

Practical Global Illumination for Hair Rendering

Cem Yuksel
Computer Science
Texas A&M University
cem@cemyuksel.com

Ergun Akleman
Visualization Sciences
Texas A&M University
ergun@viz.tamu.edu

John Keyser
Computer Science
Texas A&M University
keyser@cs.tamu.edu

Abstract

Both hair rendering and global illumination are known to be computationally expensive, and for this reason we see very few examples using global illumination techniques in hair rendering. In this paper, we elaborate on different simplification approaches to allow practical global illumination solutions for high quality hair rendering. We categorize light paths of a full global illumination solution, and analyze their costs and illumination contributions both theoretically and experimentally. We also propose two different implementation techniques using our novel projection based indirect illumination computation approach and state of the art ray tracing for hair. Our results show that by using our simplifications, a global illumination solution for hair is practical.

1. Introduction

Rendering of hair remains a particularly challenging task in both the research and industrial application domains. The high geometric complexity of realistic hair models makes even direct illumination calculations challenging. For realistic hair synthesis, simulating realistic lighting conditions using global illumination techniques is essential. Moreover, the interaction of light with hair exhibits complex behavior due to the highly anisotropic scattering properties of human hair. Thus, even though global illumination techniques are well established and have been heavily studied for decades, there have been very few attempts to apply these techniques to hair. These few attempts have generally modeled only a subset of all light paths, and they have suffered from high computational complexity, rendering them impractical for many applications.

We present a framework for rendering hair with global illumination based on a theoretical analysis of both hair light scattering behavior and the geometric properties of natural hair models. We restrict our framework to the final gathering stage used in many global illumination methods. Thus,

our framework can easily be placed into an existing global illumination technique to optimize gathering computations involving hair. We theoretically analyze and experimentally test the illumination contributions and computational complexities of different light paths in typical scenes. Based on this, we offer simplifying assumptions and fast approximation techniques for costly operations to accelerate computation of hair illumination by orders of magnitude, while maintaining the final image quality. We elaborate on two different implementation approaches using our novel projection based indirect illumination computation and state of the art fast ray tracing for hair.

2. Related Work

There is a large body of prior work in hair modeling, animation, and rendering. A recent survey [9] gives a good overview of hair methods in computer graphics.

One important visual quality of hair is its creation of semi-transparent shadows. A number of methods have been developed to handle this property both for offline [4] and interactive [3, 6] hair rendering. These methods only consider direct illumination and forward scattered direct illumination from light sources.

Considering highly realistic hair rendering, Marschner et al. have studied the scattering properties for human hair [5]. Zinke et al. have introduced a model for the local scattering effects evident in some hair types [11].

Recently, Moon and Marschner [7] developed a model for multiple scattering in hair based on photon mapping [2]. Their model effectively simulates both forward and non-forward scattered light originating from direct illumination. Unfortunately, this method directly uses the photon map for illumination estimation, which blurs the illumination as an artifact. Therefore, it is not clear whether the blurring effect created comes from multiple scattering in hair, or if it is an artifact of the estimation inherent to photon map illumination. Zinke and Weber [12] also directly used the illumination estimation from the photon map for global illumination with hair, rather than having a final gathering computation.

3. Theoretical Framework

We restrict our framework to the final gathering stage of a global illumination solution similar to the gathering in [2]. To find the illumination at any point, we compute the irradiance coming from all directions. This irradiance includes all possible light paths regardless of the importance of their illumination contribution. In this section, we explain our practical global illumination framework by categorizing these light paths and theoretically analyzing the expected contribution of each category.

In our framework, the irradiance from a certain direction may come from the following sources:

- D:** Direct illumination coming from a light source
- I:** Indirect illumination coming from non-hair points
- F:** Forward scattered from hair
- R:** Reflected from hair (all other scattering)

In this form $\mathbf{D}(\mathbf{F}^* \mathbf{R}^*)^*$ stands for all the light paths coming directly from a light source into the the hair volume, which is then scattered by the hair; and $\mathbf{I}^*(\mathbf{F}^* \mathbf{R}^*)^*$ is all the light paths coming from indirect illumination and scattered by hair. Note that **I** alone includes all the light paths that come from a non-hair point, i.e. anything originating from a light source and reflecting from one or more polygonal or hair objects in the scene.

