8 Recent Advancements in Disney's Hyperion Renderer

Brent Burley, David Adler, Matt Jen-Yuan Chiang, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, Daniel Teece[2]

Figure 7: A production frame from *Moana*, rendered using Disney's Hyperion Renderer. (Copyright © Disney Animation 2017.)

8.1 Introduction

Path tracing at Walt Disney Animation Studios began with the Hyperion renderer, first used in production on *Big Hero 6*. Hyperion is a custom, modern path tracer using a unique architecture designed to efficiently handle complexity, while also providing artistic controllability and efficiency.

The concept of physically based shading at Disney Animation predates the Hyperion renderer. Our history with physically based shading significantly influenced the development of Hyperion, and since then, the development of Hyperion has in turn influenced our philosophy towards physically based shading.

8.1.1 History of Physically Based Rendering at Disney Animation

A major theme in the past decade of rendering at Disney Animation has been the advantages of physically based solutions over biased approximations, for both visual richness and artistic controllability. Early successes with physically inspired hair shading on *Tangled* led to the development of the Disney BRDF by Burley [2012] during *Wreck-It Ralph*, and was subsequently extended into the modern Disney BSDF with subsurface scattering and refraction during *Big Hero 6* (as described by Burley [2015]). Adopting physically meaningful parameters made shader response more predictable and intuitive for artists.

At the same time, moving from REYES-style rasterization rendering to physically based path tracing has removed the considerable data management overhead imposed on artists to manage the shadow maps, point clouds, and more that rasterization rendering necessitated. We continue to strive for ease of use by simplicity and consistency—"it just works".

8.1.2 Inception of Disney's Hyperion Renderer

The Hyperion renderer was developed at Walt Disney Animation Studios with the aim of providing global illumination within a physically based framework while retaining the benefits of highly coherent shading

[2] Author names after Brent Burley are presented in alphabetical order by last name.

Figure 8: A production frame from *Zootopia*, rendered using our fully path-traced fur model. (Copyright © Disney Animation 2017.)

shown in previous production-proven rasterization-based strategies. Beginning in 2011 and continuing through 2012, an initial period of research and exploration was followed by a prototype and proof-of-concept stage where successful ideas were tested at a production scale. The results were compelling enough to drive development into a full production renderer over a short time period from 2013 to 2014, coinciding with the production of *Big Hero 6*. Hyperion has subsequently been used to render the feature films *Zootopia* and *Moana*, and is being used to render all projects currently in production at the studio.

The core of Hyperion's architecture is *sorted deferred shading*. Starting with primary rays we perform ray sorting, binning rays by direction and grouping them into large, sorted ray batches of fixed size. Next, we perform scene traversal, one sorted ray batch at a time. We use a two-level quad BVH with streaming packet cone traversal in the top level, and single-ray traversal in the bottom. We exploit the fact that our ray batches are directionally coherent to perform approximate front-to-back traversal at each node. The result of traversal is a list of hit points, one per ray. Next, hit point sorting organizes ray hits with the aim of maximizing coherent access from the texture cache. If a shading task has many hit points, it is partitioned into sub-tasks, further increasing parallelism. The shader also feeds secondary rays back into ray sorting to continue ray paths. Increasing the batch size provides improved coherence and better performance for traversal and shading. A more detailed description of this architecture is presented by Eisenacher et al. [2013]

The introduction of Hyperion has produced numerous benefits across the studio. Due to the ease of use, more departments can render full frames, and we now render shots continuously in all stages of production with full global illumination. Higher quality rendering in all stages of the production process provides a much earlier view into the look of each show. Because of the predictability of the results, artists are able to final frames in fewer iterations than before path tracing. As the complexity of our films continues to increase, the Hyperion renderer grows and evolves to meet the unique challenges of each show.

8.2 Transitioning from Multiple Scattering Approximations to Brute-force Solutions

Historically, phenomena requiring large amounts of multiple scattering were prohibitively computationally expensive to evaluate using physically correct brute-force solutions, so approximations were typically used instead. Many of our projects since *Big Hero 6* have required complex multiple-scattering effects, such as the fur in *Zootopia*, snow in *Frozen Fever*, various cases in *Moana*, and volumes in upcoming shows. Moving to path tracing has allowed us to discard previous approximate solutions and move to-

