SIGGRAPH 2021

UNBIASED EMISSION AND SCATTERING IMPORTANCE SAMPLING FOR HETEROGENEOUS VOLUMES

WEI-FENG WAYNE HUANG
PETER KUTZ,
YINING KARL LI
MATT JEN-YUAN CHIANG

THE PREMIER **CONFERENCE** & **EXHIBITION** IN
COMPUTER GRAPHICS & INTERACTIVE TECHNIQUES

Walt Disney

ANIMATION STUDIOS

©Disney

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

Hello Siggraph 2021 and thank you for having us.

Please refrain from
screen grabs or recording.

THANK YOU!

Please refrain from recording these talks. This helps us so we can continue to share our work with you.

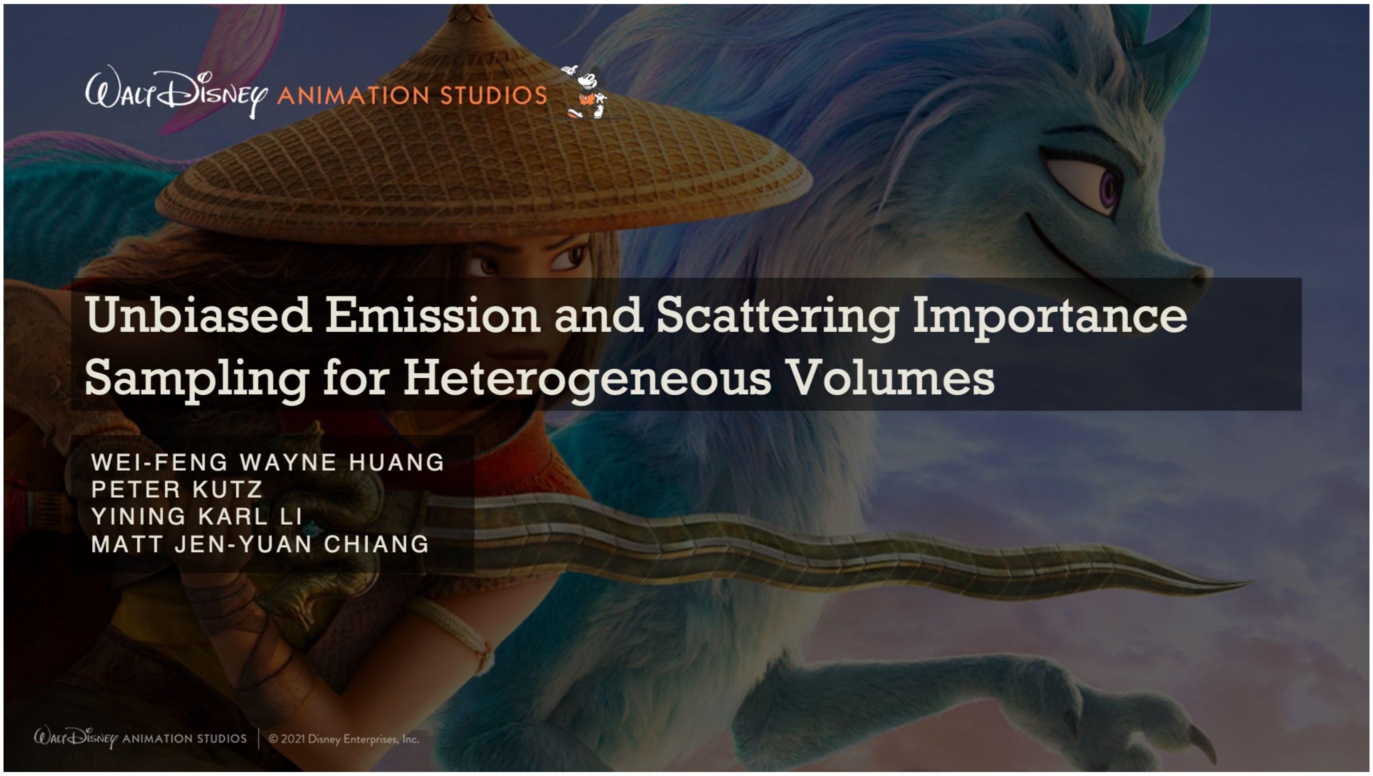

Walt Disney ANIMATION STUDIOS

Unbiased Emission and Scattering Importance Sampling for Heterogeneous Volumes

WEI-FENG WAYNE HUANG
PETER KUTZ
YINING KARL LI
MATT JEN-YUAN CHIANG

Today I would like to discuss some recent development works we did to improve the volumetric emission and scattering sampling for our in house rendering system, Hyperion.

Wei-Feng Wayne Huang

SENIOR SOFTWARE ENGINEER

RENDERING

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

I am Wayne Huang from Walt Disney Animation Studio, and I will present the work for our team

New System, New Challenges

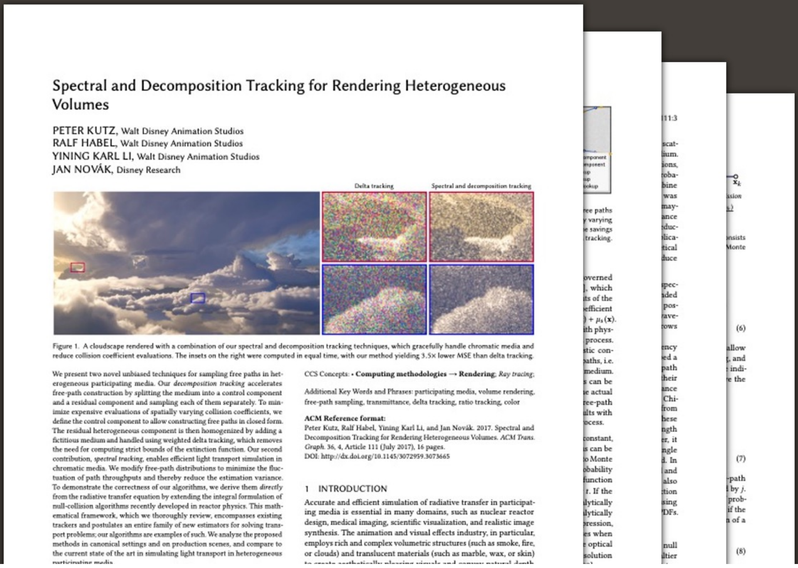

Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes

PETER KUTZ, Walt Disney Animation Studios
RALF HABEL, Walt Disney Animation Studios
YINING KARL LI, Walt Disney Animation Studios
JAN NOVÁK, Disney Research

Figure 1. A cloudscape rendered with a combination of our spectral and decomposition tracking techniques, which gracefully handle chromatic media and reduce collision coefficient evaluations. The insets on the right were computed in equal time, with our method yielding 3.5× lower MSE than delta tracking.

We present two novel unbiased techniques for sampling free paths in heterogeneous participating media. Our *decomposition tracking* accelerates free-path construction by splitting the medium into a control component and a residual component and sampling each of them separately. To minimize expensive evaluations of spatially varying collision coefficients, we define the control component to allow constructing free paths in closed form. The residual heterogeneous component is then homogenized by adding a fictitious medium and handled using weighted delta tracking, which removes the need for computing strict bounds of the extinction function. Our second contribution, *spectral tracking*, enables efficient light transport simulation in chromatic media. We modify free-path distributions to minimize the fluctuation of path throughputs and thereby reduce the estimation variance. To demonstrate the correctness of our algorithms, we derive them *directly* from the radiative transfer equation by extending the integral formulation of null-collision algorithms recently developed in reactor physics. This mathematical framework, which we thoroughly review, encompasses existing trackers and postulates an entire family of new estimators for solving transport problems; our algorithms are examples of such. We analyze the proposed methods in canonical settings and on production scenes, and compare to the current state of the art in simulating light transport in heterogeneous

CCS Concepts: • **Computing methodologies** → **Rendering**; *Ray tracing*;

Additional Key Words and Phrases: participating media, volume rendering, free-path sampling, transmittance, delta tracking, ratio tracking, color

ACM Reference format:
Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Trans. Graph.* 36, 4, Article 111 (July 2017), 16 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073665

1 INTRODUCTION

Accurate and efficient simulation of radiative transfer in participating media is essential in many domains, such as nuclear reactor design, medical imaging, scientific visualization, and realistic image synthesis. The animation and visual effects industry, in particular, employs rich and complex volumetric structures (such as smoke, fire, or clouds) and translucent materials (such as marble, wax, or skin)

At Walt Disney Animation Studios we recently revamped our in house volume rendering system. This new system allows us to render white clouds among other things in an unbiased way and we described our methods in this paper published in SIGGRAPH 2017.

However, when it was deployed, we found that some scenarios are challenging to the new system. So, today I am going to talk about some further improvements that we made along the way to overcome these challenges.

Overview

Introduction
- Volume rendering background
- Null-collision formulation
- Challenges

Emission sampling improvements

Scattering sampling improvements

I would like to first briefly introduce the null-collision formulation of volume rendering, which is what our system based on. And then I'll dive deeper into our improvements on emission and scattering sampling.

Volume Rendering Background

Let’s first talk about some background for volume rendering

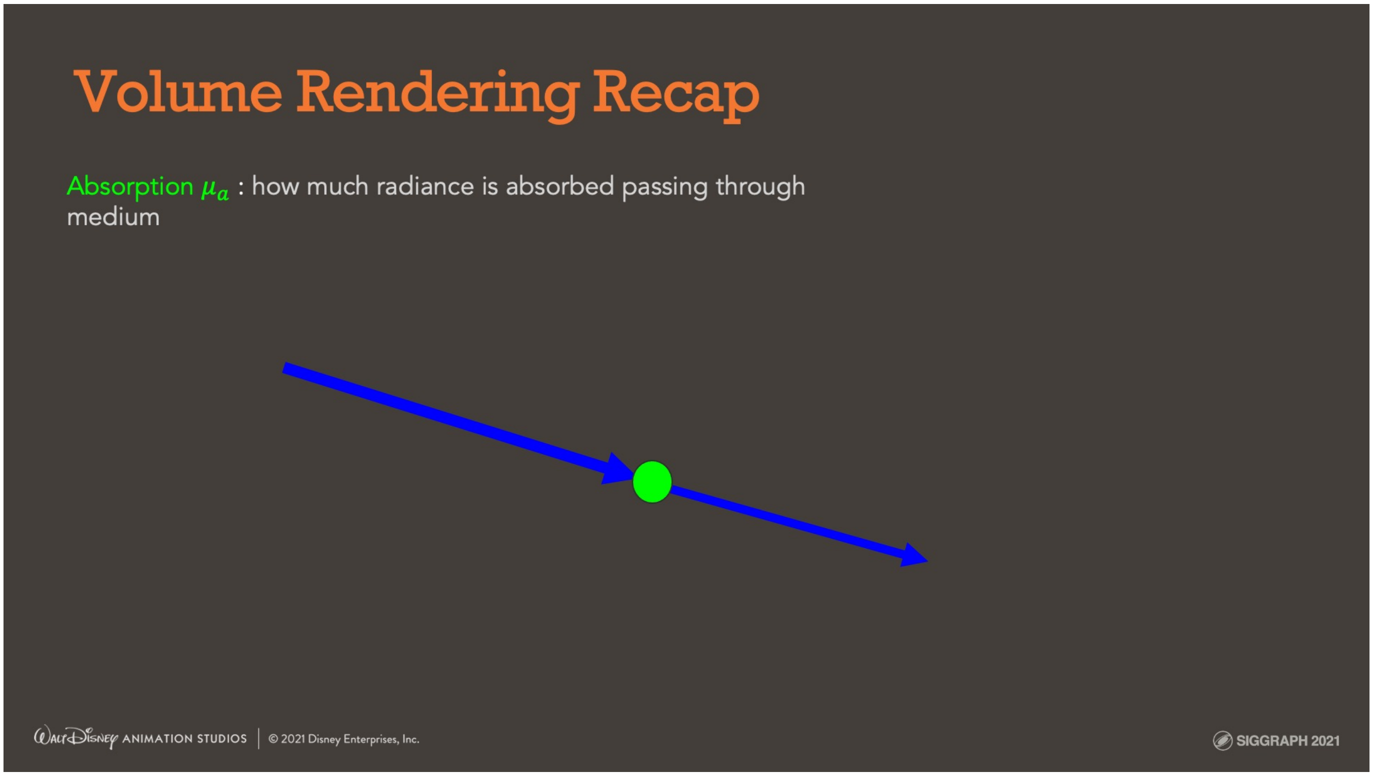

Volume Rendering Recap

Absorption μ_a : how much radiance is absorbed passing through medium

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

Volume rendering is composed of the following medium interactions: Absorption - how much energy is absorbed passing through the medium.

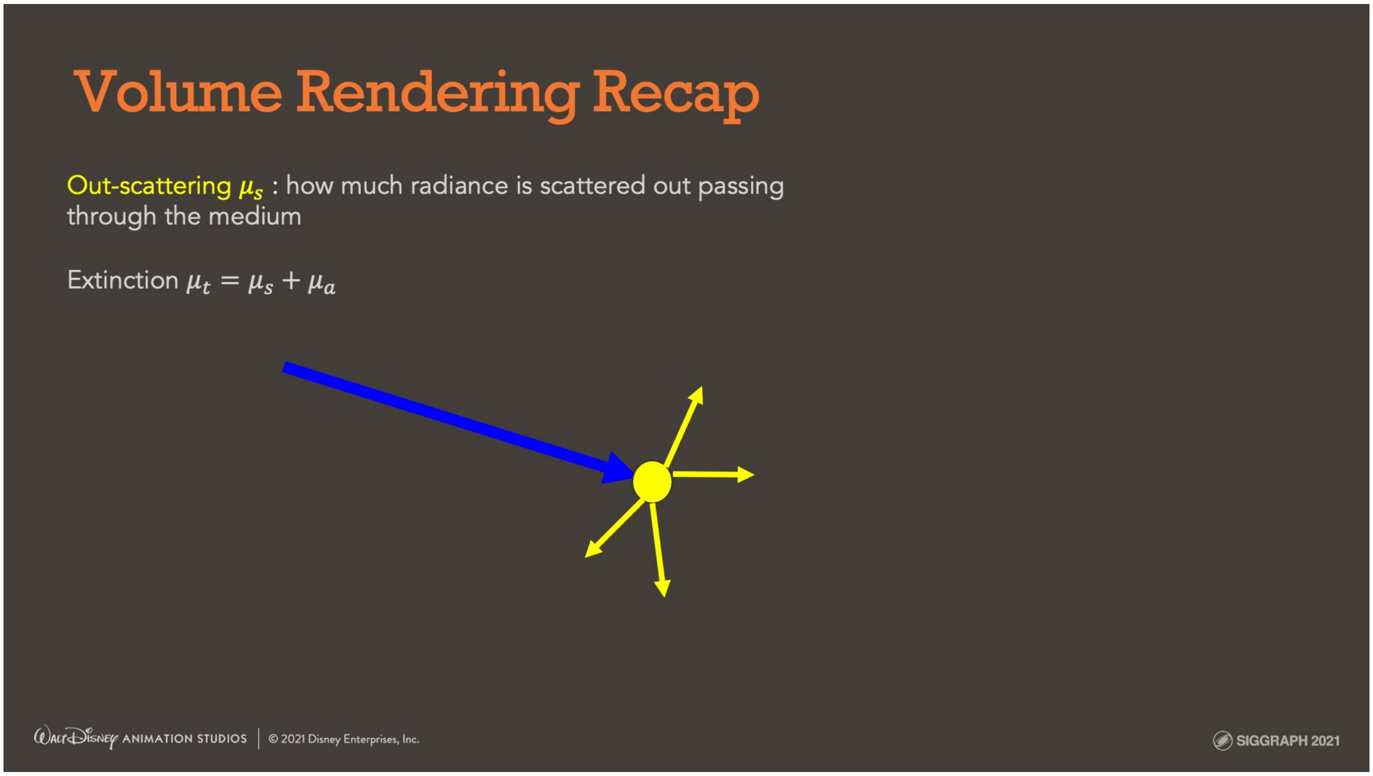

Volume Rendering Recap

Out-scattering μ_s : how much radiance is scattered out passing through the medium

Extinction $\mu_t = \mu_s + \mu_a$

Out-scattering, how much energy being scattered to other directions by the medium.

The combination of the absorption and out-scattering is commonly referred to as extinction, which accounts for the net energy loss when a ray traveling in a volume.

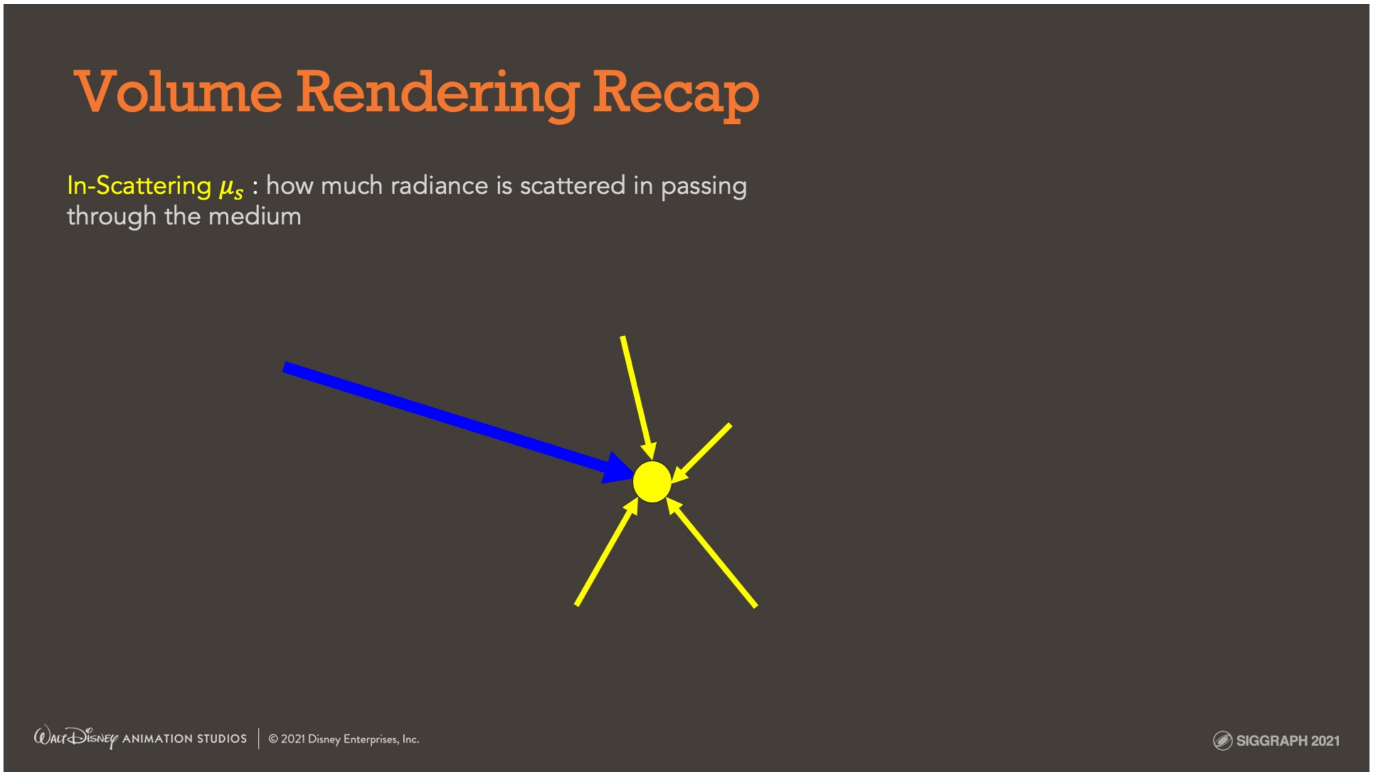

Volume Rendering Recap

In-Scattering μ_s : how much radiance is scattered in passing through the medium

In-scattering, how much energy being scattered into the current directions by the medium.

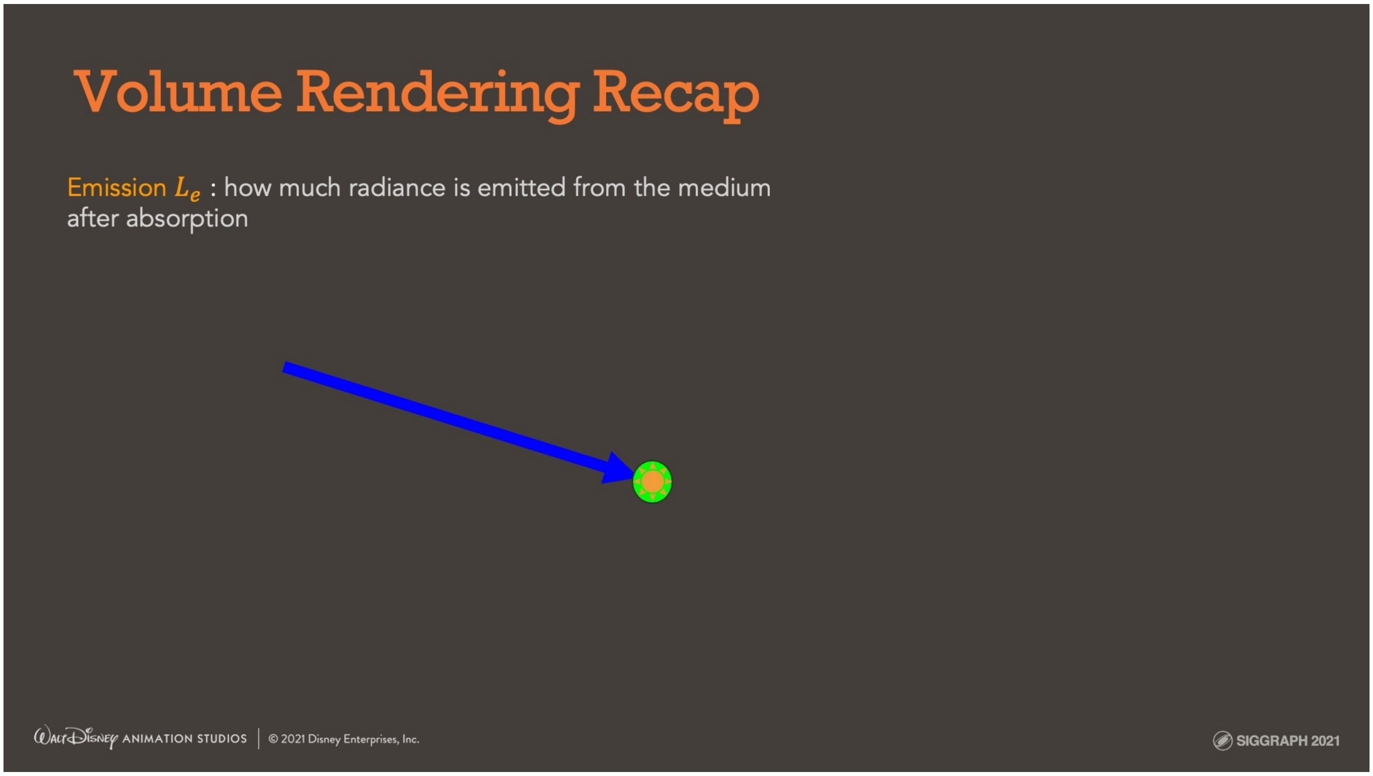

Volume Rendering Recap

Emission L_e : how much radiance is emitted from the medium after absorption

Also there is emission, describing the amount of energy emitted from a medium after absorption.

In-scattering and emission together accounts for the net energy gain when traveling through a volume.

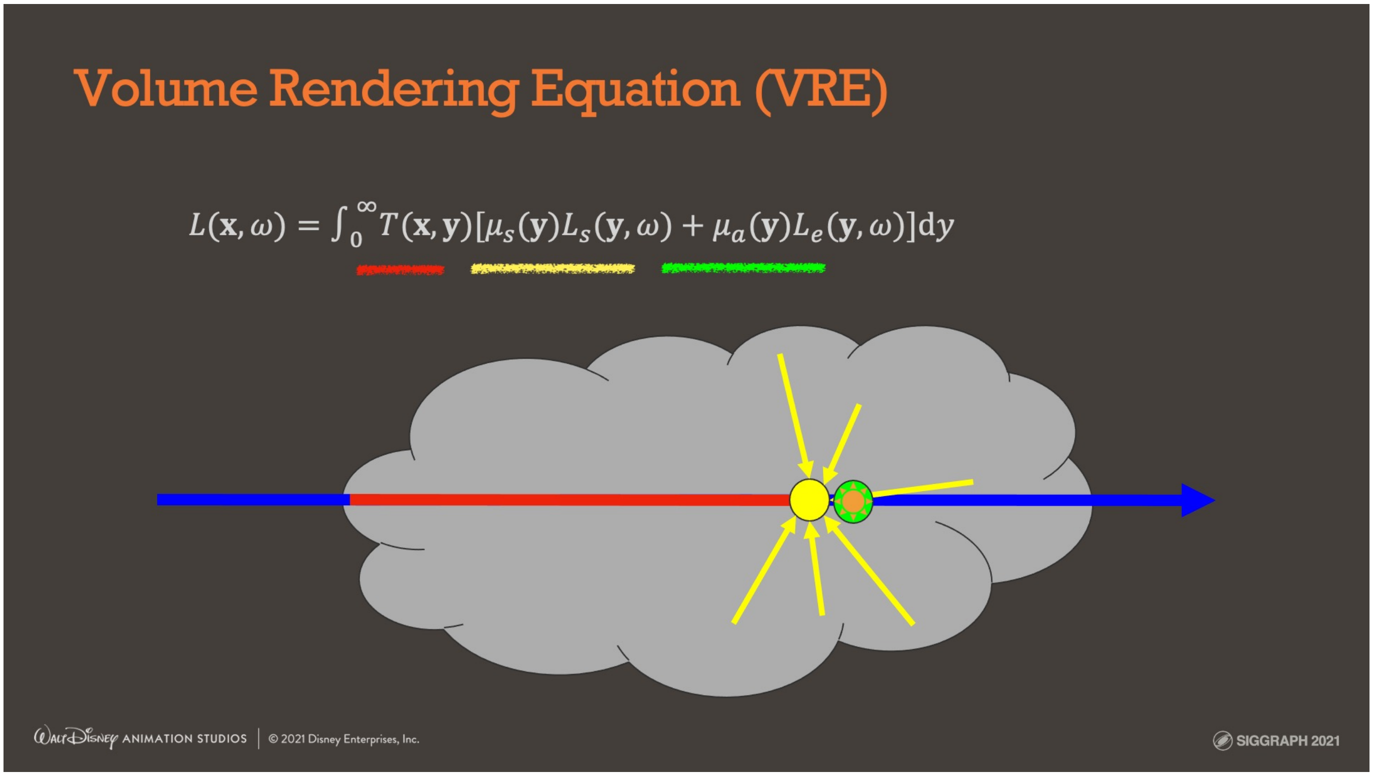

Volume Rendering Equation (VRE)

$$L(\mathbf{x},\omega) = \int_0^{\infty} T(\mathbf{x},\mathbf{y})[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega)]\mathrm{d}y$$

Putting them together we got the volume rendering equation:

To solve for the equation, for each points along the ray segment, we need to gather the emission term and the source term inside the volume and attenuate it with the transmittance term.