Since our goal is to develop a practical framework, we analyze light paths to identify the ones that have an important illumination contribution. We categorize the possible light paths into the following types:

Type1: $\mathbf{D} \mathbf{F}^*$ – These paths correspond to direct illumination and forward scattered direct illumination. Since the light originates from a light source, Type1 is expected to have the most important illumination contribution. Of all light path categories, this one is the easiest to compute, and any hair shadowing technique for point light sources is applicable [4, 3, 6]. Alternatively, ray tracing can be used by sending shadow rays to the light sources.

Type2: $\mathbf{I} \mathbf{F}^*$ – These correspond to indirect illumination and forward scattered indirect illumination. All Type2 paths come from a non-hair object (indirect illumination) and can only be forward scattered without changing direction through the hair volume. Depending on the characteristics of the scene, the illumination contribution of Type2 could be minor or could be as important as Type1. In either case, Type2 plays an important role for realistic hair illumination especially in scenes rendered using global illumination methods. Photon mapping is a widely used global illumination technique and very suitable for this computation. Please note that we are not using the photon map to directly compute illumination estimation as in [7, 12], instead we are using the final gathering approach in [2] and only compute reflected radiance from polygons using the photon

map. This lets us avoid keeping a huge number of photons in the hair volume, and gives a more accurate result.

Type3: $\mathbf{D} \mathbf{F}^* \mathbf{R} \mathbf{F}^*$ – Any Type1 path that goes through one non-forward scattering turns into a Type3 path. A Type3 path may include many instances of forward scattering both before and after the non-forward scattering. This non-forward scattering drops the intensity of the light along the path at least an order of magnitude [5]. Therefore, even if the path does not include any forward scattering, the final intensity of a Type3 path can only be as high as 10% of the direct illumination. However, at certain places the contribution of Type3 can be even more significant than a Type2 path, depending on the scene. On the other hand, computing Type3 paths is more complex than Type2.

Type4: $\mathbf{D} \mathbf{F}^* \mathbf{R} \mathbf{F}^* \mathbf{R} (\mathbf{F}^* \mathbf{R}^*)^*$ – These are the light paths that correspond to direct illumination non-forward scattered more than once. When the light goes through more than one non-forward scattering, its intensity drops below two orders of magnitude. Thus the contribution of a Type4 path is less no more than 1% of the direct illumination. On the other hand, computing Type4 paths is generally very expensive, even if accounting for only two non-forward scattering events. Therefore, in our framework, we choose to ignore Type4 paths, assuming that their contribution is negligible and their computation is not practical. One way to include these paths is using direct visualization of the photon map as suggested by [7] and [12]. However, even though these methods theoretically include Type4 paths along with Type1 and Type3, these studies do not state what percentage of the photons in the photon map actually go through more than a single non-forward scattering. Moreover, the blurring artifact of direct visualization of the photon map hides any such minor contributions even if they could be visible.

Type5: $\mathbf{I} \mathbf{F}^* \mathbf{R} (\mathbf{F}^* \mathbf{R}^*)^*$ – These correspond to indirect illumination non-forward scattered at least once. Similar to type 4 paths, the computation of type 5 is also quite expensive. Moreover, when the light from indirect illumination undergoes non-forward scattering, its intensity drops by an order of magnitude. In most cases, the intensity of the indirect illumination alone is about an order of magnitude lower than the direct illumination, though this is not always the true (e.g. when there are highly specular objects like mirrors). Therefore, the indirect illumination that undergoes non-forward scattering is expected to have at most two orders of magnitude lower intensity than the source. Moreover, computing these paths requires a path tracer or a second level (or more) Monte Carlo sampling. Therefore, we do not include these paths in our practical global illumination framework.

Note that these five light path types do not have common light paths, and the collection of these types includes all possible light paths coming to a point.

4. Implementation

We propose two different implementations for our framework. The first implementation is based on ray tracing and the second is our novel projection based indirect illumination method for hair. Both methods have their strengths and difficulties that we will discuss in this section. We also introduce simplification techniques to reduce the computation time without reducing the final image quality.

4.1. Ray Tracing Hair

In our ray tracing implementation for hair we used the techniques described in [8]. Here the accelerating data structure for space partitioning is the most important component to achieve high performance. Most of the computation time is spent traversing this data structure, which helps to eliminate unnecessary intersection computations between rays and hair strands. In our experiments, we found that parameters used for correct space partitioning in the accelerating data structure can bring more than two orders of magnitude speed up in ray tracing. We recommend that the interested reader refer to space partitioning heuristics described in [1].