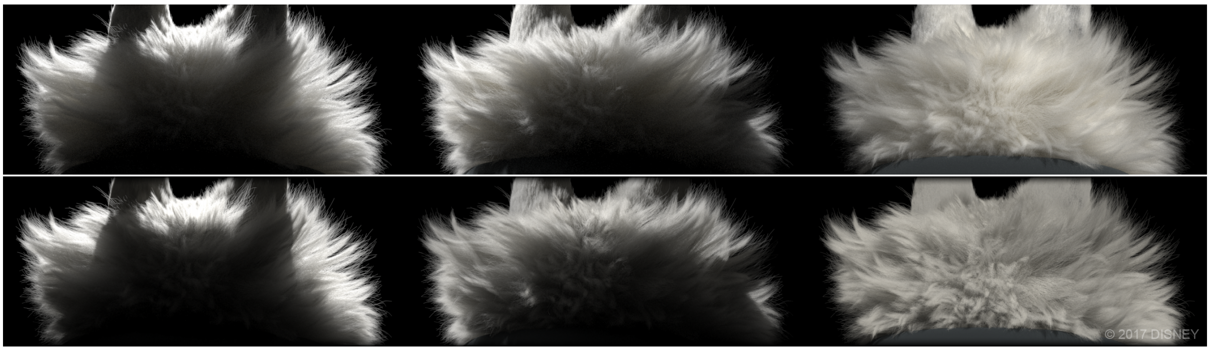

Figure 9: Fur rendered with path-traced multiple fiber scattering (top) vs. its Dual Scattering approximation, both using the same lighting setup and same absorption coefficients in the fur strands. For more details, we refer the reader to Chiang et al. [2016a]. (Copyright © Disney Animation 2017.)

wards brute-force solutions for these effects. In this section, we describe some of these phenomena, and also discuss some of the careful parameterization work that is often necessary to make these models more artist-friendly.

8.2.1 Path-traced Hair and Fur

Prior to *Zootopia*, we used an artistically controlled Dual Scattering hair model originally developed for *Tangled* by Sadeghi et al. [2010]. While this model was more physically inspired than previous ad hoc models, we found that the Dual Scattering model lacked the richness that multiple scattering already provided in other subsystems of the Hyperion renderer. The lack of multiple scattering in fur and hair often contributed to a coarse and stiff look that became amplified with the presence of high albedo fibers (Figure 9). In order to address this problem, we came up with a physically based single fiber scattering model that allows for efficient Monte Carlo rendering of path-traced multiple fiber scattering in production (Chiang et al. [2016a]).

One common sampling strategy for previous physically based fiber scattering models (such as the one proposed by d'Eon et al. [2011]) is to focus on eliminating the shading variation across the width of a fiber. However, often in production path-traced global illumination, a more prominent source of sampling variance per fiber is the illumination coming from the complex surrounding scene, as well as illumination coming from all nearby fibers. This outside illumination requires a large number of shader evaluations to fully converge. We realized that by relying on the general Monte Carlo framework to integrate over fiber width, we can greatly reduce the complexity of the per-sample evaluation. We also introduced a fourth lobe that re-injects the energy lost from only representing R, TT and TRT interactions to achieve perfect energy conservation, even for non-absorbing fibers. These improvements made brute force path tracing of fur and hair possible in production, and significantly contributed to the look of *Zootopia* (Figure 8).

One issue with a physically based model is that its parameters, such as the absorption coefficient, are often not intuitive for the artists. Also, physically based parameters can have very little visual connection with the final material appearance, which comes from the result of multiple scattering. To address artistic controllability, we re-parameterized the fiber roughness to be perceptually uniform. We also allow the artists to specify multiple scattering albedo directly, which is used internally to derive the absorption coefficient for rendering. These enable efficient artist workflows while remaining physically consistent, empowering the artists to achieve wider ranges of appearances of hair and fur with great efficiency.

We continue to use the techniques described in this section for all productions beyond *Zootopia*, to great success. For example, in *Moana*, human characters are covered in fine, groomed peach fuzz, which provides effects such as rim lighting through brute-force multiple scattering of back lighting instead of requiring a dedicated light type (Figure 10).

Figure 10: Rim lighting effects on Moana and Grandma Tala, simulated through brute-force multiple scattering of back-lighting through fine hairs. (Copyright © Disney Animation 2017.)

8.2.2 Path-traced Subsurface Scattering for Snow and Skin

Subsurface scattering is the phenomenon of light scattering inside an object and exiting at a different place than it entered. This phenomenon produces effects like softness, light bleeding, and shadow saturation. Preventing translucent materials like skin and snow from looking too opaque or hard is essential.

For years, the diffusion approximation was the method of choice. During *Big Hero 6*, we used *normalized diffusion*, introduced by Burley [2015], as our primary subsurface scattering solution. To simplify the problem, most diffusion approximations assume that the medium is a semi-infinite slab of homogeneous material. Diffusion works well even when the geometric assumption is violated, but only if the distance that the light scatters is small compared to the size of the geometric details on the surface. In geometrically small and thin regions, diffusion causes energy loss, while on convoluted surfaces, diffusion causes energy gain. Furthermore, interesting optical effects and important visual cues caused by light scattering through objects of different sizes and shapes are lost when using diffusion.