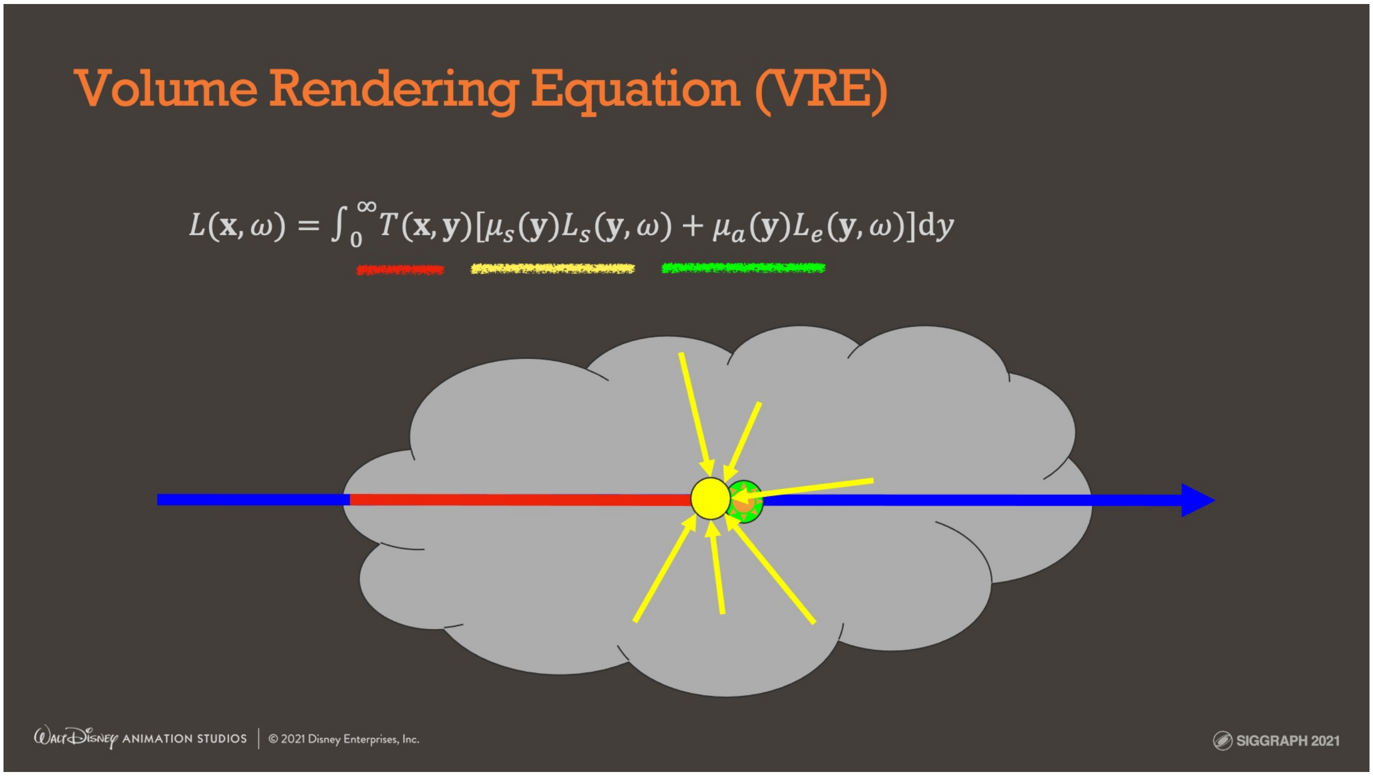

Volume Rendering Equation (VRE)

$$L(\mathbf{x},\omega) = \int_0^\infty T(\mathbf{x},\mathbf{y})[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega)]\mathrm{d}y$$

We then integrate it along the entire segment.

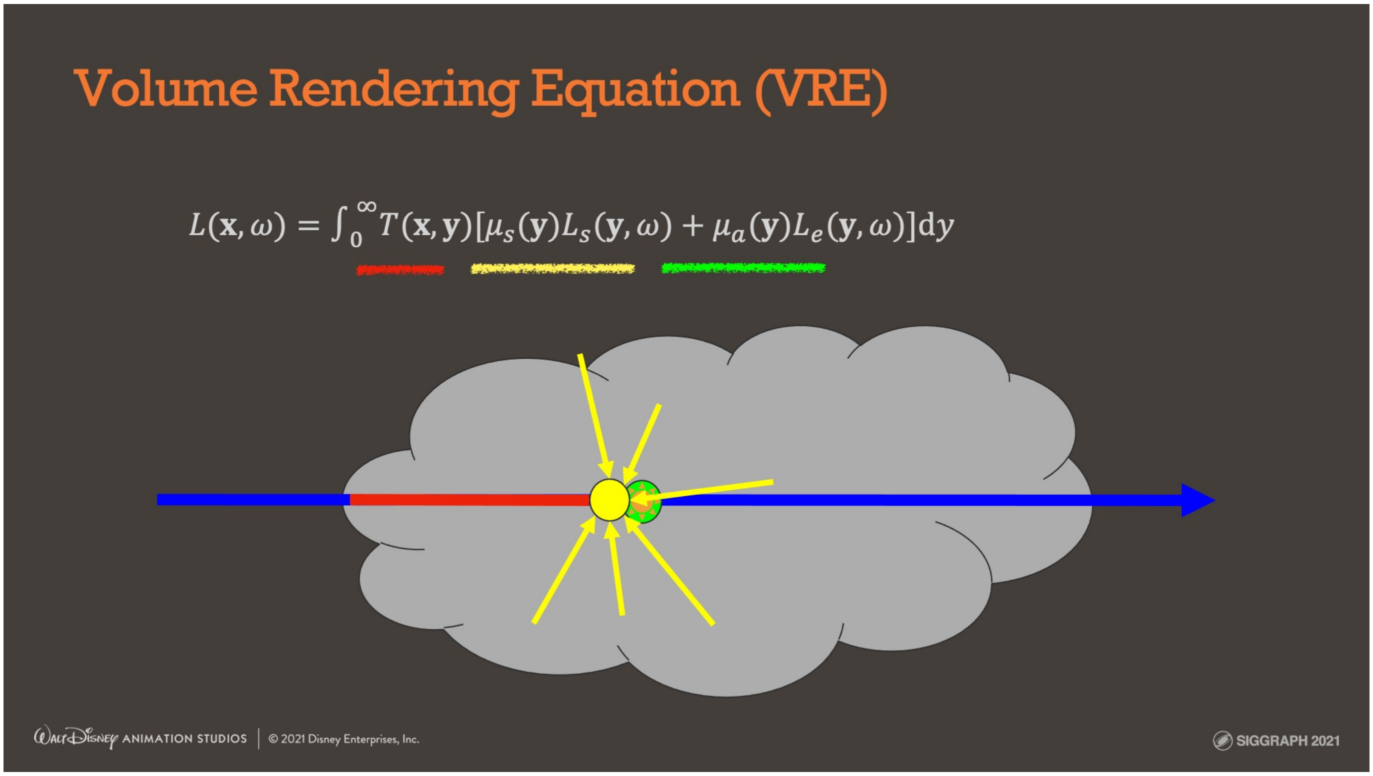

Volume Rendering Equation (VRE)

$$L(\mathbf{x},\omega) = \int_0^\infty T(\mathbf{x},\mathbf{y})[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega)]\mathrm{d}y$$

We then integrate it along the entire segment.

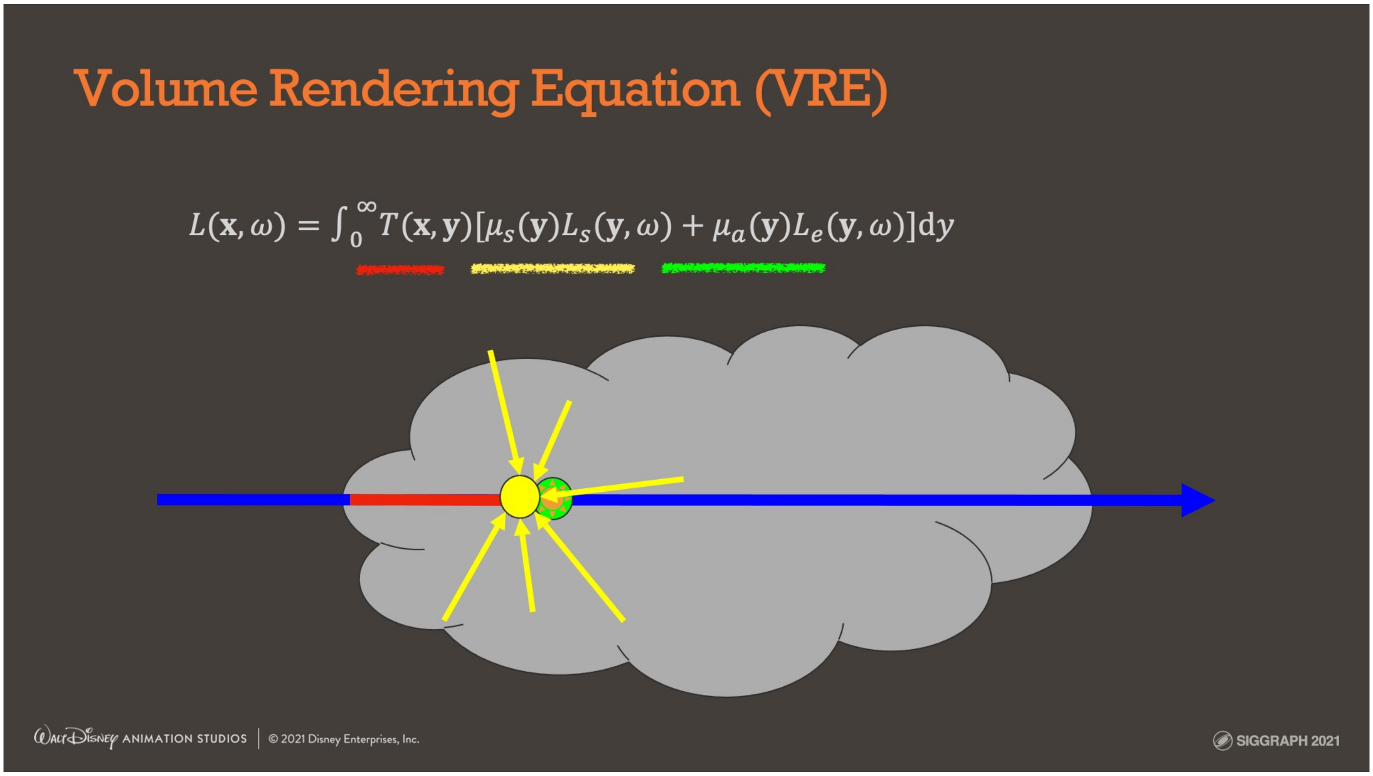

Volume Rendering Equation (VRE)

$$L(\mathbf{x},\omega) = \int_0^\infty T(\mathbf{x},\mathbf{y})[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega)]\mathrm{d}y$$

We then integrate it along the entire segment.

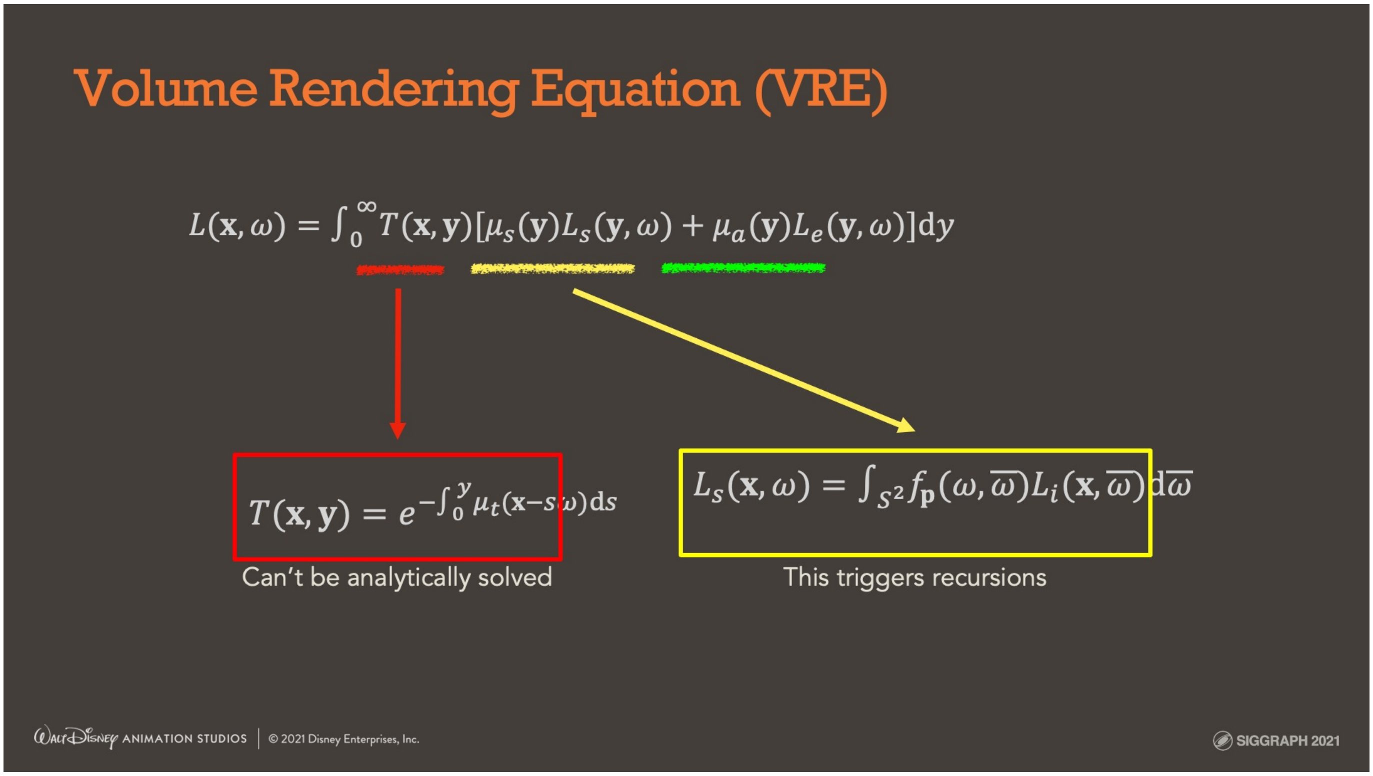

Volume Rendering Equation (VRE)

$$L(\mathbf{x},\omega) = \int_0^\infty \underline{T(\mathbf{x},\mathbf{y})}[\underline{\mu_s(\mathbf{y})L_s(\mathbf{y},\omega)} + \underline{\mu_a(\mathbf{y})L_e(\mathbf{y},\omega)}]\mathrm{d}y$$

$$\boxed{T(\mathbf{x},\mathbf{y}) = e^{-\int_0^y \mu_t(\mathbf{x}-s\omega)\mathrm{d}s}}$$

Can't be analytically solved

$$\boxed{L_s(\mathbf{x},\omega) = \int_{S^2} f_{\mathbf{p}}(\omega,\overline{\omega})L_i(\mathbf{x},\overline{\omega})\mathrm{d}\overline{\omega}}$$

This triggers recursions

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

The reason the volume rendering equation is difficult to solve is because

Evaluating the source terms requires recursively launching new rays, so a highly scattering volume often takes thousands of bounces for a path to leave the volume.

On top of that evaluating the transmittance term requires another integration which often can't be analytically solved.

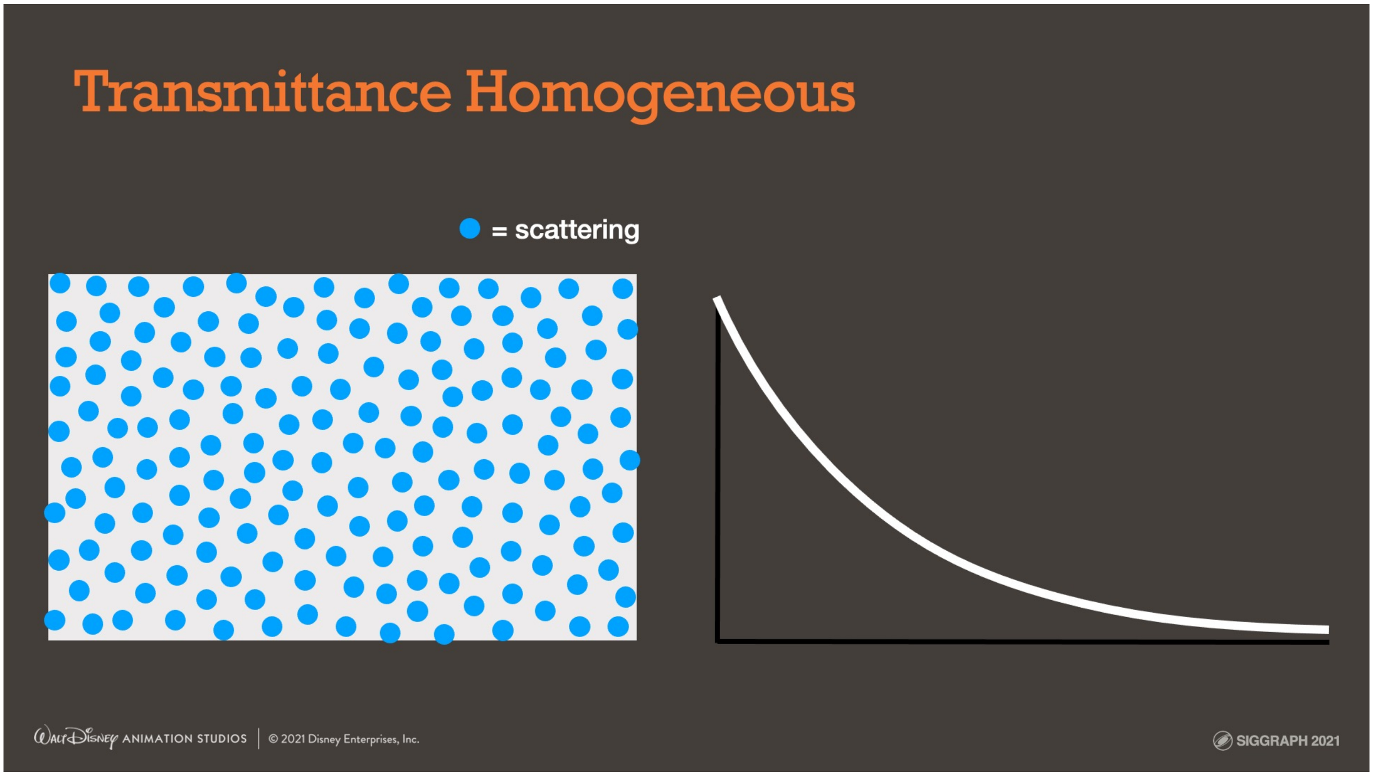

Transmittance Homogeneous

● = scattering

Homogeneous volume is an exception. It has constant extinction coefficient throughout, so that the transmittance integral can be analytically computed.

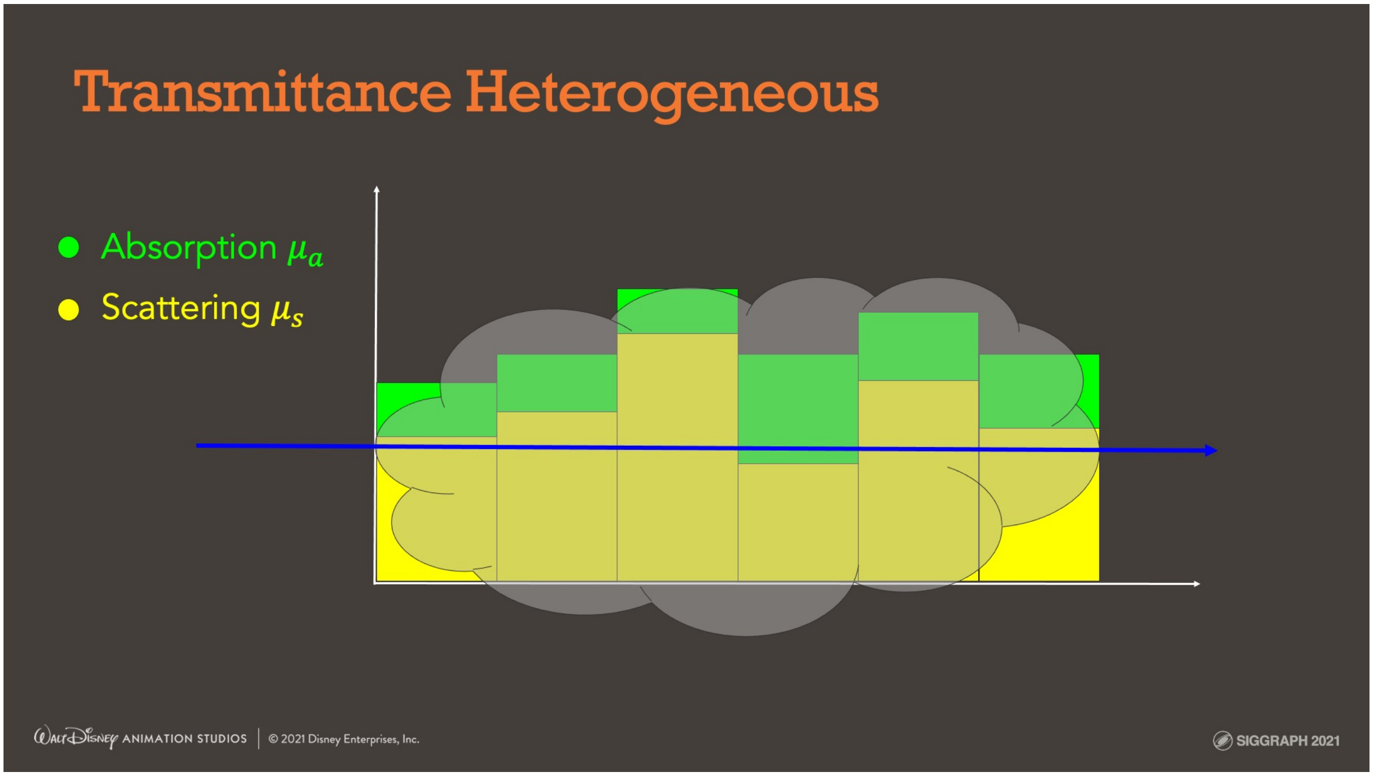

Transmittance Heterogeneous

- Absorption μ_a
- Scattering μ_s

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

However, most of the interesting volumes in film production are heterogeneous.