In our framework, we used ray tracing to compute Type2 and Type3 paths. We used Monte Carlo techniques to gather the illumination from a precomputed photon map as in [2]. While computing Type3 paths, we calculate the reflected radiance for each intersection of the ray with a hair strand and multiply it by the visibility along the ray until the intersection point.

4.2. Projection Based Indirect Illumination

Similar to ray tracing, the visibility function can be computed using a projection based technique. While computing the visibility at a point on a non-hair surface, we project all the hair strand line segments to the hemisphere centered on the point and aligned with the surface normal. Then, we sample the projection to find the occlusion for each direction. Since hair can also be illuminated from behind, to compute the visibility function at a point on a hair strand we use two hemispheres back to back to account for all directions over the full sphere.

For sampling on the hemisphere surface we subdivide the hemisphere into faces of approximately equal area. The subdivision method that we use results in an igloo-like structure (see [10] for details).

In ray tracing, the accelerating data structure removes the computation's linear dependence on the number of hair strand line segments. However, the projection based method as described has linear dependence. Here, we introduce two

approaches to speed up computation based on level of detail methods.

Hair Model Simplification: Realistic hair models inherit strong similarity between neighboring hair strands. Using this property, we can represent the hair volume only by a subset of the hair strands and lower resolution hair strand models to reduce the number of line segments.

Distance-based Level of Detail: The other simplification is to use different levels of detail based on the distance of a line segment from the point where the illumination is computed. We use a user defined distance parameter s and divide the space into a number of layers, such that a hair strand is in layer i if the distance of its beginning position to the calculation point is within the range $[(i-1)s, is)$. For each layer, we use a different level of simplification. For the i^{th} layer, we achieve this by projecting only every i^{th} hair strand line segment.

5. Experimental Analysis

We have run several tests to identify the contributions of each light path type. We have also compared the quality and the computation times of different methods. In this section, we present the results of our experiments.

Figure 1 shows the realism introduced by each light path type. The hair model consists of over one million line segments. In Figure 1a the hair model is illuminated by only Type1 paths and it appears to be too dark, because it does not have indirect illumination. Figure 1b shows the results of Type1 and Type2 paths. As can be seen from these images, Type2 paths bring a balance to the illumination and the final result is significantly more realistic as compared to only Type1 illumination. In Figure 1c, we also include Type3 paths along with Type1 and Type2. Even though the effect of Type3 paths are not as significant as Type2, they introduce considerable improvement to the realism of the final image.

All these qualitative improvements of Type2 and Type3 paths come with computational costs. These costs can be minimized by choosing the correct computational approach for each case. Table 1 shows the computation times of the scene in Figure 1 for different techniques, computing different illumination types. As can be seen from the table, a full ray tracing solution to compute Type1, Type2, and Type3 paths performs very slowly. On the other hand, using deep shadow maps for Type1 and replacing a part of ray tracing with deep shadow map lookups while computing Type3 paths significantly improves the rendering time. In this scene, totally ignoring the illumination contribution of Type3 paths improves the rendering time for ray tracing, while replacing ray tracing with our projection-based igloo method brings additional speed up.

