to increase the sampling density by uniformly subdividing triangles or increasing texture resolution. However, in most cases this wastes resources because the rate of change of the transfer operator is usually highly non-uniform over the model surface (Figure 1(b)).

We propose to adaptively subdivide meshes with precomputed radiance transfer in a way that reflects the rate of change of the transfer operator (Figure 1(c)). Our aim is to distribute the error due to interpolation uniformly over the object surface by subdividing more where the transfer operator changes rapidly. Using the fact that illumination frequency in PRT (as in any other SH-lighting technique) is bounded, we can subdivide in the preprocessing phase. This wouldn’t be possible if illumination frequency were not bounded because of possible sharp shadows. Since

*xkrivanj@fel.cvut.cz

†sumant@cs.ucf.edu

‡zara@fel.cvut.cz




Figure 1: Adaptive subdivision for the ‘column’ model. The first two rows show the model under two different illuminations, the third row shows the triangulations. (a) Original model contains 200 vertices and 246 triangles — too few to capture the shape of the column shadow correctly. (b) Uniform subdivision (1565 vertices and 2584 triangles) increases the quality, but there are still interpolation artifacts visible on the shadow boundary near the column. (c) Adaptive subdivision (1188 vertices, 2828 triangles) refines the mesh near the column and the shadow is reconstructed with less visible artifacts.

we subdivide in preprocess, we do not know the particular illumination that will be used to shade the object. We therefore estimate the worst case interpolation error for triangles over all possible illuminations and use this estimated error as a priority of triangle subdivision. To our knowledge this is the first work that addresses the problem of transfer operator sampling.

We limit ourselves to the PRT on Lambertian surfaces with self-shadowing only (no interreflections, no subsurface scattering). In this case the transfer matrix reduces to a transfer vector \mathbf{t}_p and the outgoing radiance at \mathbf{p} is computed as a dot product of transfer and illumination vectors $\langle \mathbf{t}_p | \mathbf{l} \rangle$.

2 Related work

The related work on precomputed radiance transfer is [Sloan et al. 2002; Sloan et al. 2003; Kautz et al. 2002; Lehtinen and Kautz 2003], but those papers do not address the sampling density of the radiance transfer operator.

There is much work on adaptive subdivision in the field of radiosity, e.g. [Myszkowski et al. 1994; Martin et al. 1997; Gibson and Hubbard 1997]. None of these approaches directly applies to our problem. The fundamental difference lies in the fact that in radiosity, the illumination is fixed and known at subdivision time. The subdivision criterion can therefore be based on a particular

illumination. We face a different problem: the transfer operator is fixed, but the illumination is not. We do not know the illumination at subdivision time and the criterion must therefore be based solely on the transfer operator.

Another related area is dynamic level-of-detail that takes shading into account [Xia and Varshney 1996; Cho et al. 1996; Klein et al. 1998; Williams et al. 2003]. Unlike in these approaches, our aim is to perform the subdivision in preprocess and obviate dynamic level-of-detail selection at render-time. We can do this because the illumination frequency in PRT is bounded.

3 Our Contribution

3.1 Problem statement

Given a triangle $\Delta = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ we want to assess the error introduced by interpolating the results of PRT from vertices $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ instead of computing the PRT at any point of the triangle. We are looking for the maximum of this error over all points $\mathbf{x} \in \Delta$ and all “unit power” directional lighting functions ℓ :

$$\varepsilon_\Delta = \max_{\|\ell\|=1, \mathbf{x} \in \Delta} \frac{|L(\mathbf{x}, \ell) - L_{\text{int}}(\mathbf{x}, \ell)|}{L(\mathbf{x}, \ell)}, \quad (1)$$

where $L(\mathbf{x}, \ell)$ would be the luminance at \mathbf{x} if PRT were precomputed at \mathbf{x} and $L_{\text{int}}(\mathbf{x}, \ell)$ denotes the interpolated luminance at \mathbf{x} . We divide by $L(\mathbf{x}, \ell)$ because, according to the Weber law [Ferwerda 2001], we are interested in a relative error measure rather than an absolute one.

Such an error metric can be used, e.g., in the following ways:

- Color coded ε is presented to the graphics designer so that he or she knows where the model should be refined.
- ε is used to guide an automatic mesh subdivision algorithm.

Here we deal with the automatic mesh subdivision.