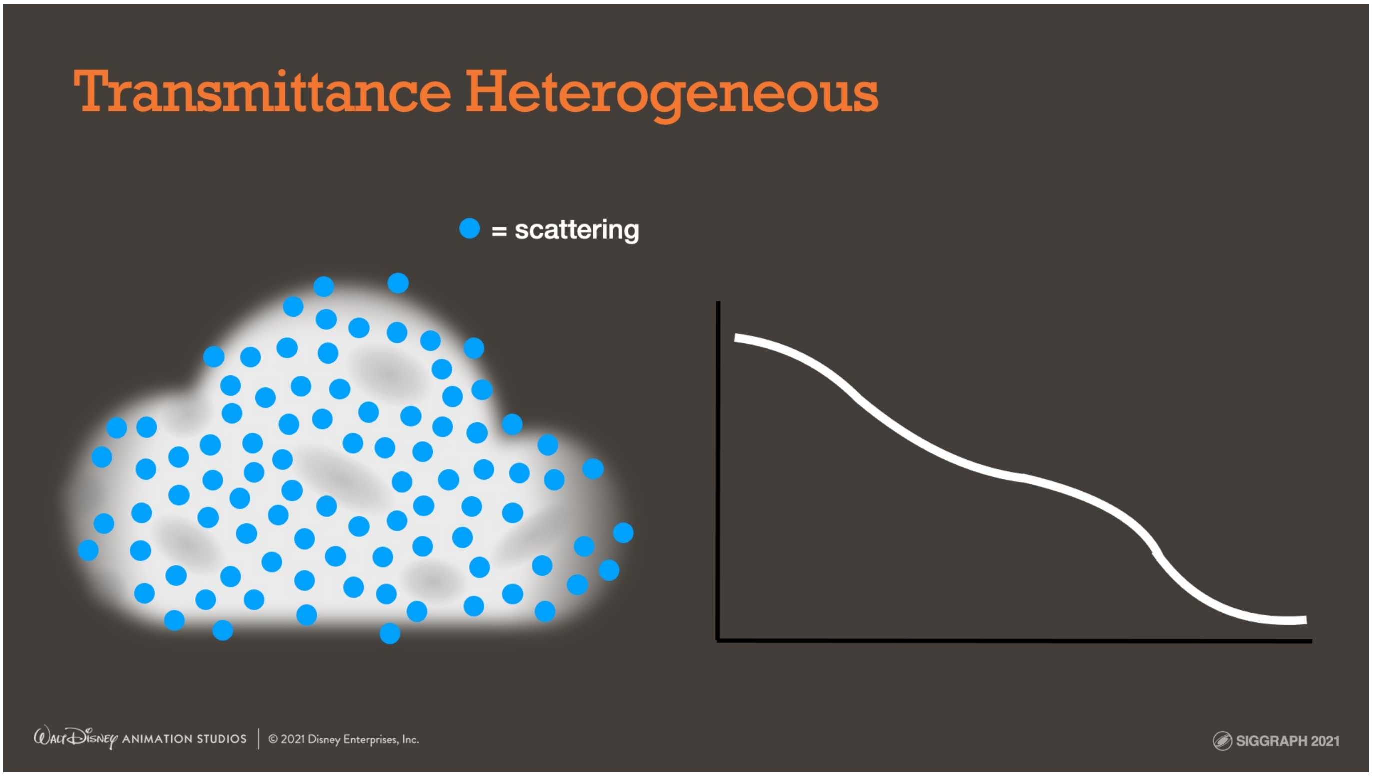

Transmittance Heterogeneous

● = scattering

And the transmittance in this case relies on numerical approximations or Monte Carlo estimations.

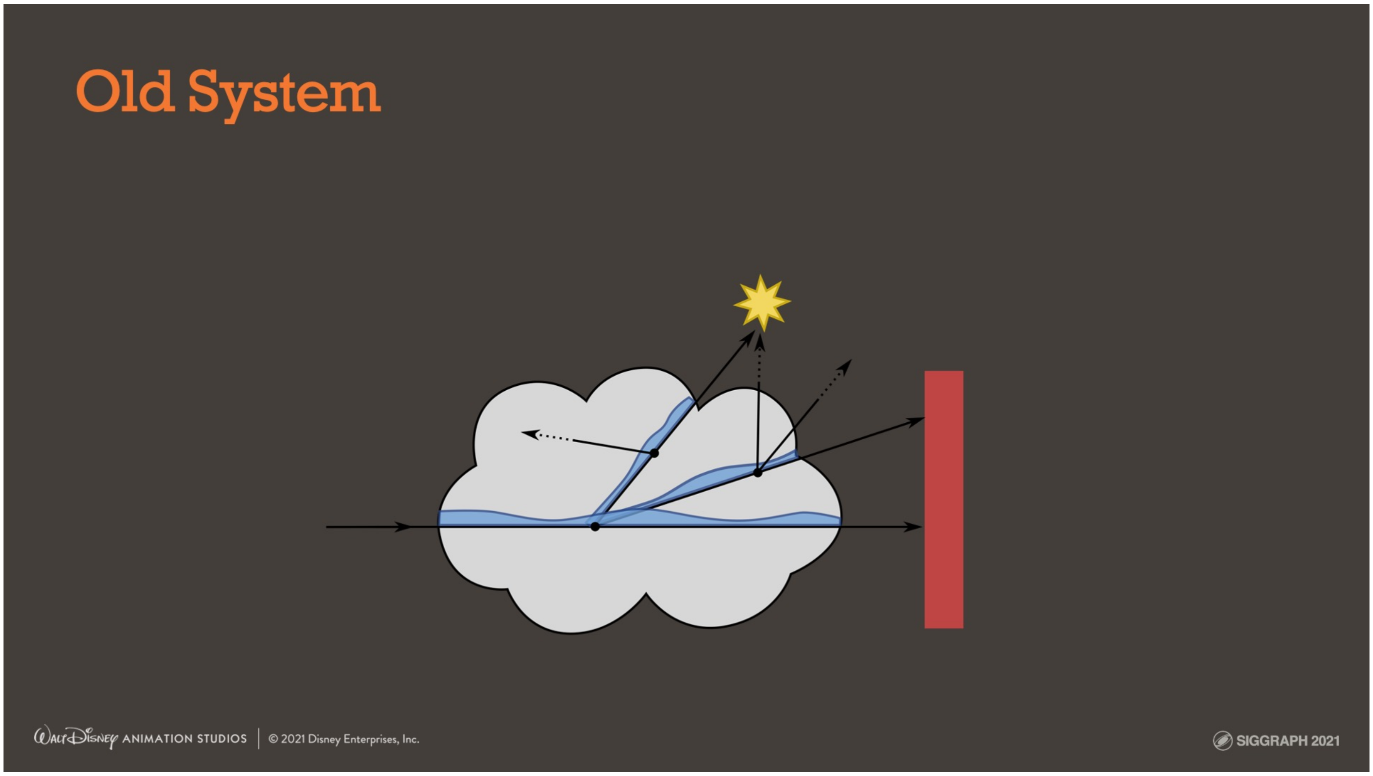

Old System

One popular technique in production volume rendering solves the equation by ray marching through volume along each ray segment to numerically approximate the transmittance. A detailed importance table is built for each ray segment to sample a location for the next emission or scattering event.

This technique works well for low-scattering volumes but its high computational cost for each sample prevents us from rendering thick volumes like a cloud.

Null-Collision Formulation

This limitation motivated us to adopt a Null-Collision based approach.

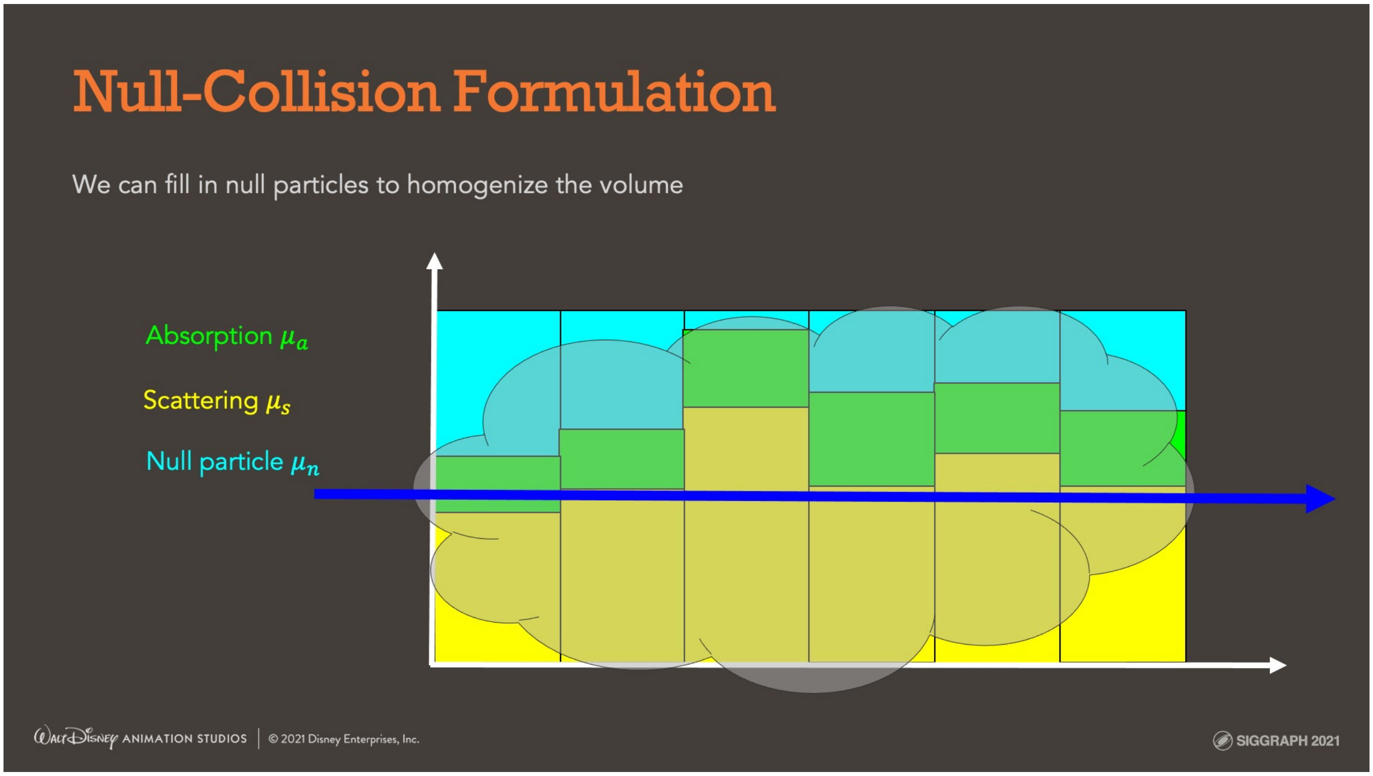

Null-Collision Formulation

We can fill in null particles to homogenize the volume

Absorption μ_a

Scattering μ_s

Null particle μ_n

The basic idea of null collision theory or a so called tracking based system is to homogenous the volume by filling it with “null” particles that don’t invoke any volume interaction.

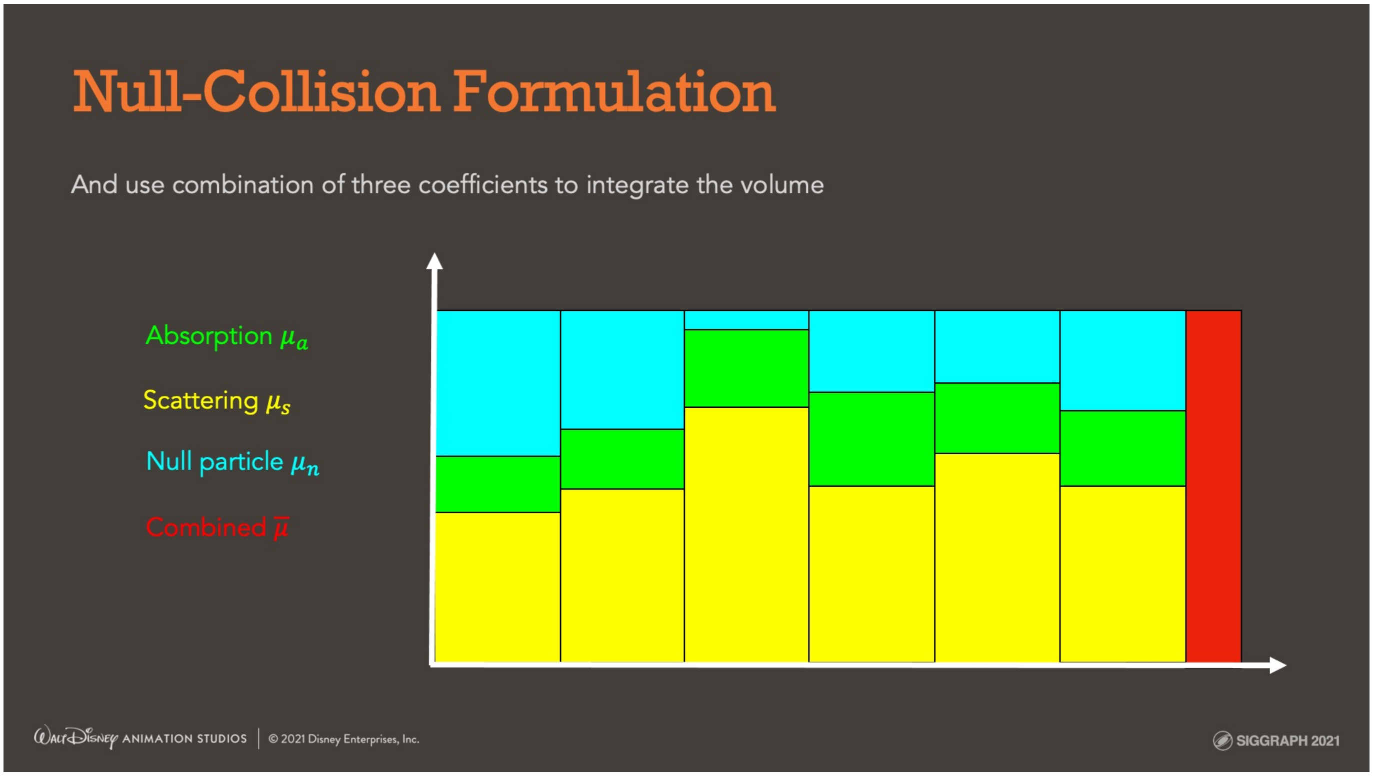

Null-Collision Formulation

And use combination of three coefficients to integrate the volume

Absorption μ_a

Scattering μ_s

Null particle μ_n

Combined $\overline{\mu}$

And in this “homogenized” volume, the combined extinction coefficient is then constant, so that we can sample transmittance analytically.

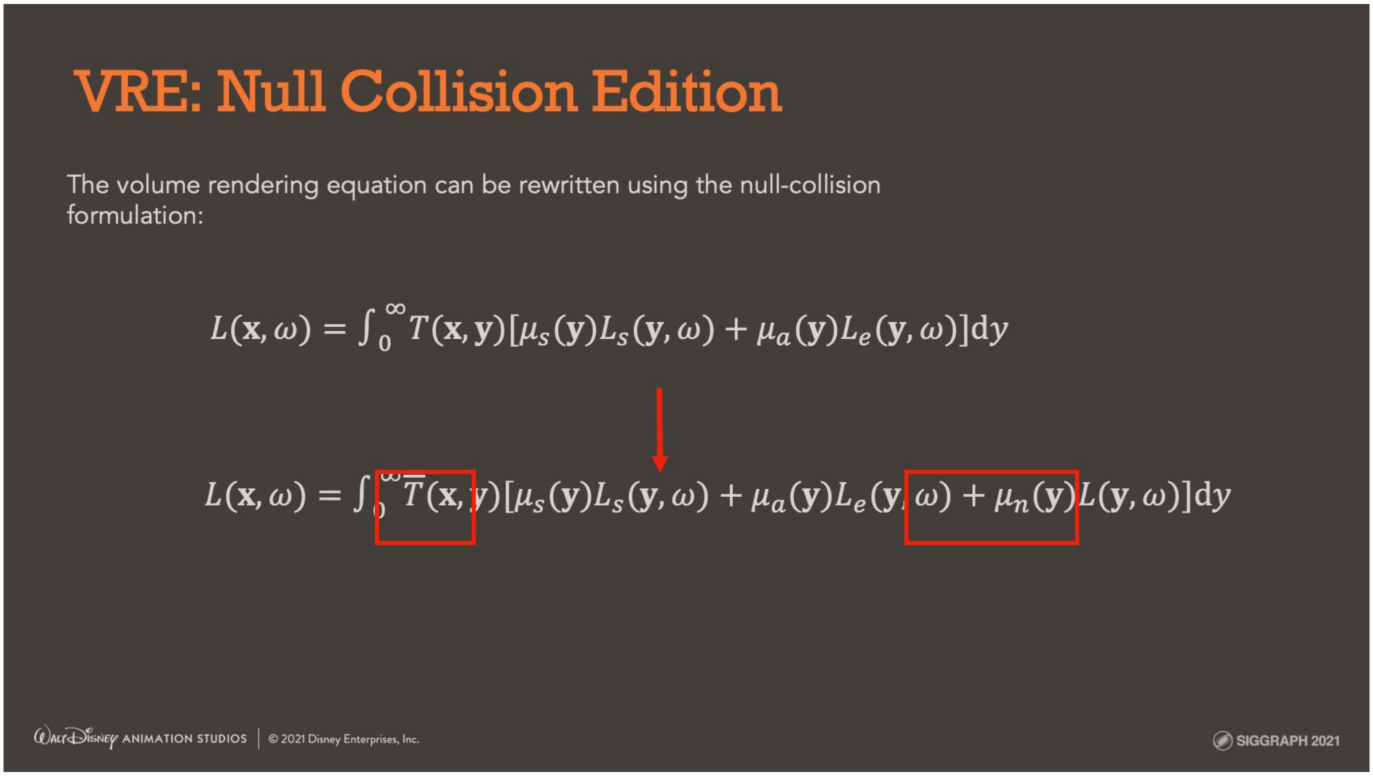

VRE: Null Collision Edition

The volume rendering equation can be rewritten using the null-collision formulation:

$$L(\mathbf{x},\omega) = \int_0^\infty T(\mathbf{x},\mathbf{y})[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega)]\mathrm{d}y$$

$$\downarrow$$

$$L(\mathbf{x},\omega) = \int_0^\infty \boxed{\bar{T}(\mathbf{x},\mathbf{y})}[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega) \boxed{+ \mu_n(\mathbf{y})} L(\mathbf{y},\omega)]\mathrm{d}y$$

We can rewrite the volume rendering equation taking all these into account. Pay attention to the updated and now analytic transmittance term and the additional null-collision term.

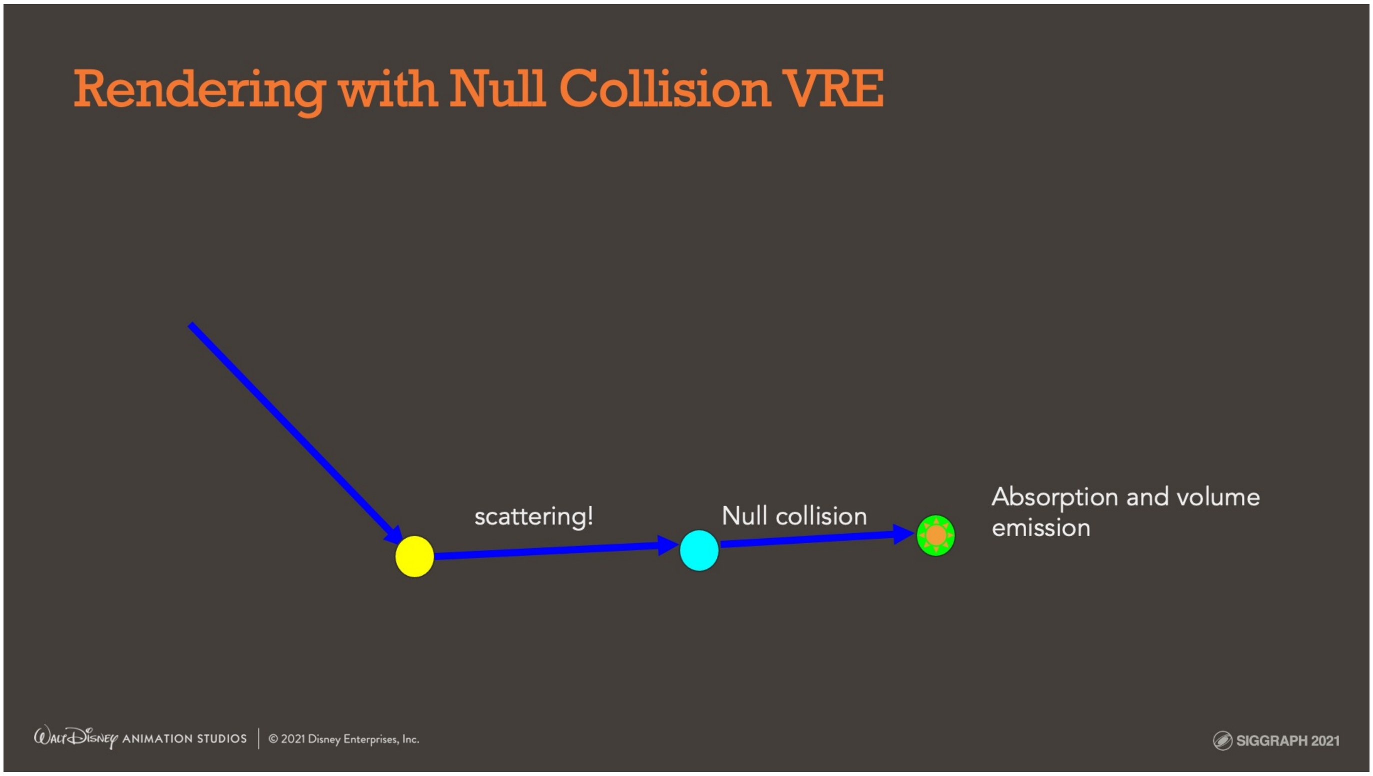

Rendering with Null Collision VRE

scattering!

Null collision

Absorption and volume emission

Walt Disney ANIMATION STUDIOS

SIGGRAPH 2021

Evaluating the Null-Collision Formulation of Volume Rendering Equation becomes a random walk process.
We first start with sampling a free path distance using the combined extinction coefficient, and then probabilistically pick an event out of the absorption, scattering and null-collision events.

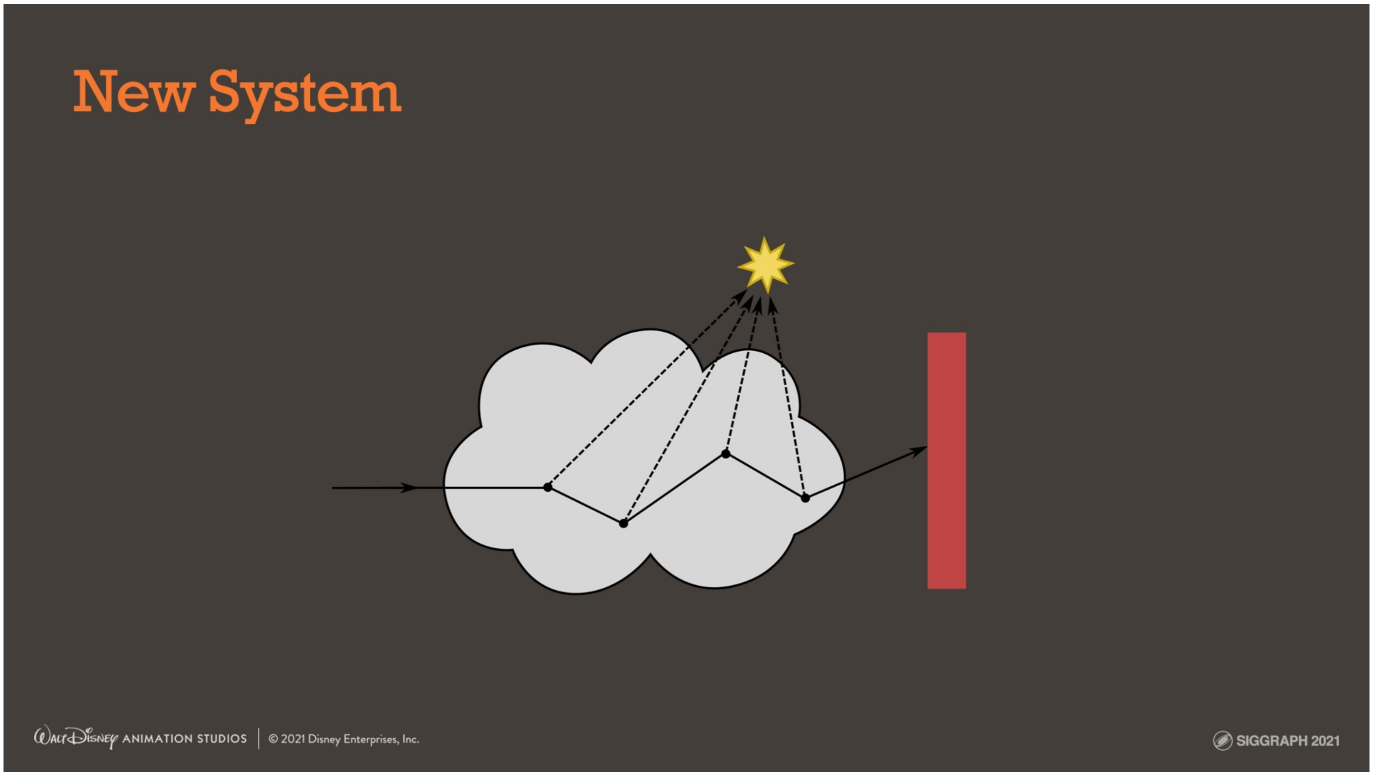

New System

Compared to the ray marching solution, such random walk tracking system no longer requires front to back importance table construction for each scattering event. With each sample having a much lower cost,

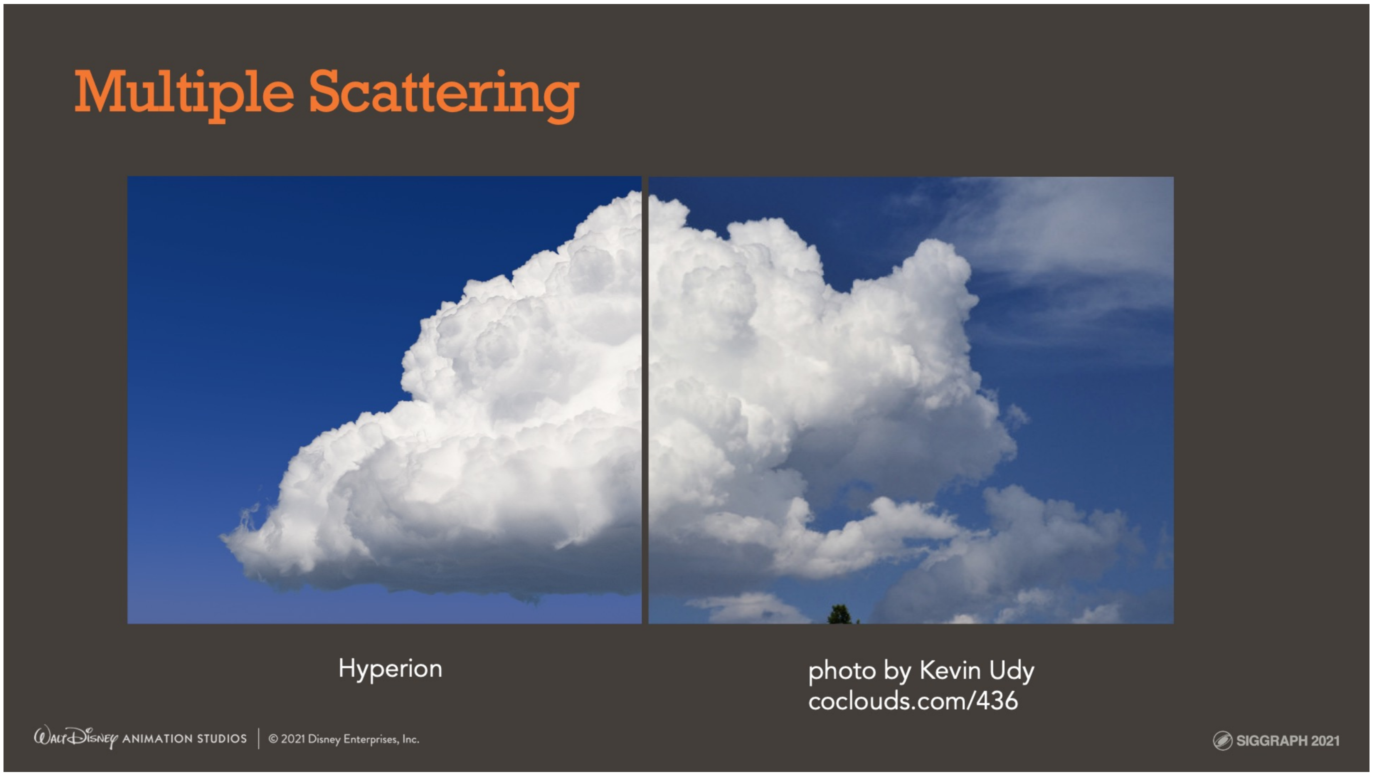

Multiple Scattering

Hyperion

photo by Kevin Udy
coclouds.com/436

It makes multiple scattering in volumes practical. And we were able to reproduce results indistinguishable from photographs.