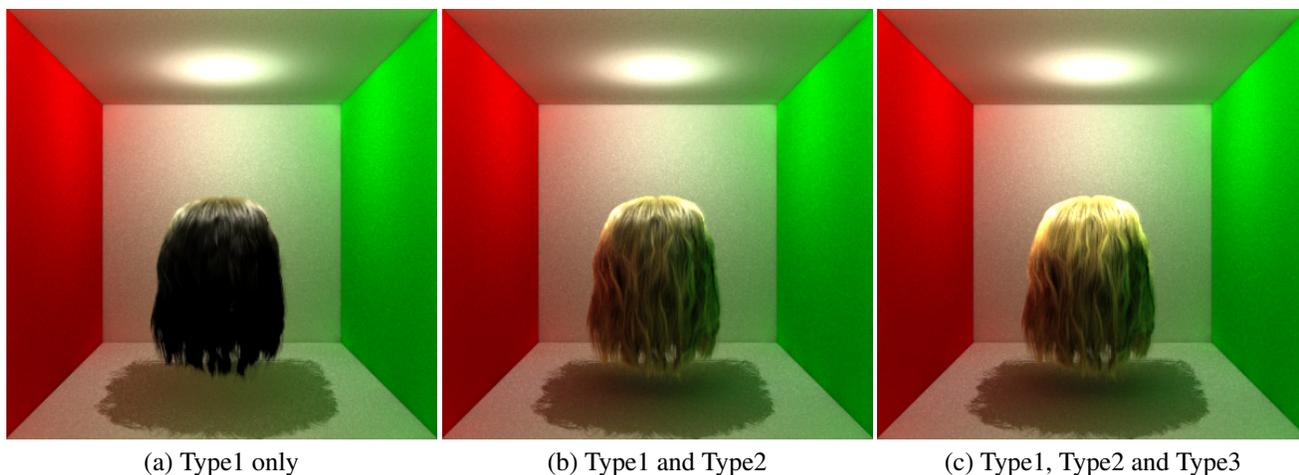


Figure 1. The effect of Type1, Type2, and Type3 paths on the final image quality

Table 1. Performance of different methods

Type1	Type2	Type3	time
RT	RT	RT	8:13:51
DSM	RT	RT+DSM	35:45
RT	RT	–	26:26
DSM	RT	–	25:45
RT	P	–	11:13
DSM	P	–	12:45

RT: ray tracing DSM: deep shadow maps P: projection

6. Conclusion

As a result of our analysis, we see that indirect illumination from Type2 paths plays an important role to achieve realism in hair rendering. Our projection-based indirect illumination method can be used to compute Type2 paths, and with our simplification techniques it generates a visually identical result to ray tracing in significantly less time. Type3 paths improve the image further with some additional cost. This additional cost is minimal when deep shadow maps are used to compute a part of the Type3 paths; however, our results show that a full ray tracing solution for all paths is too expensive.

References

- [1] V. Havran. *Heuristic Ray Shooting Algorithms*. Ph.d. thesis, Department of Computer Science and Engineering, Faculty of Electrical Engineering, Czech Technical University in Prague, November 2000.
- [2] H. W. Jensen. Global Illumination Using Photon Maps. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 21–30. Springer-Verlag, 1996.
- [3] T.-Y. Kim and U. Neumann. Opacity shadow maps. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 177–182, London, UK, 2001. Springer-Verlag.
- [4] T. Lokovic and E. Veach. Deep shadow maps. In *Proceedings of SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 385–392. ACM, ACM Press / ACM SIGGRAPH, 2000.
- [5] S. Marschner, H. W. Jensen, M. Cammarano, S. Worley, and P. Hanrahan. Light scattering from human hair fibers. In *Proceedings of SIGGRAPH 2003*, Computer Graphics Proceedings, Annual Conference Series, pages 780–791. ACM, ACM Press / ACM SIGGRAPH, 2003.
- [6] T. Mertens, J. Kautz, P. Bekaert, and F. van Reeth. A self-shadow algorithm for dynamic hair using clustered densities. In *Proceedings of Eurographics Symposium on Rendering 2004*, pages 173–178, June 2004.
- [7] J. T. Moon and S. R. Marschner. Simulating multiple scattering in hair using a photon mapping approach. In *Proceedings of SIGGRAPH 2006*, pages 1067–1074. ACM, ACM Press / ACM SIGGRAPH, 2006.
- [8] I. Wald. *Realtime Ray Tracing and Interactive Global Illumination*. PhD thesis, Computer Graphics Group, Saarland University, 2004.
- [9] K. Ward, F. Bertails, T.-Y. Kim, S. R. Marschner, M.-P. Cani, and M. Lin. A survey on hair modeling: Styling, simulation, and rendering. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 13(2):213–34, Mar-Apr 2007.
- [10] C. Yuksel and E. Akleman. Hair occlusion: Rendering hair-like objects with global illumination. Technical report, 2006. <http://www-viz.tamu.edu/faculty/ergun/research/visualization/papers/tr06.pdf>.
- [11] A. Zinke, G. Sobottka, and A. Weber. Photo-realistic rendering of blond hair. In *Vision, Modeling, and Visualization 2004*, pages 191–198, 2004.
- [12] A. Zinke and A. Weber. Global illumination for fiber based geometries. In *Ibero-American Symposium on Computer Graphics - SIACG (2006)*, 2006.