3.2 Mesh subdivision

The input of the subdivision algorithm is a mesh and the maximum allowed number of triangles (triangle budget). The output is the mesh adaptively subdivided to minimize the interpolation error. The per-vertex transfer vectors are computed concurrently with the mesh subdivision. The algorithm steps are shown in Figure 2.

```

for (every mesh triangle  $\Delta_j$ ) do
  Estimate triangle error  $\varepsilon_{\Delta_j}$ .
  if ( $\varepsilon_{\Delta_j} > \text{threshold}$ ) then
    Insert  $(\varepsilon_{\Delta_j}, \Delta_j)$  into the priority queue.
  end if
end for
while (triangle budget not reached AND queue not empty) do
  Extract triangle  $\Delta_{\text{max}}$  with maximum  $\varepsilon$ .
  Subdivide  $\Delta_{\text{max}}$ .
  for (every child  $c_j$  of  $\Delta_{\text{max}}$ ) do
    Estimate child error  $\varepsilon_{c_j}$ .
    if ( $\varepsilon_{c_j} > \text{threshold}$ ) then
      Insert  $(\varepsilon_{c_j}, c_j)$  into the priority queue.
    end if
  end for
end while

```

Figure 2: Adaptive subdivision algorithm.

The priorities are given by the estimate of error described in the next section. The subdivision we apply is a regular subdivision depicted in Figure 3(a). No triangle shape quality metric is used — it is the responsibility of the designer to provide an input model with well shaped triangles. This subdivision scheme can introduce T-vertices, which we remove in a post-processing step by additional triangle subdivision (Fig. 3(b), (c)).

We use this simple subdivision algorithm because the subdivision scheme itself is not the focus of this work. We rather concentrate on the triangle error estimate guiding the subdivision. A more sophisticated subdivision scheme is given, e.g., in [Baum et al. 1991].

3.3 Error Estimate

The error measure we are seeking to estimate is given in Equation 1. We estimate it by assuming it to be proportional to the sum of local contrasts along triangle edges, $\varepsilon_\Delta \approx \varepsilon_{(\mathbf{v}_1, \mathbf{v}_2)} + \varepsilon_{(\mathbf{v}_2, \mathbf{v}_3)} + \varepsilon_{(\mathbf{v}_3, \mathbf{v}_1)}$. The contrast measure for edge $(\mathbf{v}_1, \mathbf{v}_2)$ between vertices \mathbf{v}_1 and \mathbf{v}_2 inspired by the Michelson contrast¹ is

$$\varepsilon_{(\mathbf{v}_1, \mathbf{v}_2)} \approx \max_{\|\ell\|=1} \frac{|L(\mathbf{v}_1, \ell) - L(\mathbf{v}_2, \ell)|}{L(\mathbf{v}_1, \ell) + L(\mathbf{v}_2, \ell)}.$$

In the context of precomputed radiance transfer on Lambertian surfaces, the luminance $L(\mathbf{x}, \ell)$ is given by the dot product $L(\mathbf{x}, \ell) = \langle \mathbf{t}_\mathbf{x}, \mathbf{l} \rangle$, where $\mathbf{t}_\mathbf{x}$ is the transfer vector at \mathbf{x} and \mathbf{l} is the SH coefficient vector of illumination. The edge contrast becomes:

$$\varepsilon_{(\mathbf{v}_1, \mathbf{v}_2)} \approx \max_{\|\mathbf{l}\|=1} \frac{|\langle \mathbf{t}_{\mathbf{v}_1}, \mathbf{l} \rangle - \langle \mathbf{t}_{\mathbf{v}_2}, \mathbf{l} \rangle|}{\langle \mathbf{t}_{\mathbf{v}_1}, \mathbf{l} \rangle + \langle \mathbf{t}_{\mathbf{v}_2}, \mathbf{l} \rangle}$$

¹Defined as $\frac{L_{\text{max}} - L_{\text{min}}}{L_{\text{max}} + L_{\text{min}}}$, where L_{max} is the luminance of a lighter area of interest and L_{min} is the luminance of an adjacent darker area.




Figure 3: Triangle subdivision used in the algorithm. (a) Regular subdivision. (b) Subdivision used in post-process to remove T-vertices. (c) Example of a T-vertex removal using the subdivision form (b).