We were also able to render full cloudscapes with unbounded path lengths…

Challenges

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

However, as I mentioned in the very beginning of the talk, we learned from the production of Ralph Breaks the Internet and Frozen 2, that such volume system also had its disadvantages.

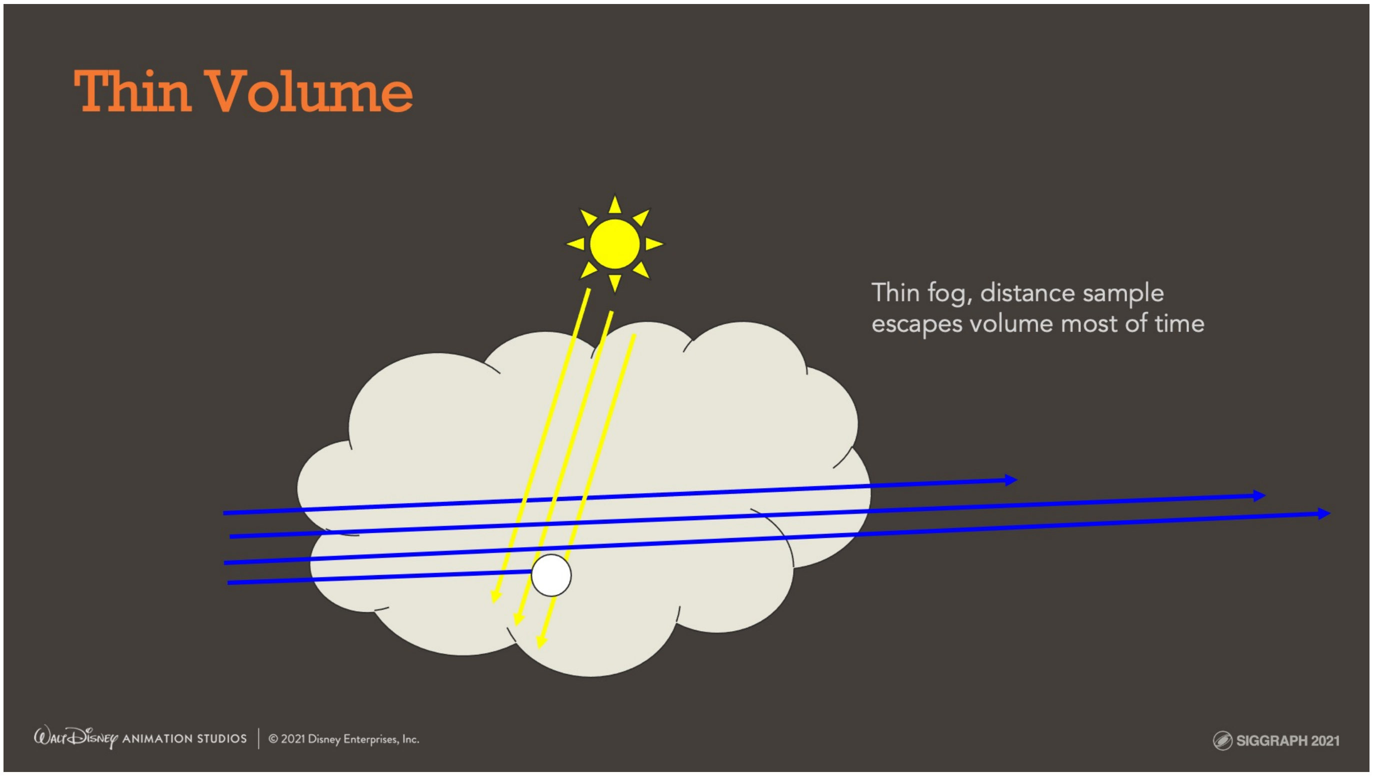

Thin Volume

Thin fog, distance sample escapes volume most of time

For example, when rendering thin fog, the average sampled distance is quite large. As most of the rays go right through the volume, there are few chances to sample a scattering event to evaluate the illumination.

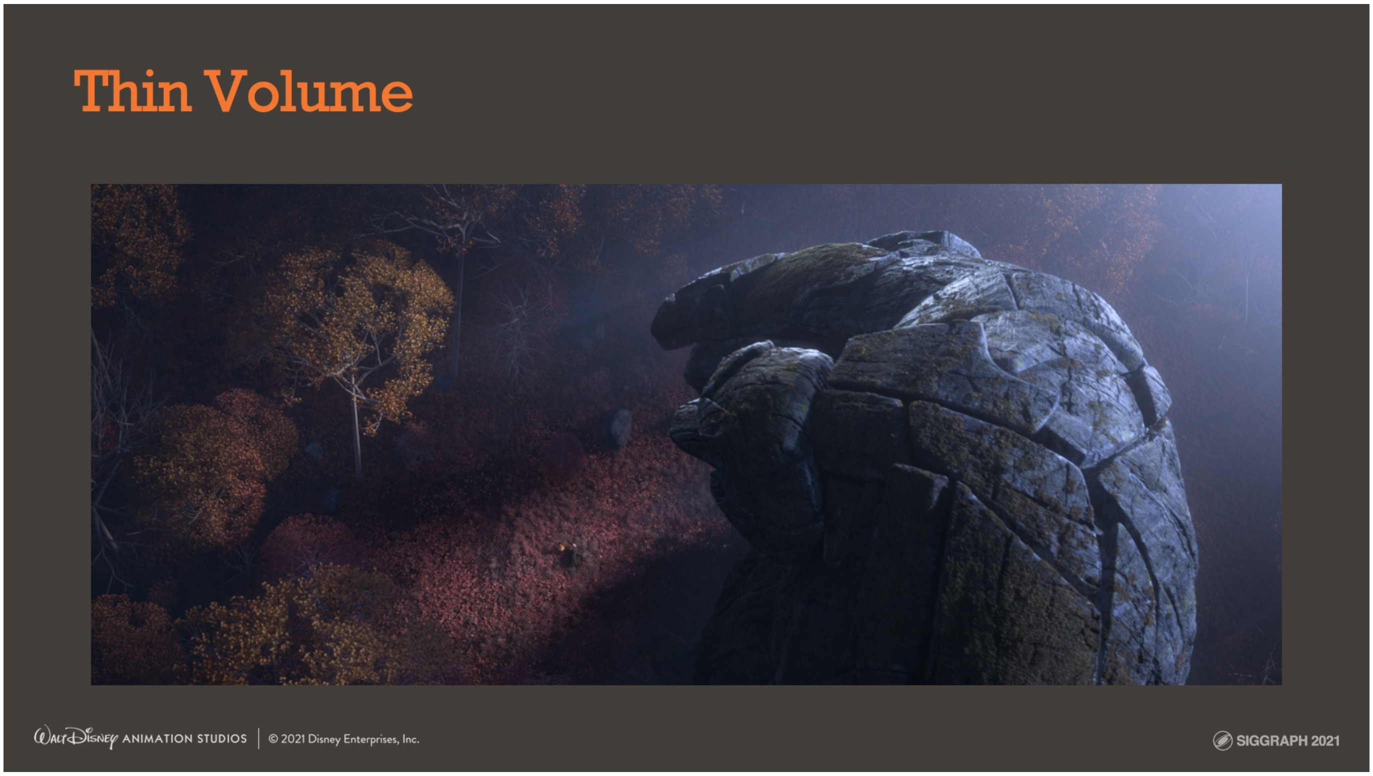

Thin Volume

As a result, effects such as god ray or light shaft were often noisy.

Fire

Same issue applies to fire as well. As fire usually has low extinction, the random walk process again takes too big of a step, passing through fire easily.

Fire

So, capturing details in flame is quite challenging.

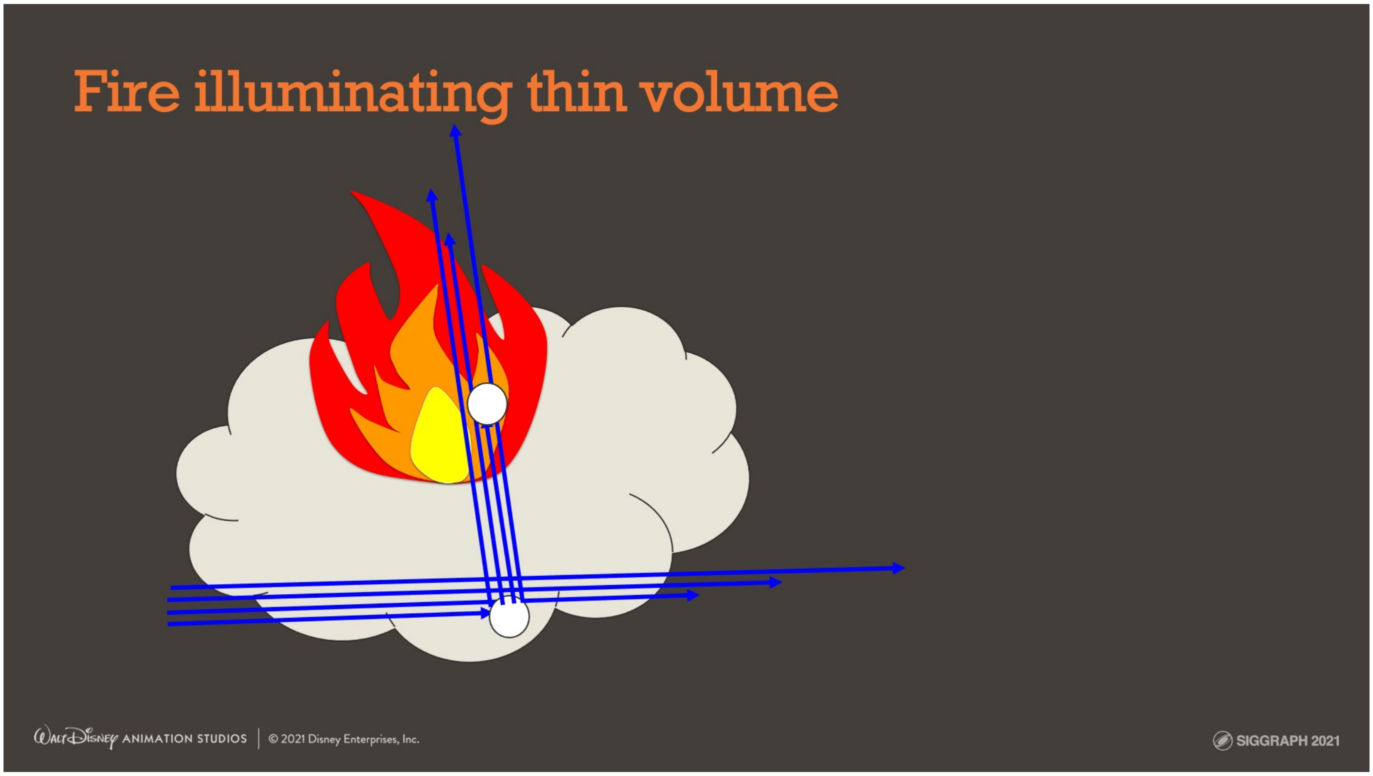

Fire illuminating thin volume

WALT DISNEY ANIMATION STUDIOS |

SIGGRAPH 2021

And not too surprisingly, illuminating a thin volume with fire makes render efficiency even worse.

Fire illuminating thin volume

Walt Disney Animation Studios

SIGGRAPH 2021

Here is a production example of such case, and it often resulted in noisy renders.

Fire behind thick volume

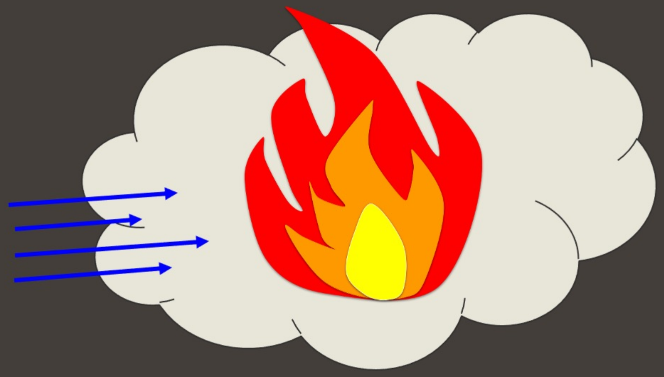

Moreover, on the opposite case, where we have fire inside a thick volume, it is equally challenging since the average sampled distance in this case is now too short that a ray rarely makes to the fire.

Fire behind thick volume

This is problematic not just for rendering an explosion, but also when such explosion illuminates the environment.

Raya Visual Development

Fast forward to the pre-production of the film Raya and the Last Dragon.
From the early visual development work, we saw there were god rays, light shafts,

Raya Visual Development

A misty tunnel lit by torches,

Raya Visual Development

A city full of torch lights covered in thin fog,

Raya Visual Development

And even characters made of emissive volumes wrapped in dense smoke…

It became apparent that some improvements to the existing volume system would have to be made in order to render the film.

Emission Sampling

WALT DISNEY ANIMATION STUDIOS |

SIGGRAPH 2021

We first take a look at what we improved on the emission sampling side:

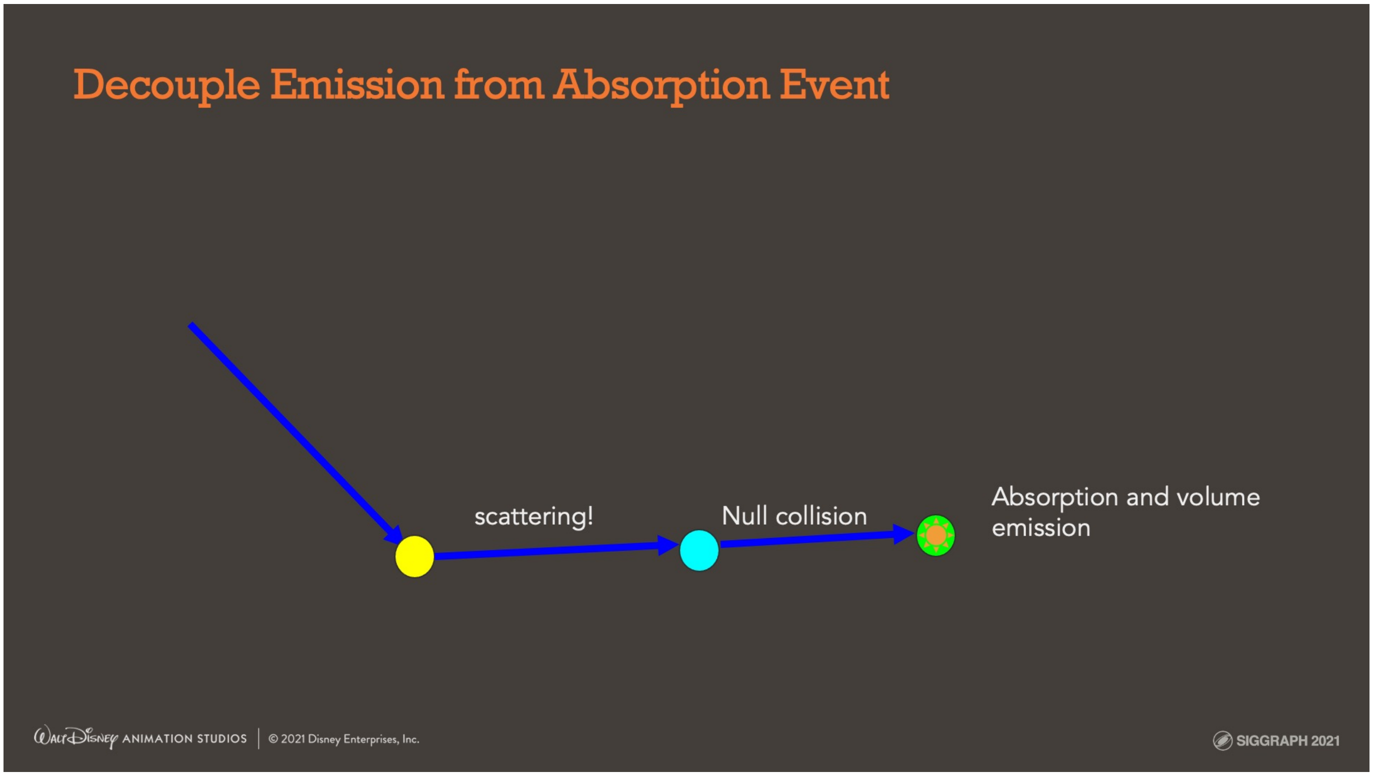

Decouple Emission from Absorption Event

scattering!

Null collision

Absorption and volume emission

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

We noticed that the emission term was only evaluated when absorption events happen. This is an intuitive interpretation of the random walk since we can associate each interaction to one particular event.

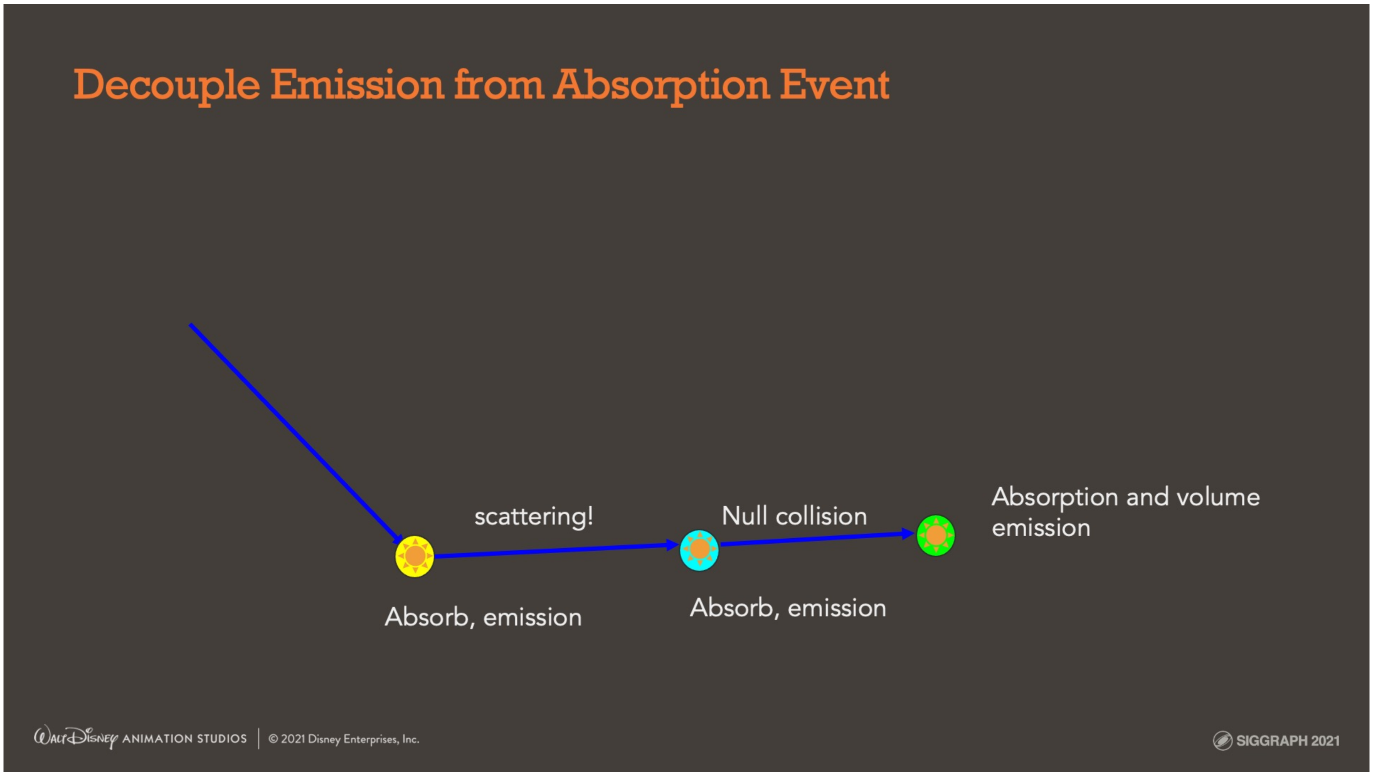

Decouple Emission from Absorption Event

scattering!

Null collision

Absorption and volume emission

Absorb, emission

Absorb, emission

However, we realized that we can actually always evaluate absorption and emission in all three types of events and still be mathematically correct. We don't want to evaluate scattering and null collision events together since it would cause rays to split into two different directions, but absorption won't split the ray and therefore we can always evaluate it. This increased evaluation frequency is compensated for by the increased probability, which makes the estimator converge to the same result faster.

Decouple Emission

Gathering emission at:
absorption only

This is a test render where the renderer only evaluates emission according to probability Pa

Decouple Emission

Gathering emission at:
all steps

And this is an equal sample test render that evaluates emission on all types of events. The noise reduction is quite noticeable.