The physical process of subsurface scattering through arbitrary geometry can be simulated much more accurately using volumetric path tracing. However, since light can potentially scatter hundreds or thousands of times inside an object before exiting, this approach has the potential be extremely computationally expensive. Furthermore, highly directional scattering and small bright light sources increase noise.

We experimented with several approaches to improve the appearance of snow in *Frozen Fever*, and ended up implementing and using a limited brute-force volumetric-path-tracing solution (Figure 11). A number of design decisions were made to circumvent performance problems:

- We performed simulated scattering inside only a user-defined subset of the scene surfaces.
- We assumed that the volumes were completely homogeneous.
- We performed free-flight sampling according to a monochromatic scattering coefficient and calculated the amount of chromatic absorption based on the full path length.

Figure 11: A production frame from *Frozen Fever* with characters made from path-traced subsurface-scattering snow. (Copyright © Disney Animation 2017.)

- We only supported isotropic phase functions and only index-matched diffusely-transmitting interfaces.
- We introduced a hack to increase scatter distances after several hundred bounces.

Although this system worked well, it was very difficult to achieve a desired overall color and desired scatter distances for each color channel. After a good deal of research (which resulted in Koerner et al. [2016]), we came up with an artist-friendly parameterization for setting chromatic scattering and absorption coefficients and a sampling strategy that efficiently handled these chromatic coefficients (Chiang et al. [2016b]). The parameterization is based on fitting curves to sets of simulated results. The sampling strategy uses the path throughput and single-scattering albedo to bias free-flight distributions. We also introduced internal reflection to make results more accurate and reduce unrealistic brightening of edges.

We first used our current path-traced subsurface scattering solution on selected prop elements in *Moana*, although skin for characters continued to use normalized diffusion. When comparing normalized diffusion and path-traced subsurface scattering on *Moana* characters, we discovered that details such as creases and wrinkles had been modeled deeper than expected to compensate for detail loss that occurs from diffusion, which produced different looks with path-traced subsurface scattering. All of our current productions have fully switched over to path-traced subsurface scattering for everything from snow to skin, which has simplified modeling and shading workflows because compensations for diffusion artifacts no longer need to be modeled into geometry. We show an example of a test character rendered with path-traced subsurface scattering skin in Figure 12.

8.2.3 Volume Rendering

Hyperion's sorted deferred architecture provides a significant challenge for implementing volumetric rendering. During *Big Hero 6*, a volume-rendering system was developed that was heavily designed around the sorted deferred concept. Up until recently, one fundamental requirement of the sorted deferred architecture was that a ray must hit a surface before a new ray could be fired. In our current volume-rendering system, Hyperion traces a ray completely through a heterogeneous volume, calculating a transmittance estimate using residual-ratio tracking (Novák et al. [2014]) which is followed by constructing PDFs to sample in-scattering and emission. In-scattering rays are generated to be treated like any other ray and are added to the list of rays for processing in the sorted deferred shading queue. Multiple scattering is achieved by recursing on the procedure.

Figure 12: Skin rendered using path-traced subsurface scattering. For similar render times as our old normalized diffusion technique, our new path-traced subsurface scattering provides richer visual quality and better predictability. (Copyright © Disney Animation 2017.)

Figure 13: A production frame from *Moana* demonstrating large, complex, dense smoke plumes dominated by low-order scattering. Our current residual-ratio tracking based volume rendering system is not as efficient with high-order multiple scattering, but handles low-order scattering very efficiently. (Copyright © Disney Animation 2017.)

Figure 14: A cloudscape with high-albedo clouds and thousands of multiple-scattering bounces, rendered using our new brute-force volume-rendering system. Our new volume-rendering system makes use of our spectral and decomposition tracking techniques, introduced by Kutz et al. [2017]. (Copyright © Disney Animation 2017.)

This solution produces high quality estimates at high compute costs per sample and is optimized for low-order scattering, such as in smoke plumes and dust clouds with low albedos (Figure 13). This approach avoided major modifications to the core sorted deferred architecture. However, building a high quality PDF per ray makes high-order multiple scattering with potentially tens of thousands of bounces unfeasibly expensive. Initially, efforts were made to rely on only Hyperion's existing volume-rendering solution, and fill in missing energy from high-order scattering using various approximations and cheats. These efforts failed; artists found these approximate techniques difficult to control and were not able to achieve the highly realistic look they were targeting.

Starting in 2016, a project arose within the studio that required rendering enormous quantities of high-albedo clouds with very high-order multiple scattering. We are in the process of transitioning into a new volume system that allows us to render thousands of scattering events per path. To make such long paths more practical to render, we derived new, advanced versions of tracking (Kutz et al. [2017]). Architecturally, this new volume renderer is made possible by changes within the sorted deferred system that remove the requirement for a ray to always end at a surface, meaning that the volume renderer can immediately re-scatter a ray upon finding a scattering event. The brute-force and therefore predictable nature of our new volume-rendering system has allowed artists to hit their target looks with significantly greater ease than before, while also providing significantly faster iteration and feedback loops. In Figure 14, we demonstrate an example of a scene with many clouds that Hyperion's new volume rendering system handles with ease.