$$\begin{aligned}
&= \max_{\|\mathbf{l}\|=1} \frac{|\langle \mathbf{t}_{v_1} - \mathbf{t}_{v_2}, \mathbf{l} \rangle|}{|\langle \mathbf{t}_{v_1} + \mathbf{t}_{v_2}, \mathbf{l} \rangle|} \\
&= \max_{\|\mathbf{l}\|=1} \frac{\|\mathbf{t}_{v_1} - \mathbf{t}_{v_2}\| \|\mathbf{l}\| \cos \angle(\mathbf{t}_{v_1} - \mathbf{t}_{v_2}, \mathbf{l})}{\|\mathbf{t}_{v_1} + \mathbf{t}_{v_2}\| \|\mathbf{l}\| \cos \angle(\mathbf{t}_{v_1} + \mathbf{t}_{v_2}, \mathbf{l})} \\
&= \max_{\|\mathbf{l}\|=1} \frac{\|\mathbf{t}_{v_1} - \mathbf{t}_{v_2}\| \cos \angle(\mathbf{t}_{v_1} - \mathbf{t}_{v_2}, \mathbf{l})}{\|\mathbf{t}_{v_1} + \mathbf{t}_{v_2}\| \cos \angle(\mathbf{t}_{v_1} + \mathbf{t}_{v_2}, \mathbf{l})}.
\end{aligned}$$

An additional constraint on the illumination vector \mathbf{l} is that it represents a non-negative function (light is never negative). We want to maximize the function $|\cos \angle(\mathbf{t}_{v_1} - \mathbf{t}_{v_2}, \mathbf{l})| / |\cos \angle(\mathbf{t}_{v_1} + \mathbf{t}_{v_2}, \mathbf{l})|$ over all non-negative unit illuminations. We haven't succeeded in solving this problem analytically and instead we assume that the maximum is a constant independent of \mathbf{t}_{v_1} and \mathbf{t}_{v_2} . We used

$$\mathcal{E}_{(v_1, v_2)} = \frac{\|\mathbf{t}_{v_1} - \mathbf{t}_{v_2}\|}{\|\mathbf{t}_{v_1} + \mathbf{t}_{v_2}\|} \quad (2)$$

as the first measure of the edge error. The disadvantage of this error measure is that it promotes subdivision in shadowed areas too much. This is because of the denominator $\|\mathbf{t}_{v_1} + \mathbf{t}_{v_2}\|$ in Equation 2; the more shadowed the triangle is, the lower $\|\mathbf{t}_{v_1} + \mathbf{t}_{v_2}\|$ is and $\mathcal{E}_{(v_1, v_2)}$ becomes very high even for low difference between the transfer vectors \mathbf{t}_{v_1} and \mathbf{t}_{v_2} (Figure 4(a)).

This observation motivated us to remove the problematic denominator and use the error measure

$$\mathcal{E}_{(v_1, v_2)} = \|\mathbf{t}_{v_1} - \mathbf{t}_{v_2}\|, \quad (3)$$

which is simply the L_2 -distance between vectors \mathbf{t}_{v_1} and \mathbf{t}_{v_2} . The results obtained with this error estimate (Figure 4(b)) correspond to our expectations better than those




Figure 4: Comparison of the error measures (2) and (3). (a) Error measure (2): division by $\|\mathbf{t}_{v_1} + \mathbf{t}_{v_2}\|$ promotes subdivision in shadowed areas. The plane under the box and the bottom of the box are subdivided too much, even though the gradient of illumination in these parts cannot be high. (b) Error measure (3): the denominator from (2) is removed. The bottom of the box is less subdivided. There is more subdivision on the sides of the box and on the plane around the box, where the illumination gradient can be high.

obtained with Eq. 2. By further experimentation we have found that error measure (3) gives good results for flat surfaces, but promotes too many subdivisions on curved surfaces in comparison to flat ones. To rectify this problem we divide measure 3 by the divergence of triangle normals. The final error measure for triangle $\Delta = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ is

$$\mathcal{E}_{\Delta} = \frac{\|\mathbf{t}_{v_1} - \mathbf{t}_{v_2}\| + \|\mathbf{t}_{v_2} - \mathbf{t}_{v_3}\| + \|\mathbf{t}_{v_3} - \mathbf{t}_{v_1}\|}{1 + \alpha \angle \left(\mathbf{n}_1, \frac{\mathbf{n}_1 + \mathbf{n}_2 + \mathbf{n}_3}{\|\mathbf{n}_1 + \mathbf{n}_2 + \mathbf{n}_3\|} \right)}, \quad (4)$$

where α is a user specified constant and $\mathbf{n}_1, \mathbf{n}_2, \mathbf{n}_3$ are triangle normals.




Figure 5: (a) Uniform subdivision (1953 vertices and 3504 triangles). (b) Adaptive subdivision (1540 vertices, 3720 triangles).