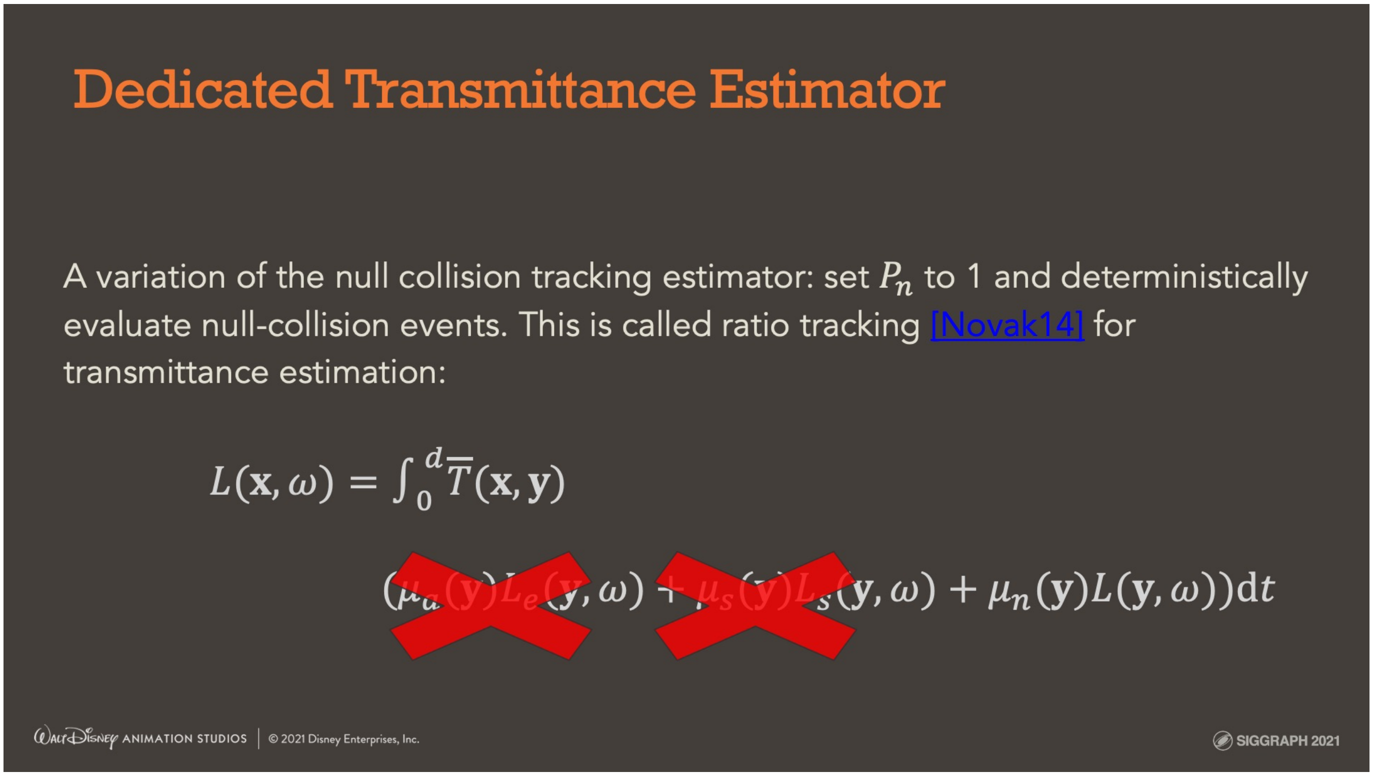

Dedicated Transmittance Estimator

A variation of the null collision tracking estimator: set P_n to 1 and deterministically evaluate null-collision events. This is called ratio tracking [Novak14] for transmittance estimation:

$$L(\mathbf{x},\omega) = \int_0^d \overline{T}(\mathbf{x},\mathbf{y}) \left(\cancel{\mu_a(\mathbf{y})L_e(\mathbf{y},\omega)} + \cancel{\mu_s(\mathbf{y})L_s(\mathbf{y},\omega)} + \mu_n(\mathbf{y})L(\mathbf{y},\omega)\right)\mathrm{d}t$$

Ratio tracking is a dedicated null-collision formulation transmittance estimator that works by ignoring the emission and scattering terms from the equation and allowing the null-collision event to be carried out at all time…

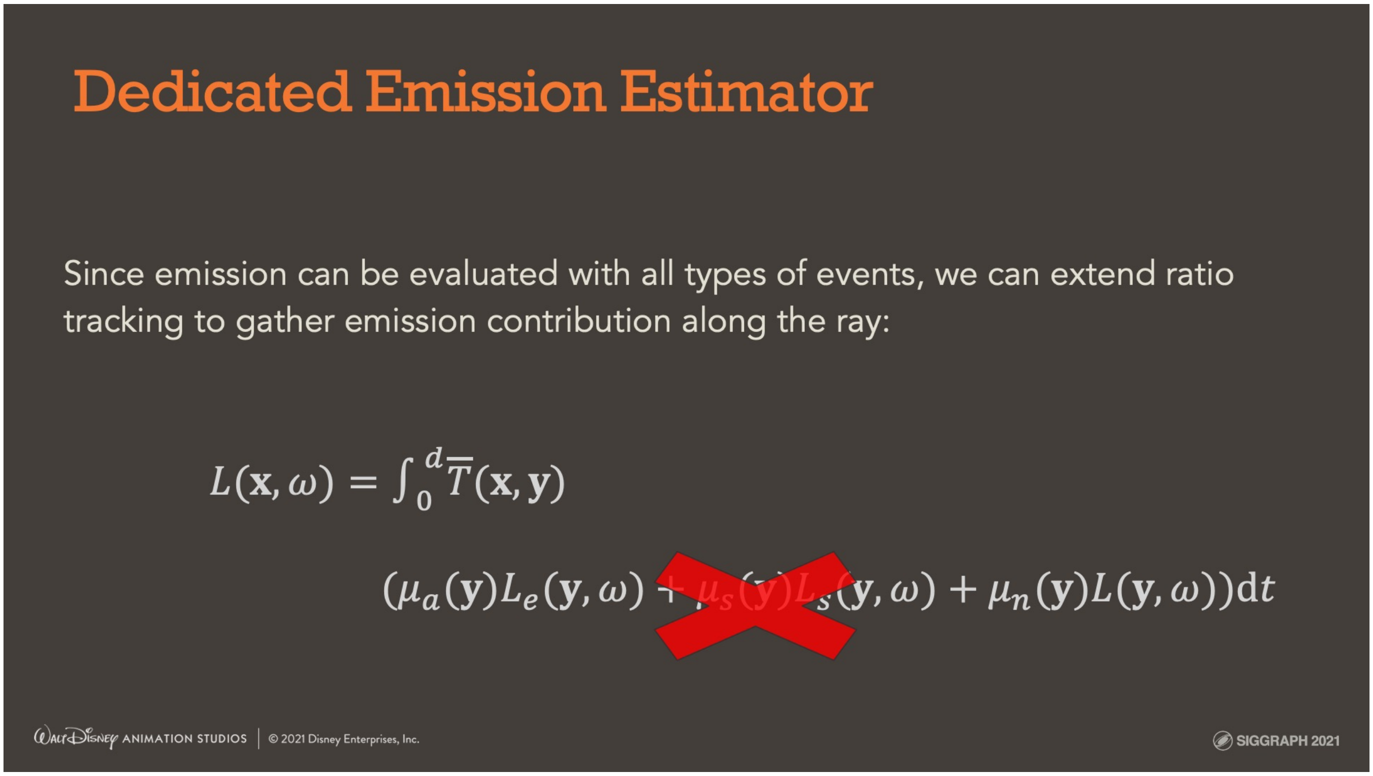

Dedicated Emission Estimator

Since emission can be evaluated with all types of events, we can extend ratio tracking to gather emission contribution along the ray:

$$L(\mathbf{x},\omega) = \int_0^d \overline{T}(\mathbf{x},\mathbf{y})$$

$$(\mu_a(\mathbf{y})L_e(\mathbf{y},\omega) \xcancel{+ \mu_s(\mathbf{y})L_s}(\mathbf{y},\omega) + \mu_n(\mathbf{y})L(\mathbf{y},\omega))\mathrm{d}t$$

As we discussed earlier, emission can be evaluated along with other volume interaction events. If we split the scattering and emission integration into two estimators, we can extend ratio tracking to gather emission contributions along the ray.

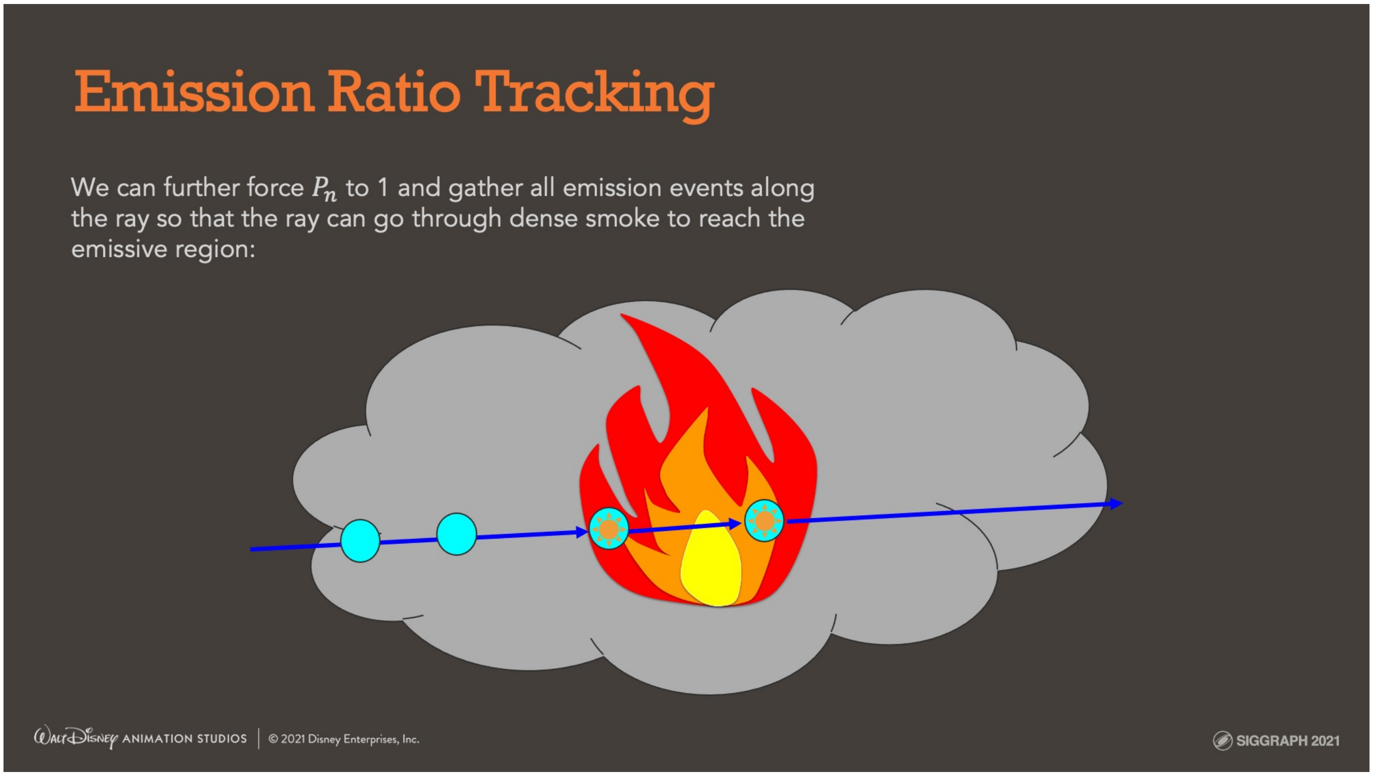

Emission Ratio Tracking

We can further force P_n to 1 and gather all emission events along the ray so that the ray can go through dense smoke to reach the emissive region:

WALT DISNEY ANIMATION STUDIOS

SIGGRAPH 2021

We call this extension “emission ratio tracking”. By setting the null collision probability to 1, this emission tracker can go through dense volumes to reach the emissive region without being absorbed or scattered too early.

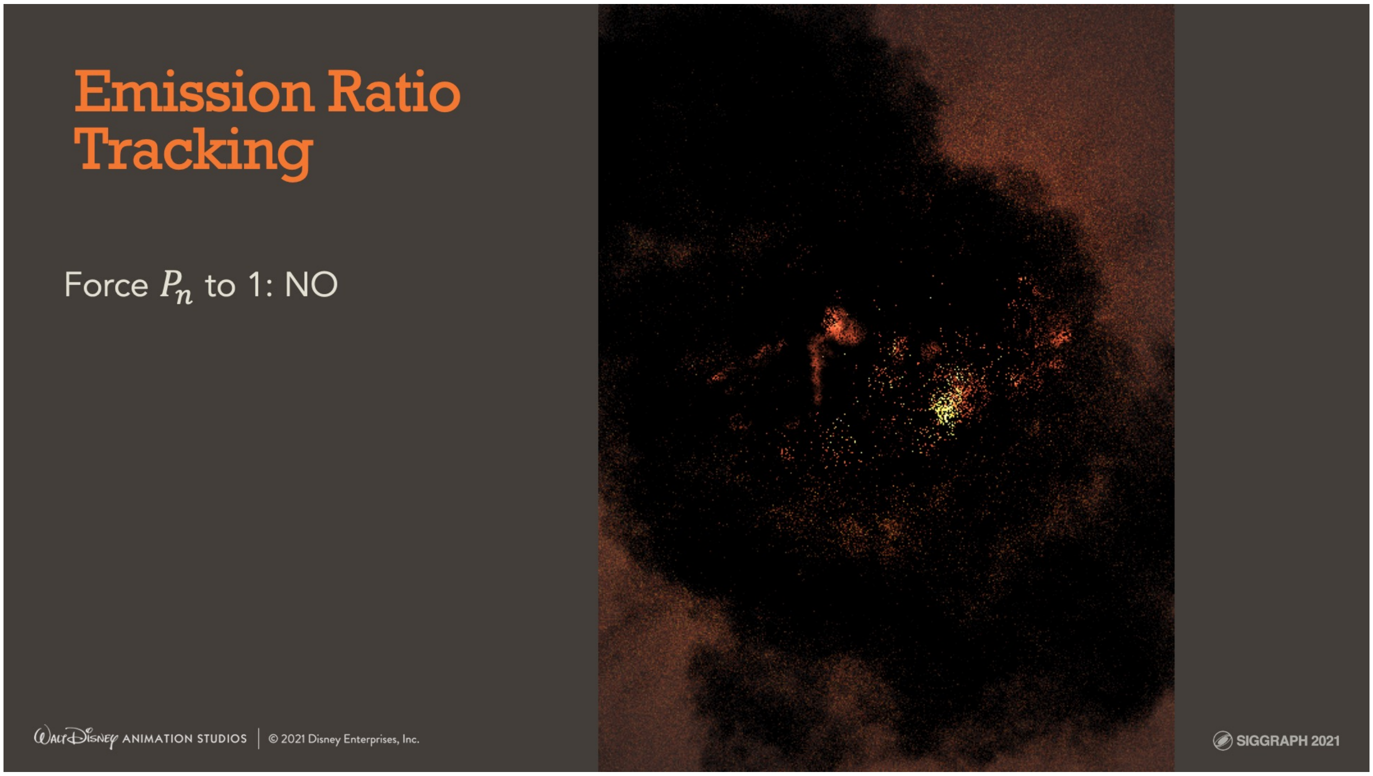

Emission Ratio Tracking

Force P_n to 1: NO

This is a test render that evaluates emission on all types of events. The emission term is difficult to gather since a lot of samples get scattered or absorbed by dense smoke before reaching the emissive region.

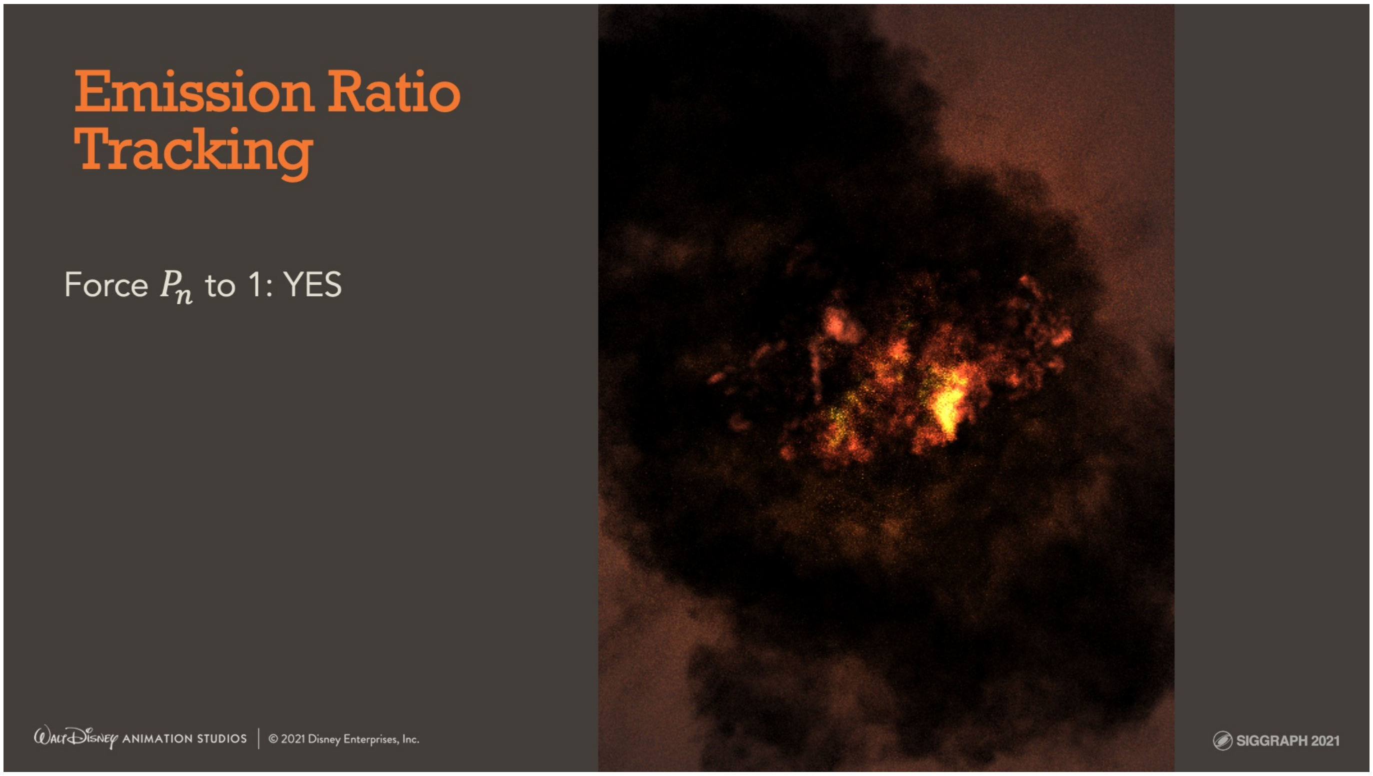

Emission Ratio Tracking

Force P_n to 1: YES

And this is an equal sample test render using emission ratio tracking. The ray now can go through the entire volume and gather the emission term that is buried in the heavy smoke.

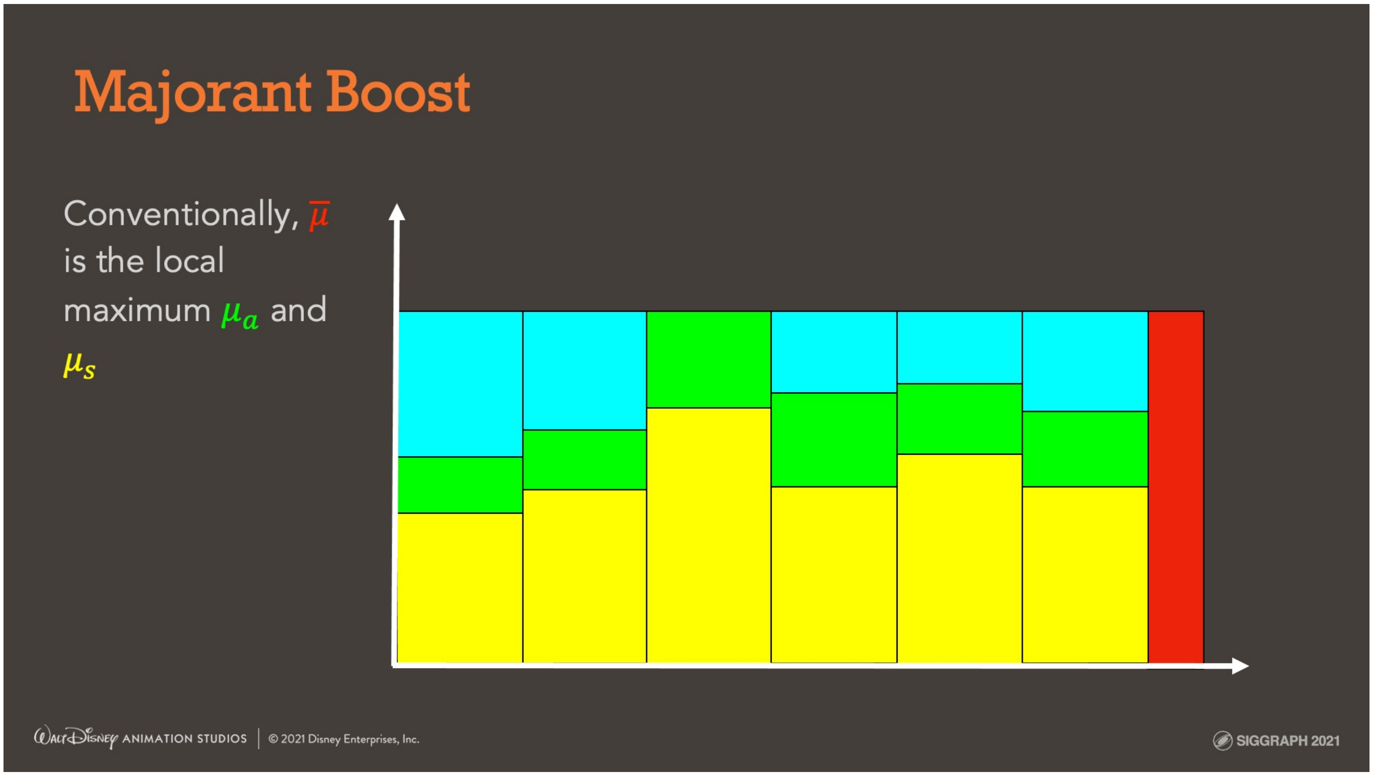

Majorant Boost

Conventionally, $\overline{\mu}$ is the local maximum μ_a and μ_s

For null-collision tracking performance reasons, mu_bar is usually based on a local maximum of the absorption and scattering coefficients

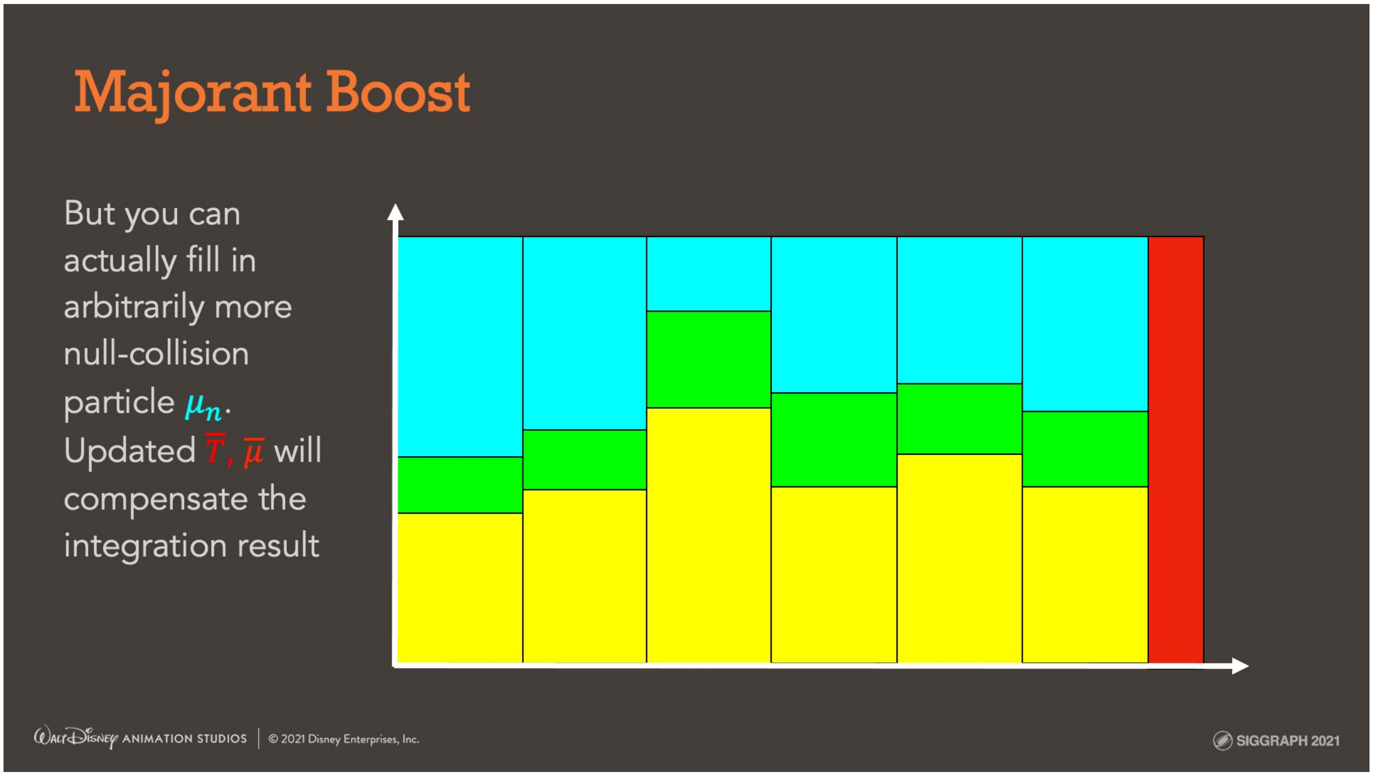

Majorant Boost

But you can actually fill in arbitrarily more null-collision particle μ_n. Updated $\overline{T}, \overline{\mu}$ will compensate the integration result

But the null-collision formulation puts no limitations on the amount of null-collision particles mu_n to homogenize the volume.