8.3 Future Directions

Over the course of three feature films and several more short films, we have gained significant experience with both general path tracing, and the practicalities of our sorted deferred architecture. Using lessons learned from production, we continue to evolve Hyperion's architecture going forward.

8.3.1 Unbounded Path Lengths

Hyperion has allowed us to universally adopt multi-bounce global illumination studio-wide with significant efficiency, up to a certain number of bounces. However, Hyperion's architecture imposes some interesting restrictions that shaped our light sampling strategy, along with our ability to support truly unbounded path lengths efficiently.

Hyperion originally only supported a single ray type, which makes direct light sampling via next-event-estimation difficult to implement. Instead of using next event estimation, we relied on explicit light sampling, splitting rays at each scattering event so that we could shoot separate samples towards lights and along BSDFs. This splitting approach alone results in a geometric increase in the number of rays at each bounce, which places considerable memory pressure on a system that already keeps tens of millions of rays in flight at once. To keep the total number of rays in flight manageable, we rely on an aggressive Russian Roulette strategy to cull rays.

While this approach has generally worked well for us, it has some consequences at higher bounces. As we reach higher and higher bounces, our aggressive Russian Roulette begins to dominate over splitting, resulting in large drops in ray counts. We eventually reach a point where there are too few rays in flight simultaneously to justify the overhead of our sorted deferred ray batches, meaning that paths with extremely high lengths can become inefficient to compute.

One current active area of development is loosening our requirement for a single ray type, allowing us to replace our existing sampling strategy with a more conventional next-event-estimation approach. Changing the sampling strategy changes the relationship between samples-per-pixel and variance, which presents an interesting user-education topic. To make unbounded paths more efficient, we are also examining ways to loosen the definition of a ray batch, along with different methods for scheduling ray batches. We plan to have more to report on this topic in the near future.

8.3.2 Leveraging Increasing Core Counts

As the complexity of our films continues to grow, we anticipate our studio's rendering needs to outstrip the increases in computational power predicted by Moore's Law in the near future. As a result, we are concerned not just about scaling Hyperion to increasingly more cores on a single render node, but also about distributing and scaling both memory and compute for a single render job beyond the bounds of a single render node. This topic continues to be an active area of research that we hope to be able to report more on in the near future.

8.4 Conclusion

We presented a series of recent advancements made in Disney's Hyperion Renderer, with a particular focus towards replacing multiple scattering approximations with true, brute-force path-traced solutions. Adopting a path-tracing renderer studio-wide has allowed us to pursue effects and features that were previously considered completely infeasible to solve using brute-force methods in production, all while also providing simpler, more intuitive controls for artists. We continue to investigate ways to evolve our architecture and increase the scalability and efficiency of the Hyperion renderer even further, in support of even more future advancements to production quality and artist controllability.

References

Brent Burley. 2012. Physically Based Shading at Disney. *Practical Physically-Based Shading in Film and Game Production, SIGGRAPH 2012 Course Notes* (July 2012).

Brent Burley. 2015. Extending Disney's Physically Based BRDF with Integrated Subsurface Scattering. *Physically Based Shading in Theory and Practice, SIGGRAPH 2015 Course Notes* (July 2015).

Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. 2016a. A Practical and Controllable Hair and Fur Model for Production Path Tracing. *Computer Graphics Forum (Proc. of Eurographics)* 35, 2 (May 2016), 275–283.

Matt Jen-Yuan Chiang, Peter Kutz, and Brent Burley. 2016b. Practical and Controllable Subsurface Scattering for Production Path Tracing. In *SIGGRAPH 2016 Talks*. 49:1–49:2.

Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An Energy-Conserving Hair Reflectance Model. *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 30, 4 (June 2011), 1181–1187.

Christian Eisenacher, Gregory Nichols, Andrew Selle, and Brent Burley. 2013. Sorted Deferred Shading for Production Path Tracing. *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 32, 4 (July 2013), 125–132.

David Koerner, Jan Novák, Peter Kutz, Ralf Habel, and Wojciech Jarosz. 2016. Subdivision Next-Event Estimation for Path-Traced Subsurface Scattering. In *Eurographics Symposium on Rendering 2016: Experimental Ideas and Implementations*. 91–96.

Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 36, 4 (July 2017), 111:1–111:16.

Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual Ratio Tracking for Estimating Attenuation in Participating Media. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 33, 6 (Nov. 2014), 179:1–179:11.

Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. 2010. An Artist Friendly Hair Shading System. *ACM Transactions on Graphics* 29, 4 (July 2010), 56:1–56:10.