4 Results

Figure 1 shows renderings of a model with three different subdivision levels. The first two rows show the model under two different illuminations, the third row shows the triangulations. We used spherical harmonics of order 6 for precomputed radiance transfer (length of the transfer vector is 36). The model is illuminated by an approximation of a point light source. Figure 1(a) shows the original model containing 200 vertices and 246 triangles. This is not sufficient to sample the transfer operator adequately and results in severe interpolation artifacts in the shadow. Figure 1(b) shows uniform subdivision of the original model (1565 vertices and 2584 triangles). This leads to a sufficient sampling of the transfer

operator farther from the column, but there are still visible artifacts on the shadow boundary near the column. Figure 1(c) shows the results of our adaptive subdivision (1188 vertices, 2828 triangles). The plane is subdivided much more near the column where the shadow can be quite sharp. The artifacts near the column are much less visible.

Figure 5 compares uniform and adaptive subdivision in a more complicated configuration. With the uniform subdivision, the transfer operator is undersampled on the floor in proximity of the knot object. Adaptive subdivision refines the mesh more in this area and removes the most severe artifacts from the uniform subdivision.

5 Conclusions and Future Work

We have presented an approach to adaptive subdivision of meshes with precomputed radiance transfer. The objective is to decrease the interpolation error that can occur if the radiance transfer operator is undersampled. The subdivision is performed in preprocess and no dynamic subdivision is therefore required at render-time. The most important part of the subdivision algorithm is the error measure which estimates the worst case interpolation error for a triangle over all possible illuminations. We proposed an empirical measure that was inspired by estimating worst case contrast over triangle edges. For our test scenes the rendering quality with adaptively subdivided models is superior to that of uniformly subdivided models.

In future work we plan to extend the approach to precomputed radiance transfer with general BRDFs. We also want to relate the error measure on curved and flat surfaces in a more principled way.

Acknowledgements

This work has been partially supported by the Ministry of Education, Youth and Sports of the Czech Republic under the research program No. Y04/98: 212300014 (Research in the area of information technologies and communications), Office of Naval Research, and ATI Research.

References

BAUM, D. R., MANN, S., SMITH, K. P., AND WINGET, J. M. 1991. Making radiosity usable: automatic preprocessing and meshing techniques for the generation of accurate radiosity solutions. In *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, ACM Press, 51–60.

CHO, Y., NEUMANN, U., AND WOO, J. 1996. Improved specular highlights with adaptive shading. In *Proceedings of Computer Graphics International (CGI '96)*.

FERWERDA, J. A. 2001. Elements of early vision for computer graphics. *IEEE Computer Graphics and Applications* 21, 5, 22–33.

GIBSON, S., AND HUBBOLD, R. 1997. Perceptually-driven radiosity. *Computer Graphics Forum* 16, 2 (June), 129–140.

KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics workshop on Rendering*, Eurographics Association, 291–296.

KLEIN, R., SCHILLING, A., AND STRASSER, W. 1998. Illumination dependent refinement of multiresolution meshes. In *Proceedings of Computer Graphics International (CGI '98)*, IEEE Computer Society Press, Los Alamitos, CA, 680–687.

LEHTINEN, J., AND KAUTZ, J. 2003. Matrix radiance transfer. In *Proceedings of the 2003 symposium on Interactive 3D graphics*, ACM Press, 59–64.

MARTIN, I., PUEYO, X., AND TOST, D. 1997. An image-space refinement criterion for linear hierarchical radiosity. In *Graphics Interface '97*, Canadian Human-Computer Communications Society, W. A. Davis, M. Mantei, and R. V. Klassen, Eds., 26–36.

MYSZKOWSKI, K., WOJDALA, A., AND WICYNSKI, K. 1994. Non-uniform adaptive meshing for global illumination. *Machine GRAPHICS & VISION* 3, 4, 601–609.

RAMAMOORTHI, R. 2002. *A Signal-Processing Framework for Forward and Inverse Rendering*. PhD thesis, Stanford University.

SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, ACM Press, 527–536.

SLOAN, P.-P., HALL, J., HART, J., AND SNYDER, J. 2003. Clustered principal components for precomputed radiance transfer. *ACM Trans. Graph.* 22, 3, 382–391.

WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, ACM Press, 255–264.

WILLIAMS, N., LUEBKE, D., COHEN, J. D., KELLEY, M., AND SCHUBERT, B. 2003. Perceptually guided simplification of lit, textured meshes. In *Proceedings of the 2003 symposium on Interactive 3D graphics*, ACM Press, 113–121.

XIA, J. C., AND VARSHNEY, A. 1996. Dynamic view-dependent simplification for polygonal models. In *IEEE Visualization '96*, R. Yagel and G. M. Nielson, Eds., 335–344.