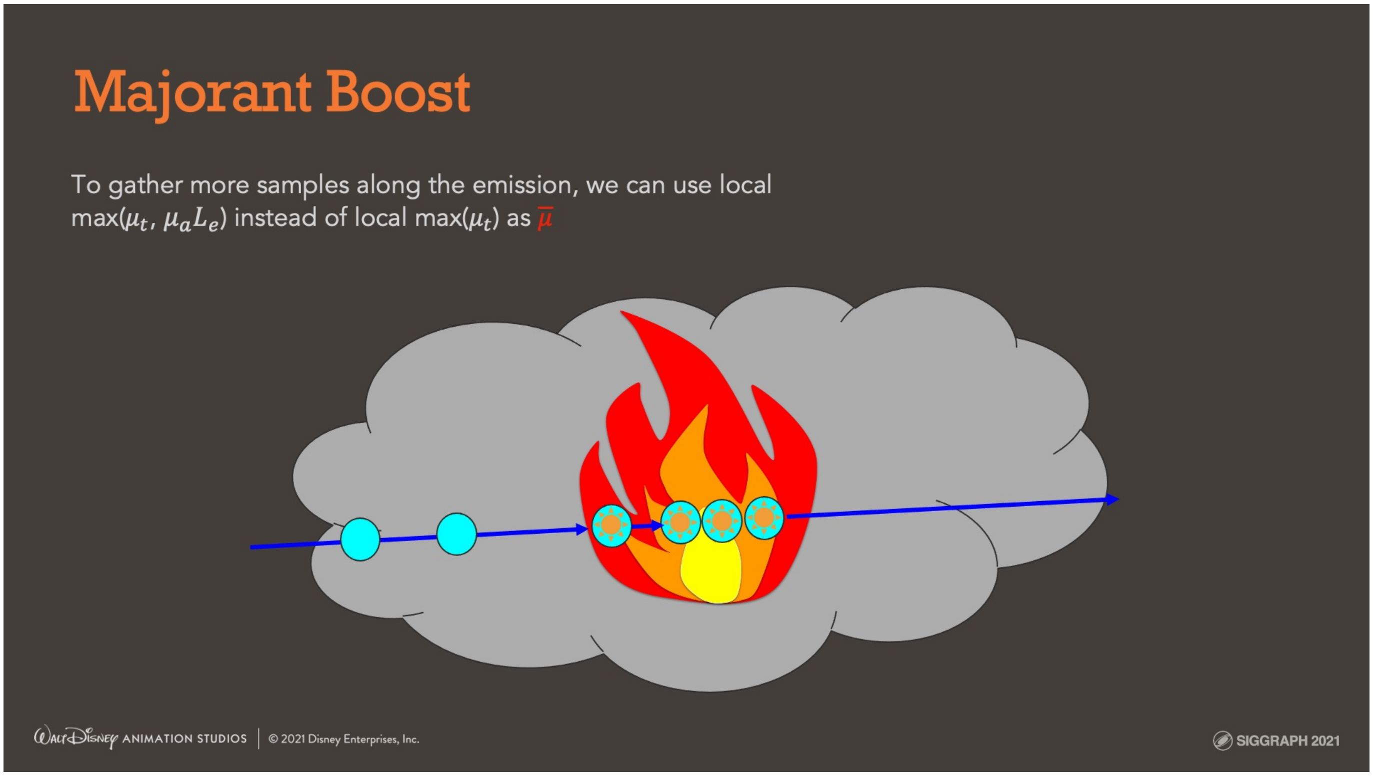

Majorant Boost

To gather more samples along the emission, we can use local max(μ_t, $\mu_a L_e$) instead of local max(μ_t) as $\overline{\mu}$

Since a higher number of mu_bar will result in smaller tracking steps, instead of just using the local maximum of mu_t as mu_bar, we use the local maximum of mu_t “AND” mu_a by Le as mu_bar, so we can take more steps in highly emissive regions.

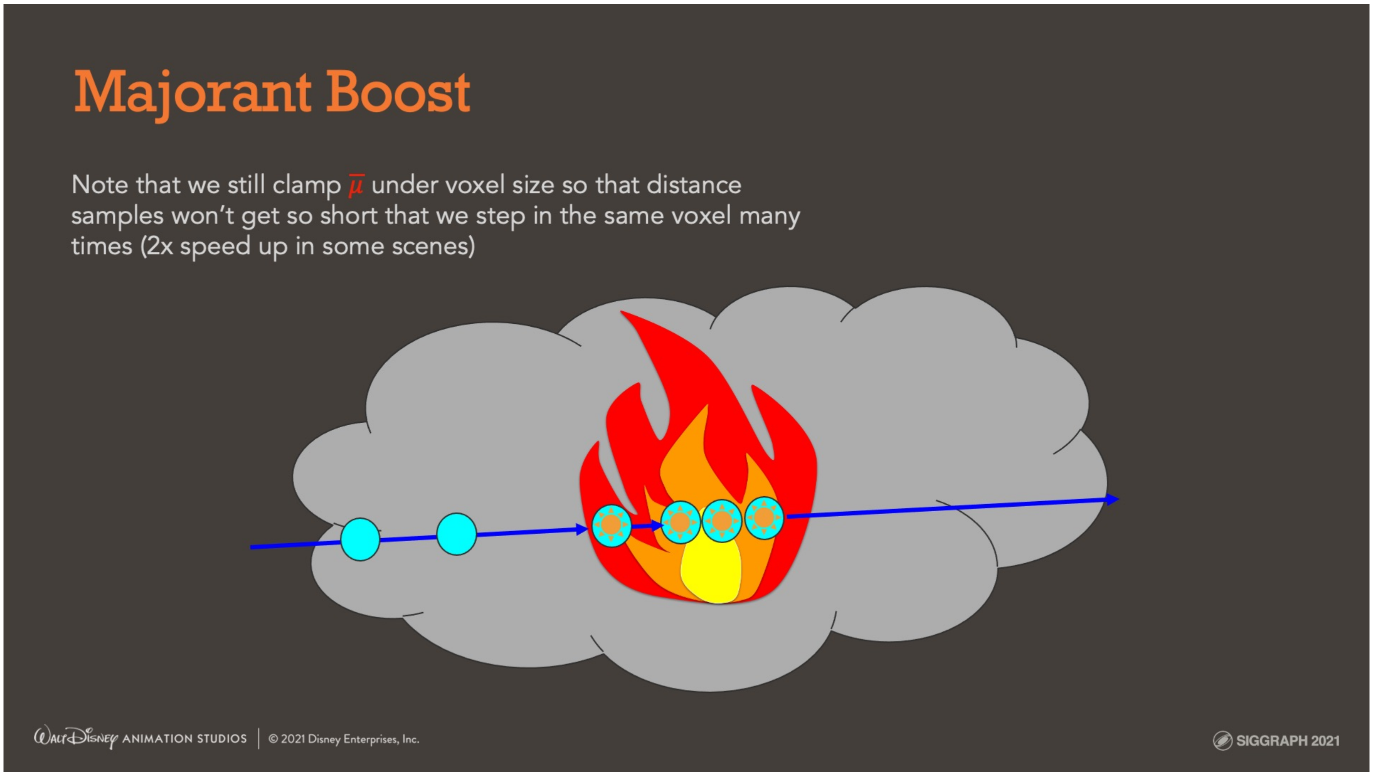

Majorant Boost

Note that we still clamp $\bar{\mu}$ under voxel size so that distance samples won't get so short that we step in the same voxel many times (2x speed up in some scenes)

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

We learned from production experience that mu_bar also needs to be clamped under the volume grid voxel size to prevent stepping through the same voxel too many times. We think that finding a more principled way to determine the optimal mu_bar value is a good topic for future research.

Majorant Boost

Free-path-sampling coefficient:
max extinction

This is a test render using emission ratio tracking and using the mu_t local maximum as mu_bar

Majorant Boost

Free-path-sampling coefficient: max extinction and emission

And this is a test render also using emission ratio tracking, but using the local maximum of mu_t AND mu_a by Le as mu_bar

Emissive Volume as Light Source

With these techniques, we can now efficiently render emissive volumes that are directly visible to camera, but we would also like to illuminate the scene with these emissive volumes.

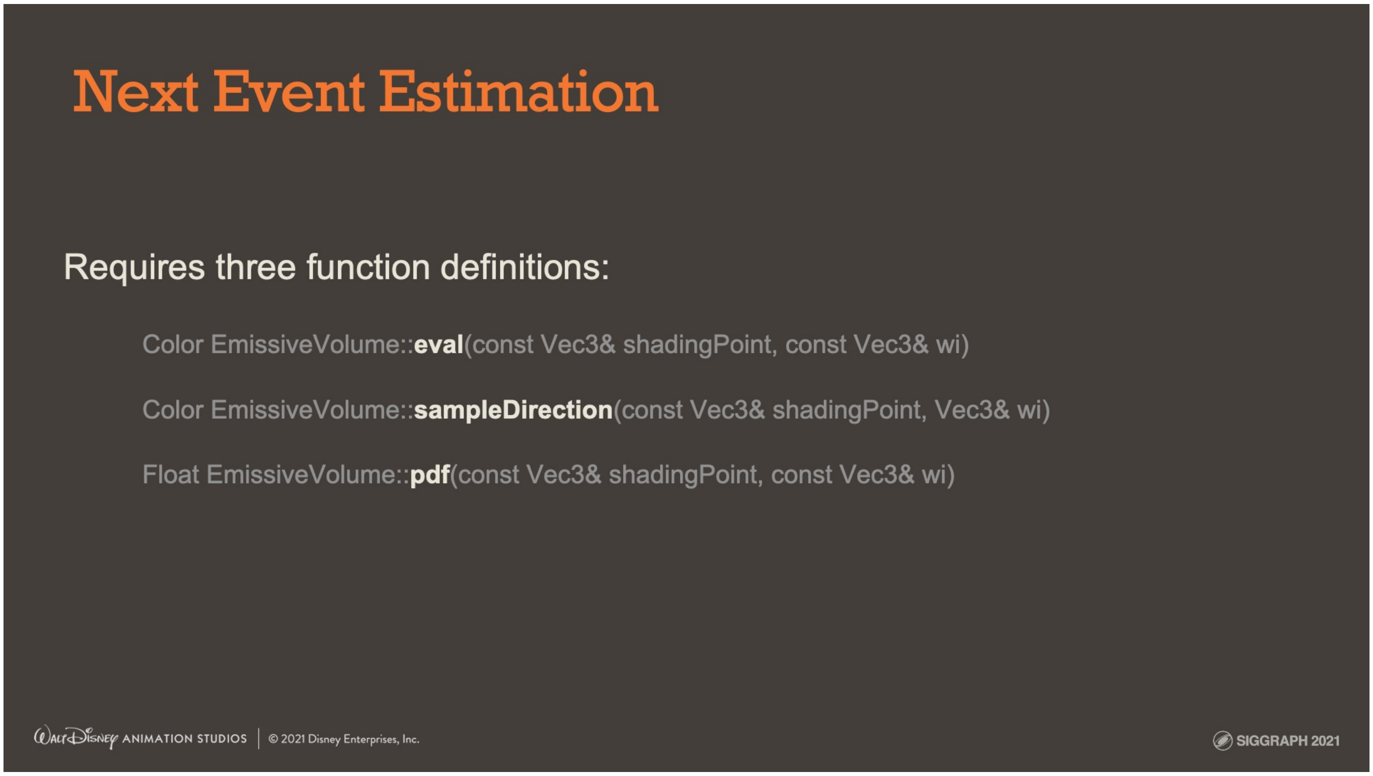

Next Event Estimation

Requires three function definitions:

Color EmissiveVolume::**eval**(const Vec3& shadingPoint, const Vec3& wi)

Color EmissiveVolume::**sampleDirection**(const Vec3& shadingPoint, Vec3& wi)

Float EmissiveVolume::**pdf**(const Vec3& shadingPoint, const Vec3& wi)

In path tracing, this is typically done by Next Event Estimation. For next event estimation, we already know how to evaluate the emissive volume in a given direction, but we still need some way to sample the emissive volume like a light source and some way to evaluate its PDF along a direction.

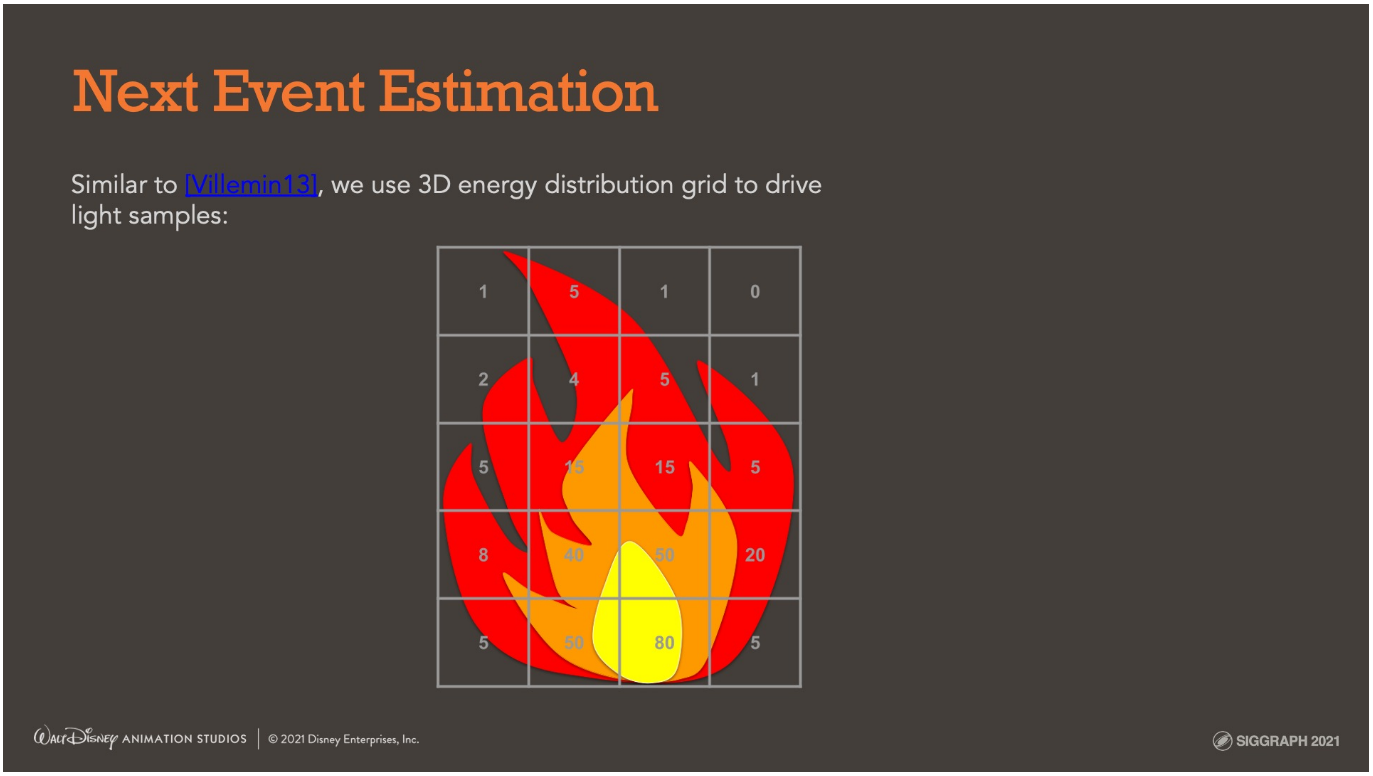

Next Event Estimation

Similar to [Villemin13], we use 3D energy distribution grid to drive light samples:

1	5	1	0
2	4	5	1
5	15	15	5
8	40	50	20
5	50	80	5

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

We extended Villemin’s solution for emissive volume next event estimation. We use a coarse grid representing the 3D energy distribution to make sure that hotter regions have a higher chance of receiving light samples.

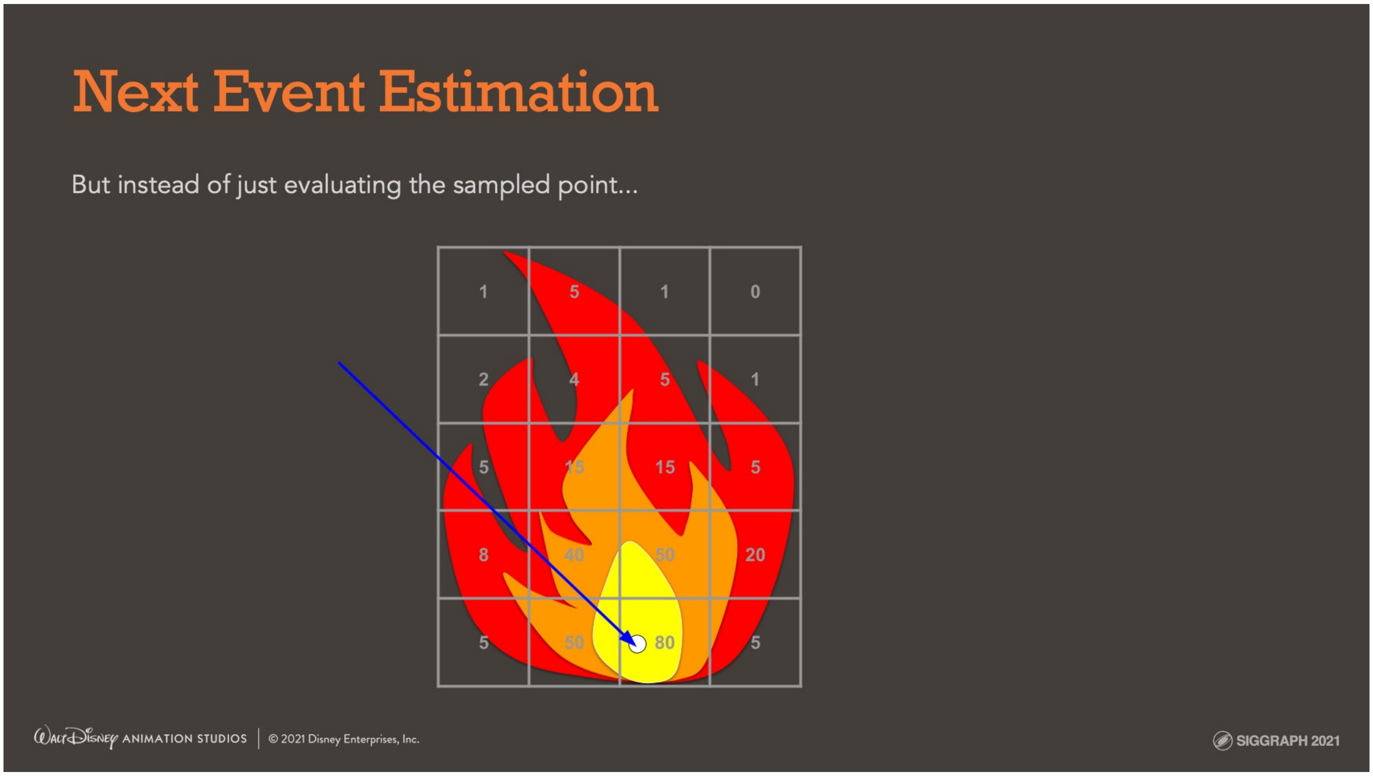

Next Event Estimation

But instead of just evaluating the sampled point…

1	5	1	0
2	4	5	1
5	15	15	5
8	40	50	20
5	50	80	5

Villemin’s technique used point sampling, which has difficulty when emission is obscured by heavy smoke or when the emissive region is large. These cases require very high sample counts to capture emission details in glossy reflections.

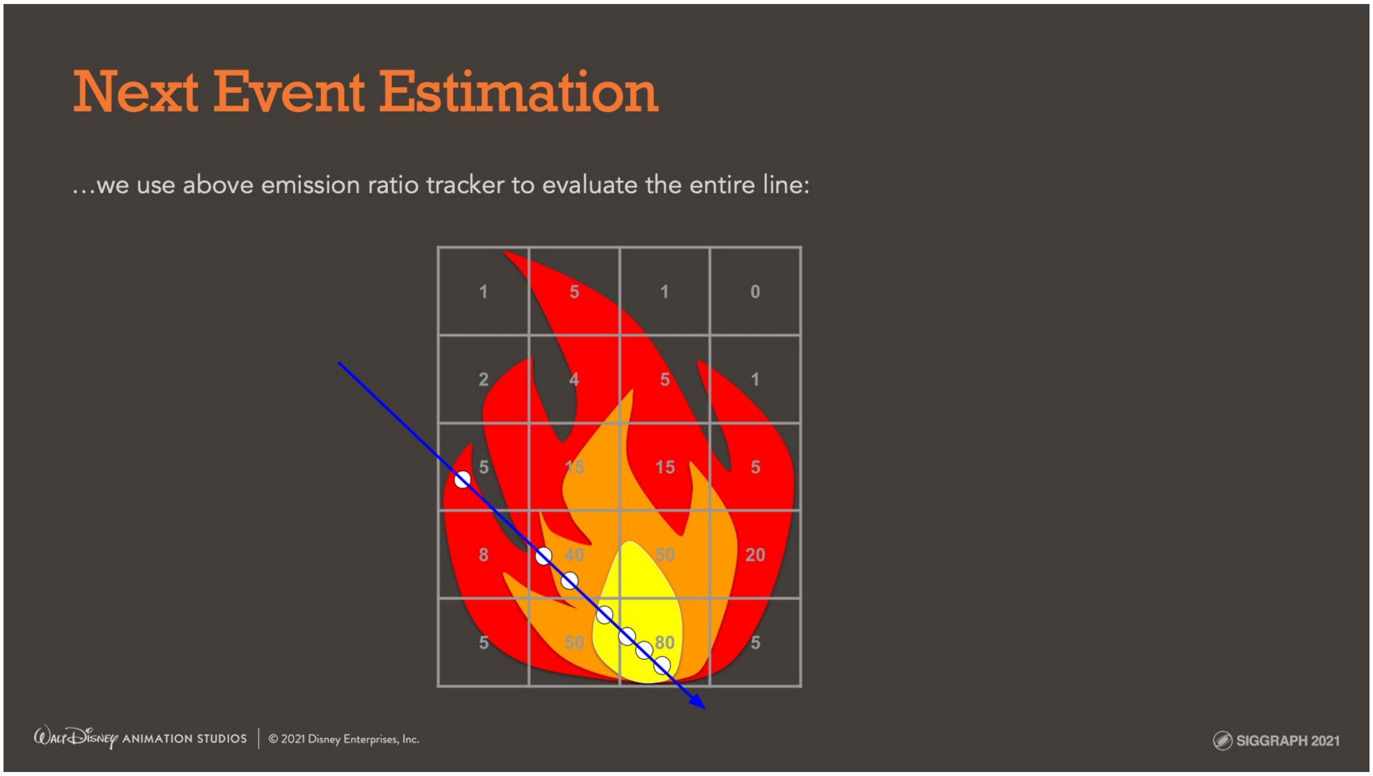

Next Event Estimation

…we use above emission ratio tracker to evaluate the entire line:

1	5	1	0
2	4	5	1
5	15	15	5
8	40	50	20
5	50	80	5

Therefore, we use the emission ratio tracker introduced earlier to evaluate every tracking point along the ray, gathering more information in each light sample.

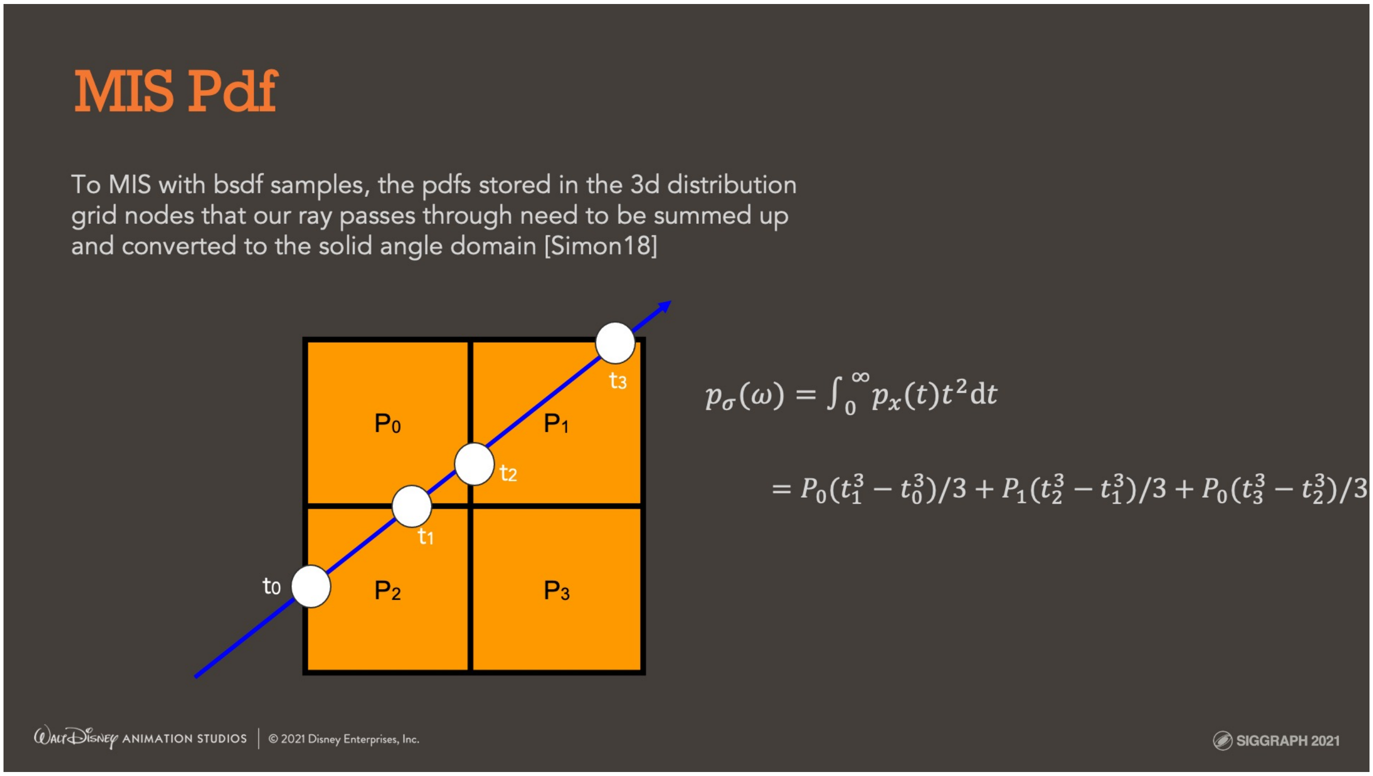

MIS Pdf

To MIS with bsdf samples, the pdfs stored in the 3d distribution grid nodes that our ray passes through need to be summed up and converted to the solid angle domain [Simon18]

$$p_\sigma(\omega) = \int_0^\infty p_x(t)t^2\mathrm{d}t$$

$$= P_0(t_1^3 - t_0^3)/3 + P_1(t_2^3 - t_1^3)/3 + P_0(t_3^3 - t_2^3)/3$$

Walt Disney ANIMATION STUDIOS

SIGGRAPH 2021

Finally, to use multiple importance sampling to combine bsdf samples with light samples in the solid angle domain, we need to integrate pdfs stored in distribution grid cells that the light sample ray passed through using a Jacobian transform.

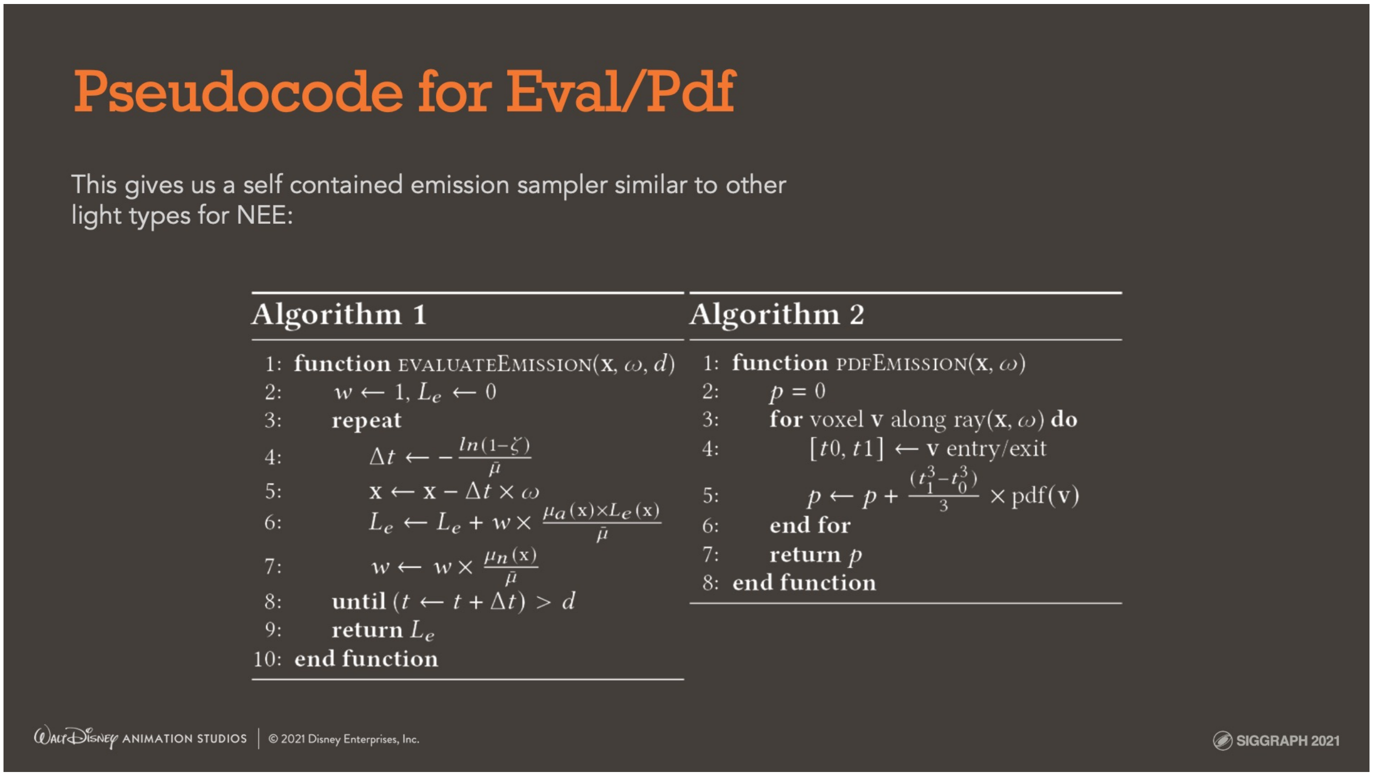

Pseudocode for Eval/Pdf

This gives us a self contained emission sampler similar to other light types for NEE:

Algorithm 1

```
function EVALUATEEMISSION($\mathbf{x}, \omega, d$)
    $w \leftarrow 1, L_e \leftarrow 0$
    repeat
        $\Delta t \leftarrow -\frac{ln(1-\zeta)}{\bar{\mu}}$
        $\mathbf{x} \leftarrow \mathbf{x} - \Delta t \times \omega$
        $L_e \leftarrow L_e + w \times \frac{\mu_a(\mathbf{x}) \times L_e(\mathbf{x})}{\bar{\mu}}$
        $w \leftarrow w \times \frac{\mu_n(\mathbf{x})}{\bar{\mu}}$
    until $(t \leftarrow t + \Delta t) > d$
    return $L_e$
end function
```

Algorithm 2

```
function PDFEMISSION($\mathbf{x}, \omega$)
    $p = 0$
    for voxel $\mathbf{v}$ along ray($\mathbf{x}, \omega$) do
        $[t0, t1] \leftarrow \mathbf{v}$ entry/exit
        $p \leftarrow p + \frac{(t_1^3 - t_0^3)}{3} \times \mathrm{pdf}(\mathbf{v})$
    end for
    return $p$
end function
```

This gives us a self-contained emission sampler like any other type of light source. The pseudo code can also be found in the talk abstract.

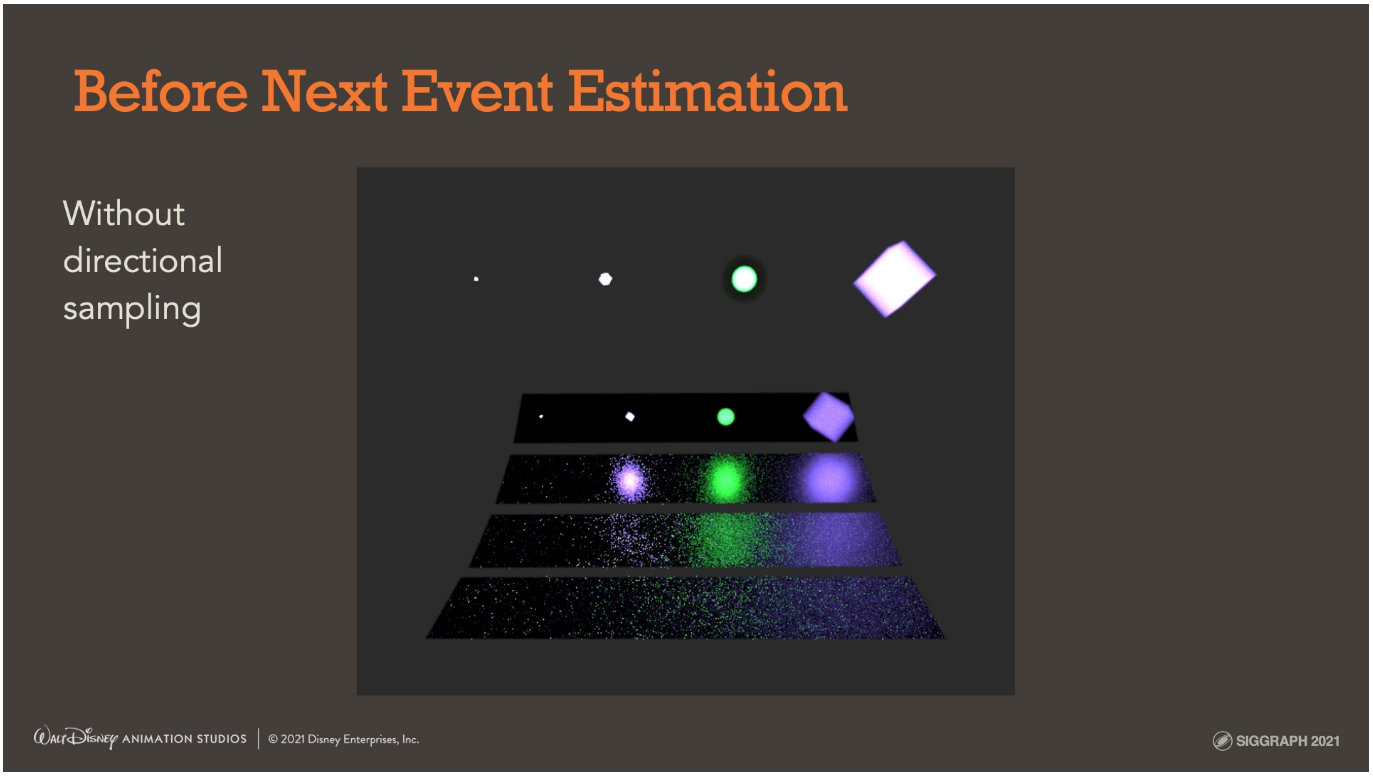

Before Next Event Estimation

Without directional sampling

This is a test render using only bsdf samples to capture the emissive volume contribution; the diffuse and high roughness surfaces are noisy since few samples are lucky enough to hit the emissive volume.

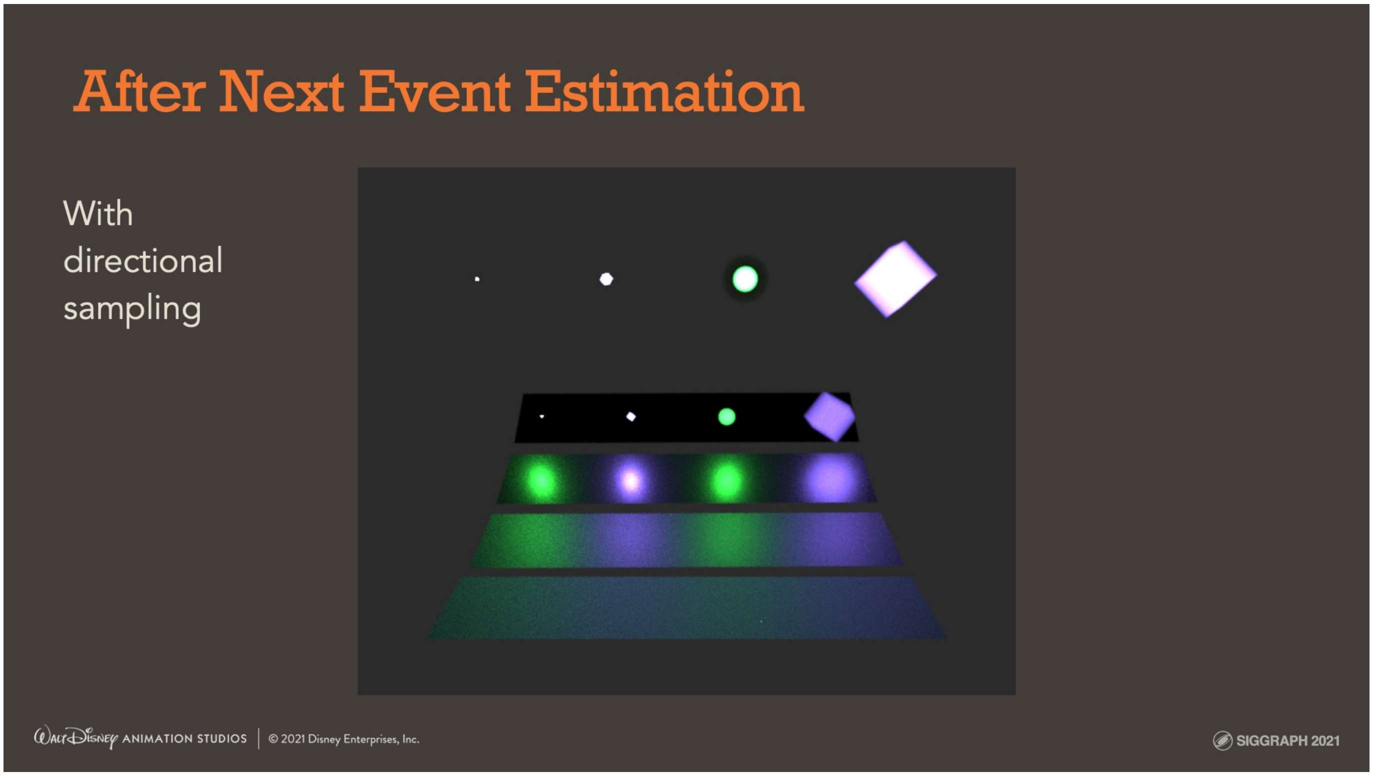

After Next Event Estimation

With directional sampling

And this is an equal sample test render combining next event estimation with bsdf sampling using multiple importance sampling.

Original approach
Optimized emission sampling

We put the above emission sampling improvements together in our volume renderer. In this before and after equal-sample-count comparison from production, our emission sampling improvements capture noticeably more detail in the surface/volume illumination.

Scattering Sampling

Now we switch gears to look into the scattering sampling side of the problem.

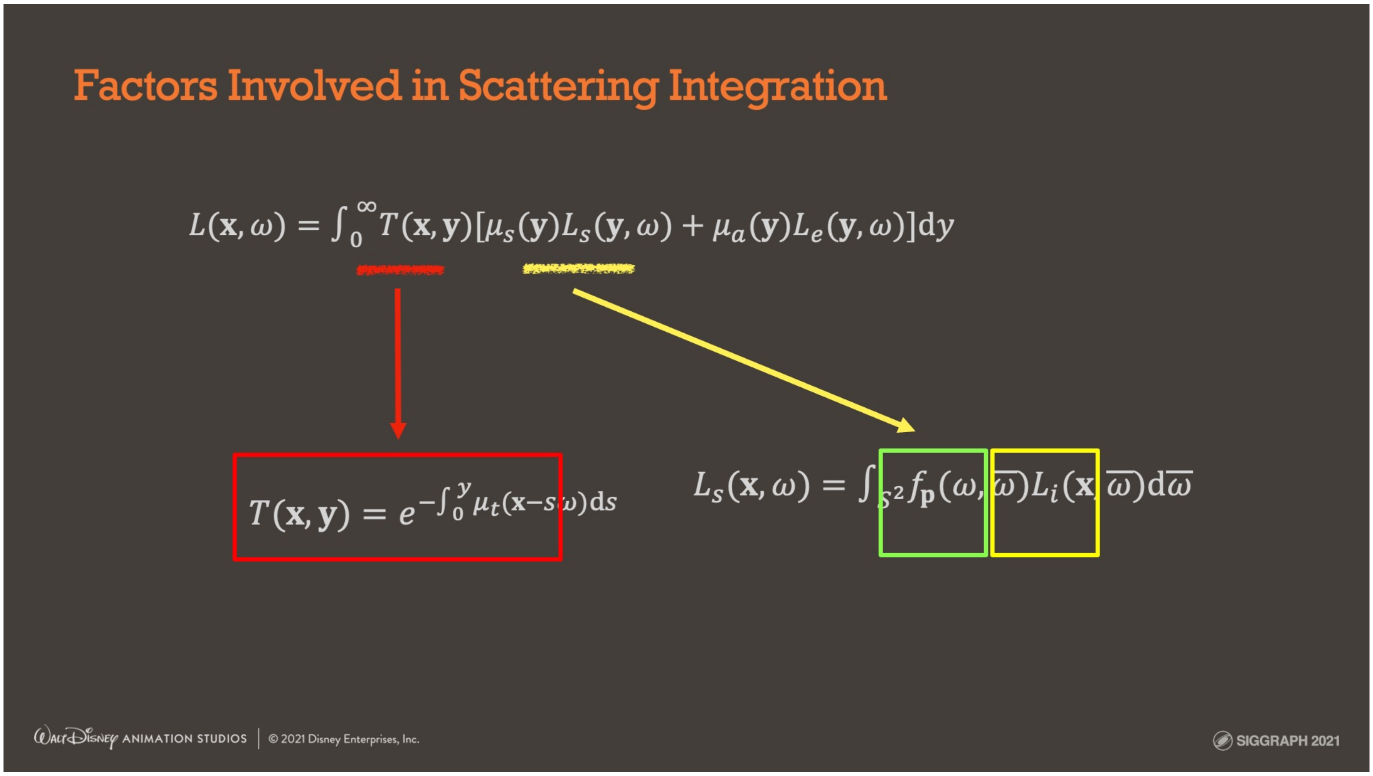

Factors Involved in Scattering Integration

$$L(\mathbf{x},\omega) = \int_0^\infty T(\mathbf{x},\mathbf{y})[\mu_s(\mathbf{y})L_s(\mathbf{y},\omega) + \mu_a(\mathbf{y})L_e(\mathbf{y},\omega)]\mathrm{d}y$$

$$T(\mathbf{x},\mathbf{y}) = e^{-\int_0^y \mu_t(\mathbf{x}-s\omega)\mathrm{d}s}$$

$$L_s(\mathbf{x},\omega) = \int_{S^2} f_{\mathbf{p}}(\omega,\overline{\omega})L_i(\mathbf{x},\overline{\omega})\mathrm{d}\overline{\omega}$$

Volume scattering is mainly affected by the following factors in the volume rendering equation:
Transmittance, Radiance and Phase function

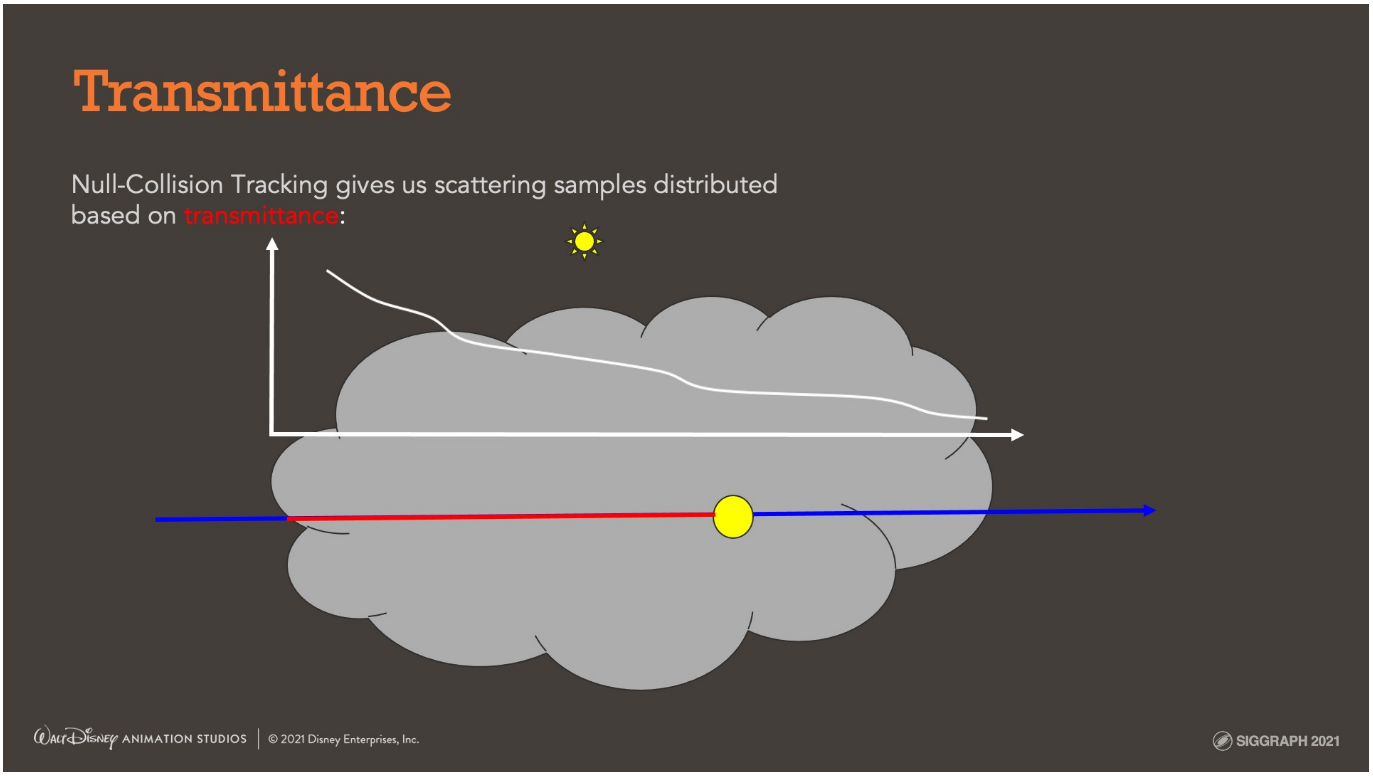

Transmittance

Null-Collision Tracking gives us scattering samples distributed based on transmittance:

Null-collision distance sampling gives us scattering samples based on the transmittance term

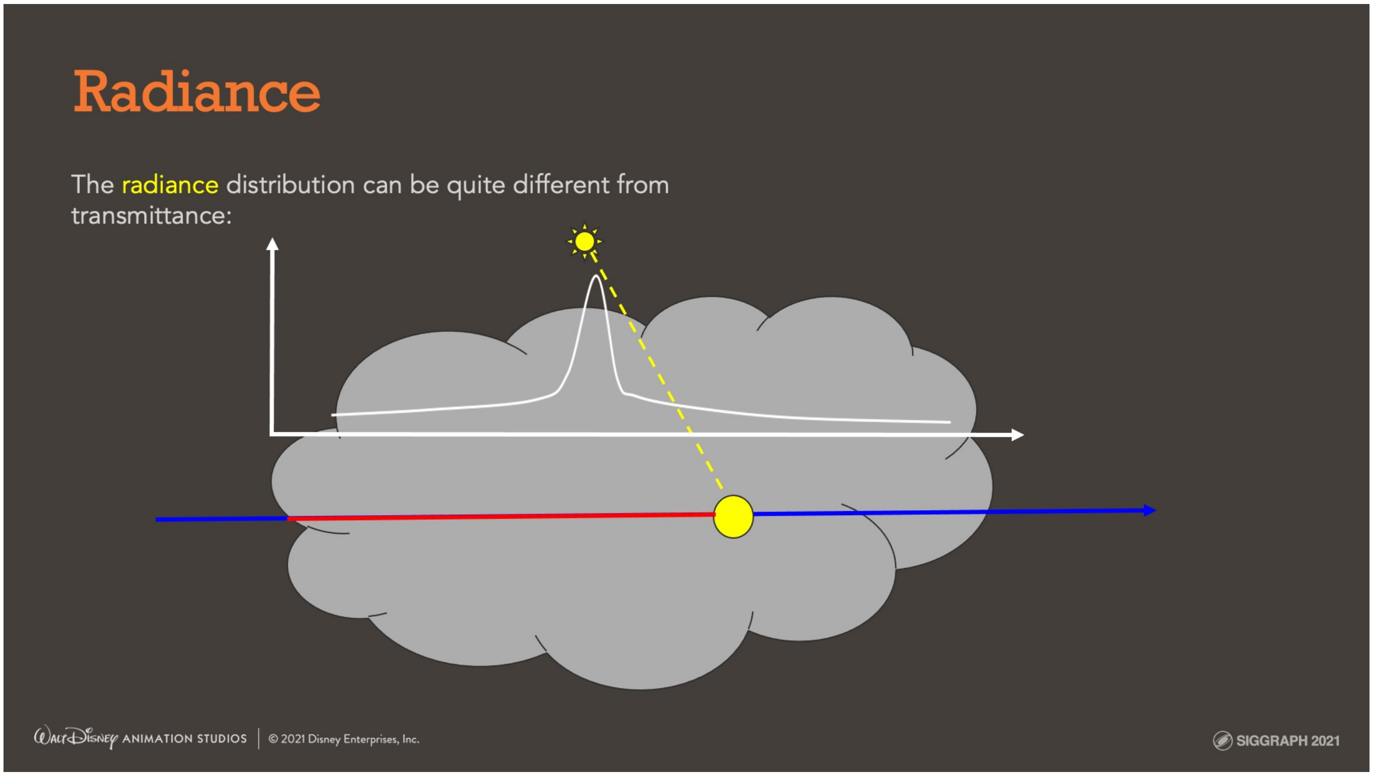

Radiance

The radiance distribution can be quite different from transmittance:

The radiance term distribution can be quite different from the transmittance, and there are techniques like equi-angular sampling for cases dominated by this term...

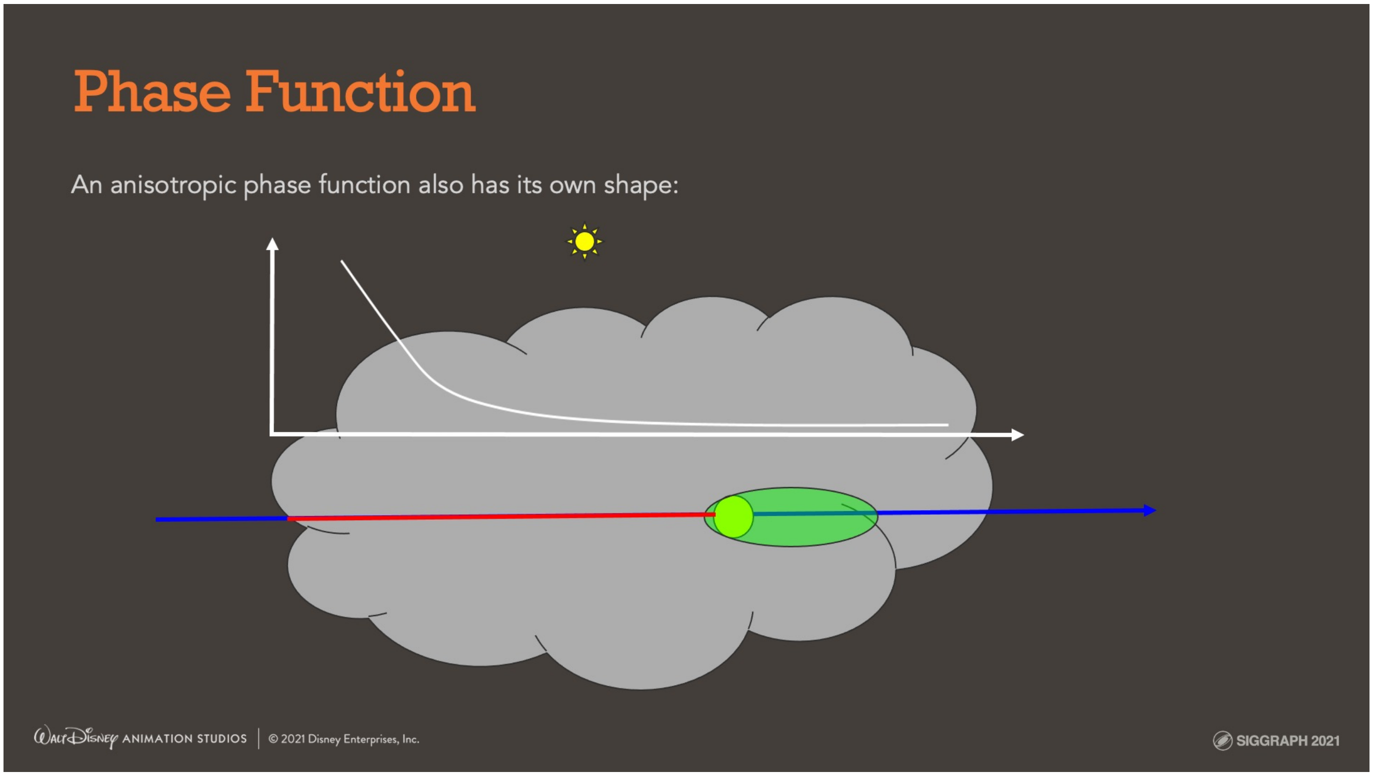

Phase Function

An anisotropic phase function also has its own shape:

A forward scattering phase function will have a front peak distribution.

Common solution to handle these different distributions is to use multiple importance sampling to combine different sample strategies

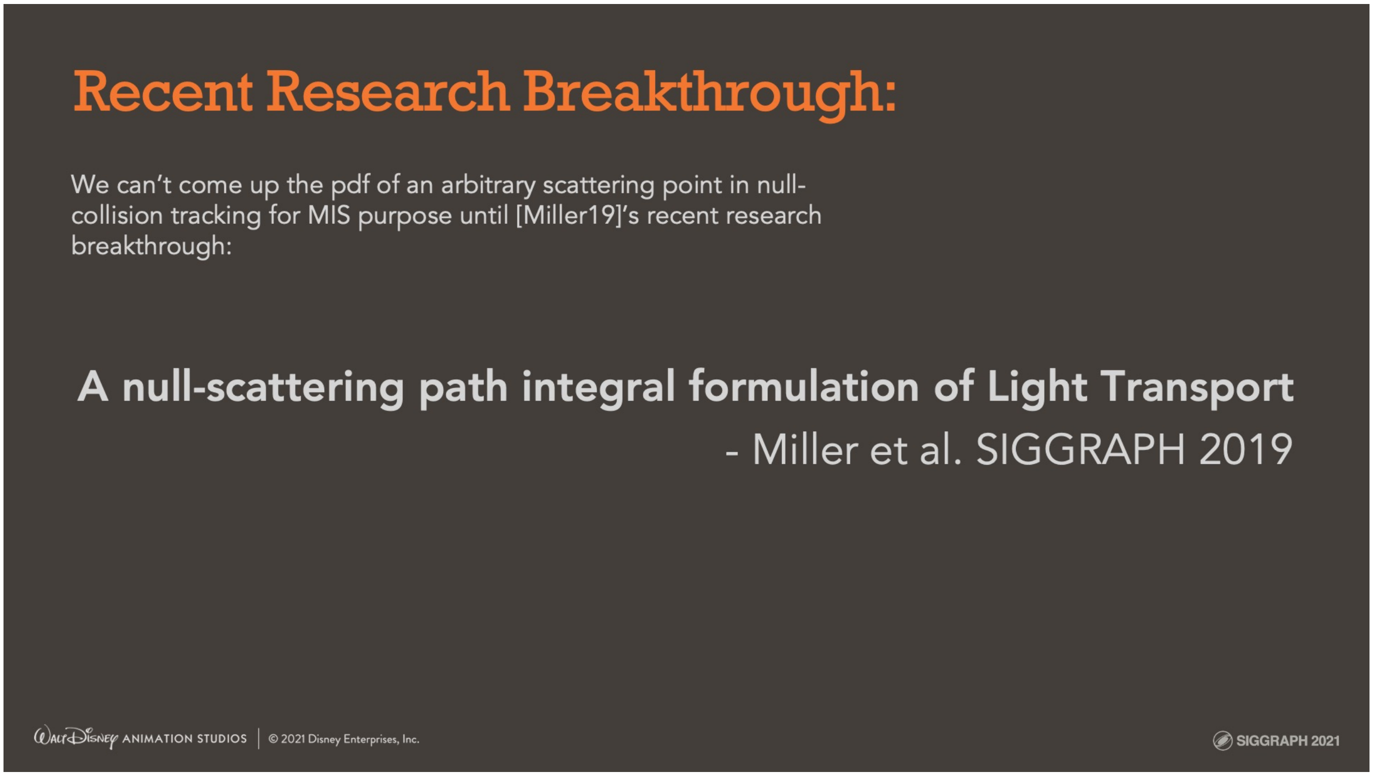

Recent Research Breakthrough:

We can't come up the pdf of an arbitrary scattering point in null-collision tracking for MIS purpose until [Miller19]'s recent research breakthrough:

A null-scattering path integral formulation of Light Transport

- Miller et al. SIGGRAPH 2019

However, compute distance sampling pdf for an arbitrary scattering point for MIS was an unsolved problem until the recent research breakthrough. After Miller et al. derived the analytical pdf for null-collision path formulations, we started experimenting with combining other strategies with null collision tracking.

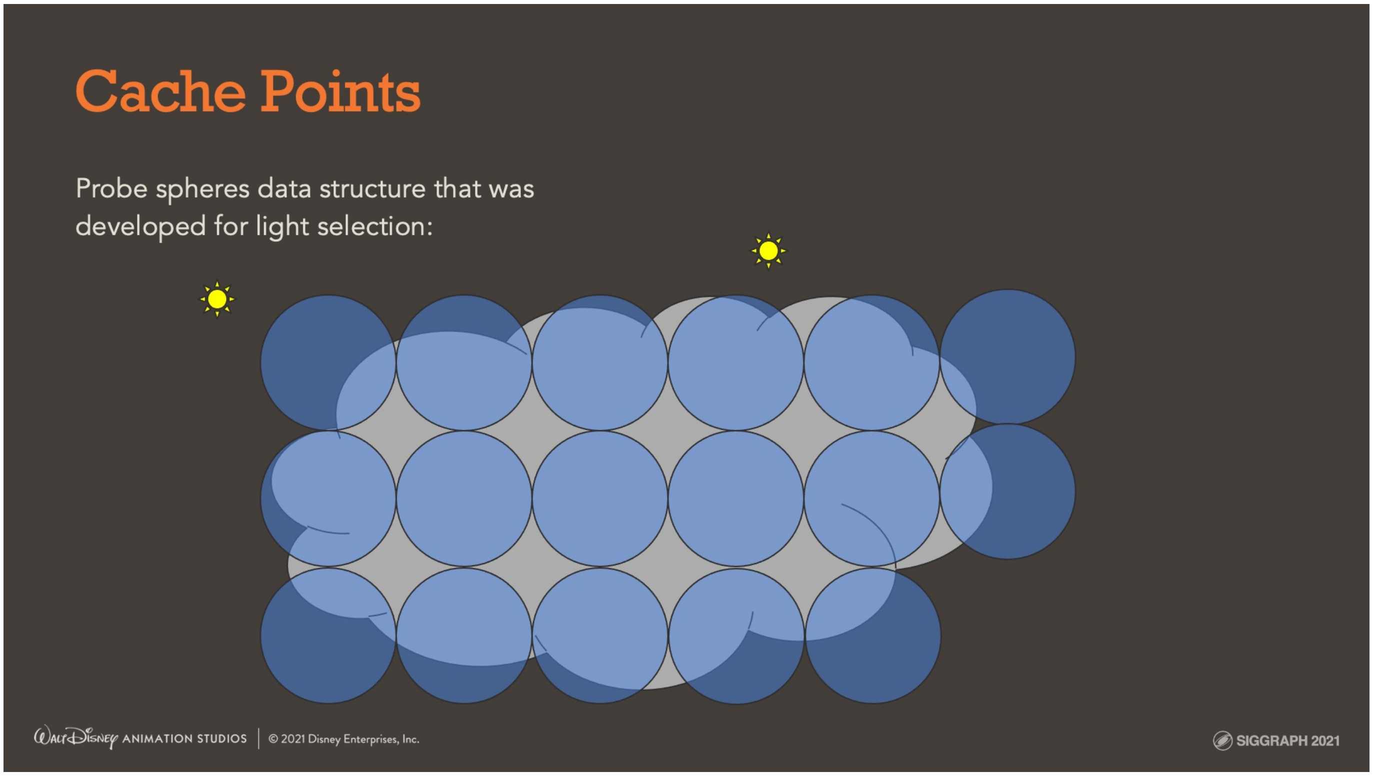

Cache Points

Probe spheres data structure that was developed for light selection:

Hyperion uses a probe-spheres data structure called cache points to handle many-light selection, and we thought it is possible to reuse them to guide volumetric scattering distance samples.

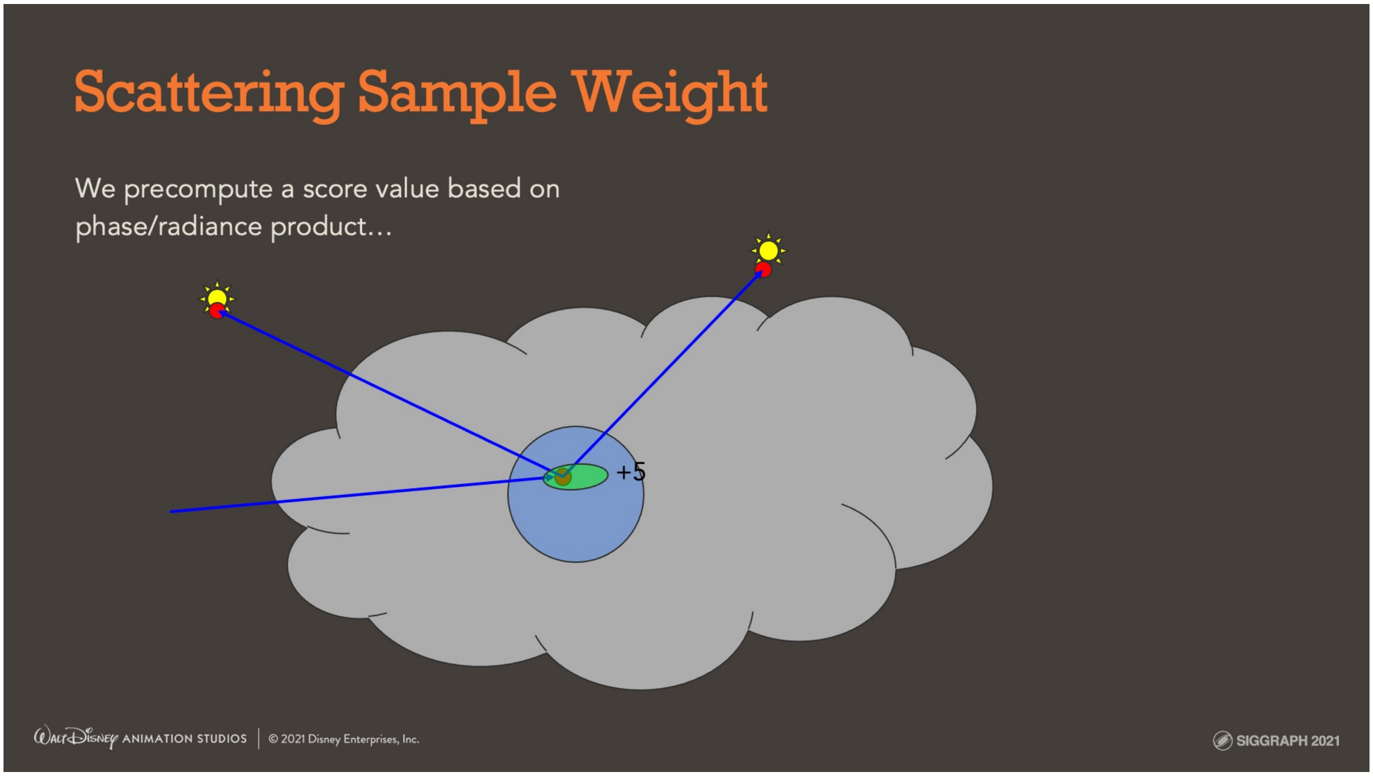

Scattering Sample Weight

We precompute a score value based on phase/radiance product…

For each probe sphere populated in volume, we go through emissive sources and draw samples from the emissive sources, from the scattering point within the probe sphere, and from the camera lens in order to approximate the scattering source term. We use this approximation to assign a score value, which we call the scattering sample weight.

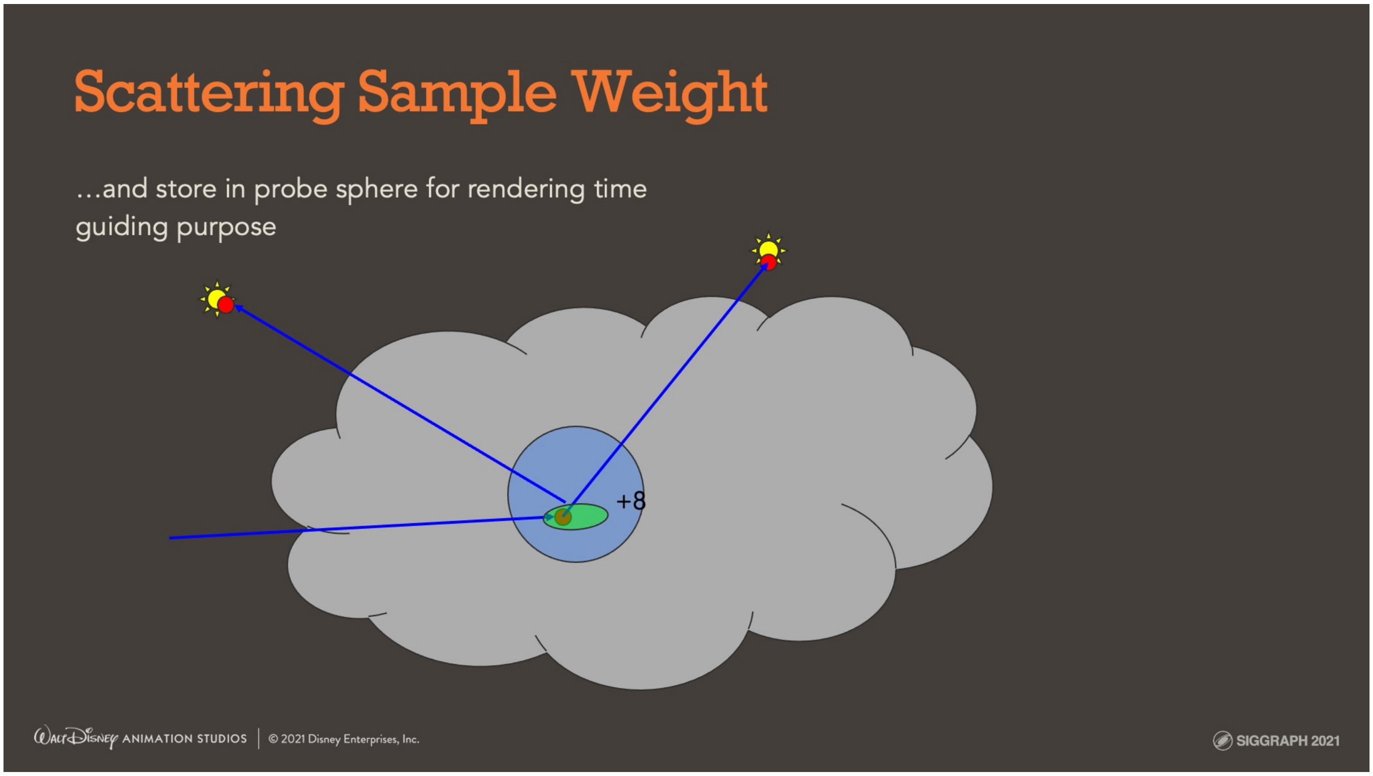

Scattering Sample Weight

…and store in probe sphere for rendering time guiding purpose

WALT DISNEY ANIMATION STUDIOS |

SIGGRAPH 2021

We then store the score value into the probe sphere for later guiding purpose

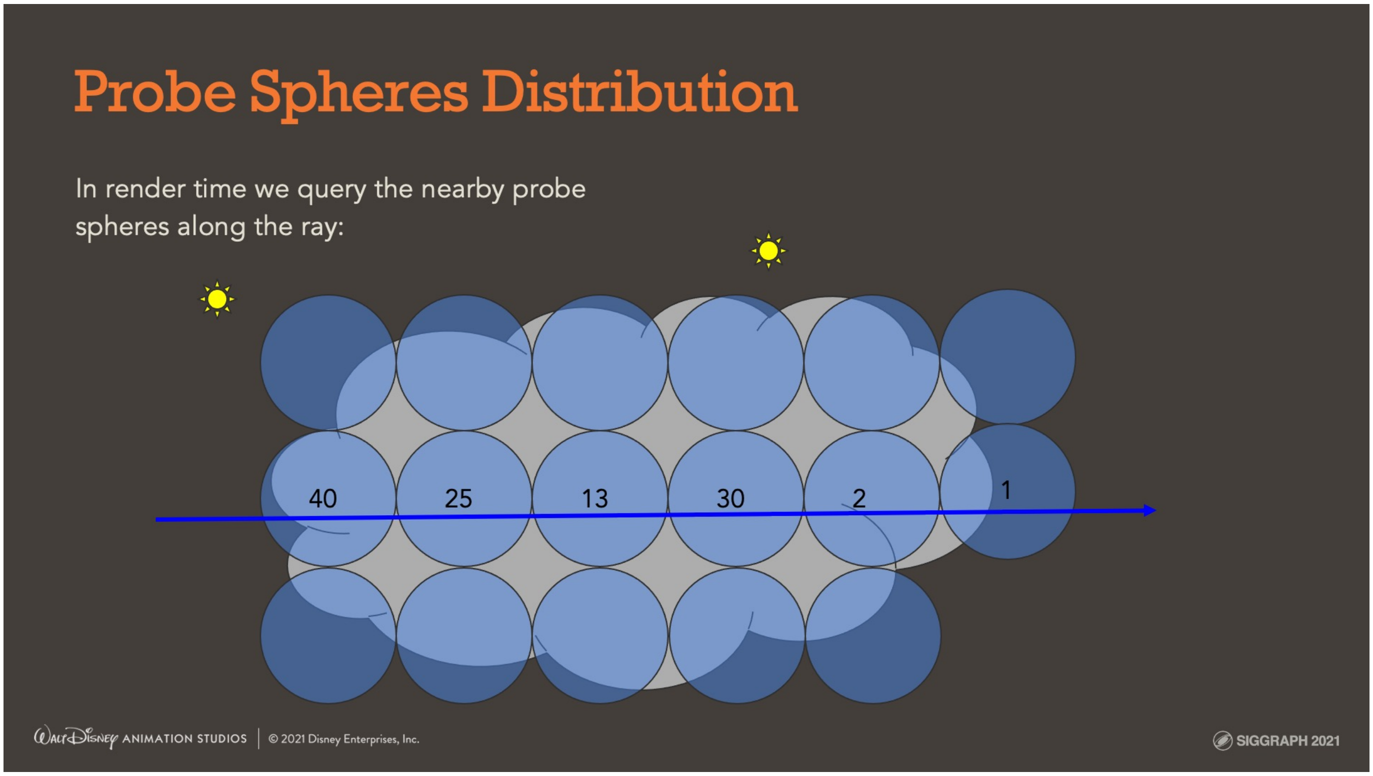

Probe Spheres Distribution

In render time we query the nearby probe spheres along the ray:

During the ray tracing stage, we query the nearby probe spheres along the ray.

Probe Spheres Distribution

In render time we query the nearby probe spheres along the ray:

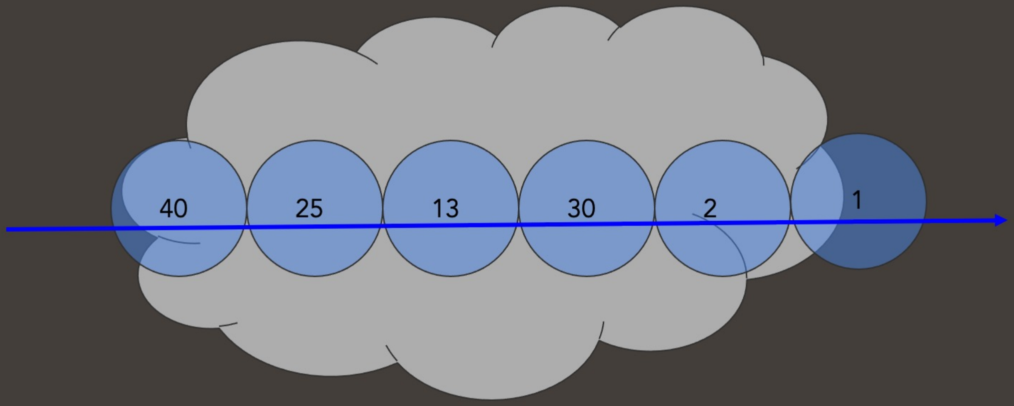

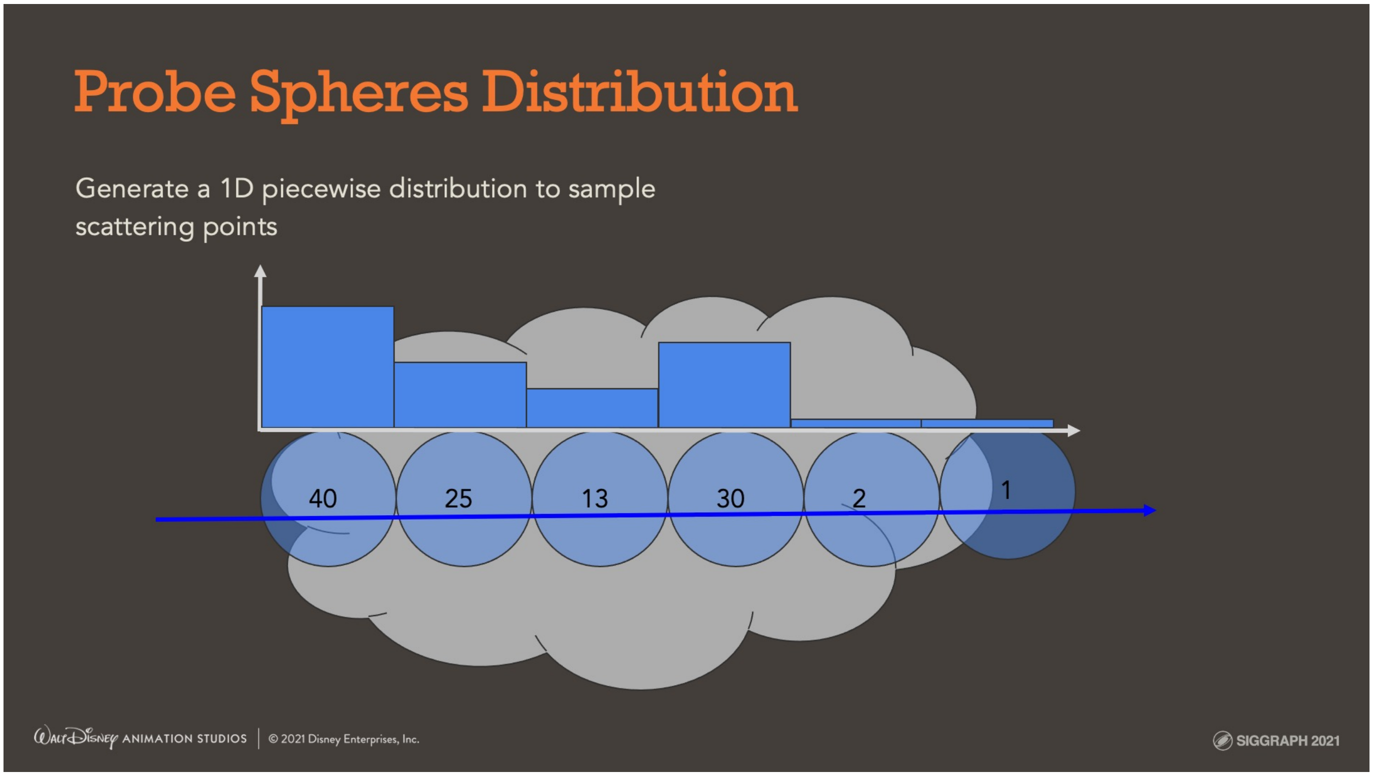

Probe Spheres Distribution

Generate a 1D piecewise distribution to sample scattering points

then use these score values to form a 1D distribution to draw scattering samples. this sampling strategy is only used to generate scattering points for primary-ray direct lighting now to keep the overhead down.

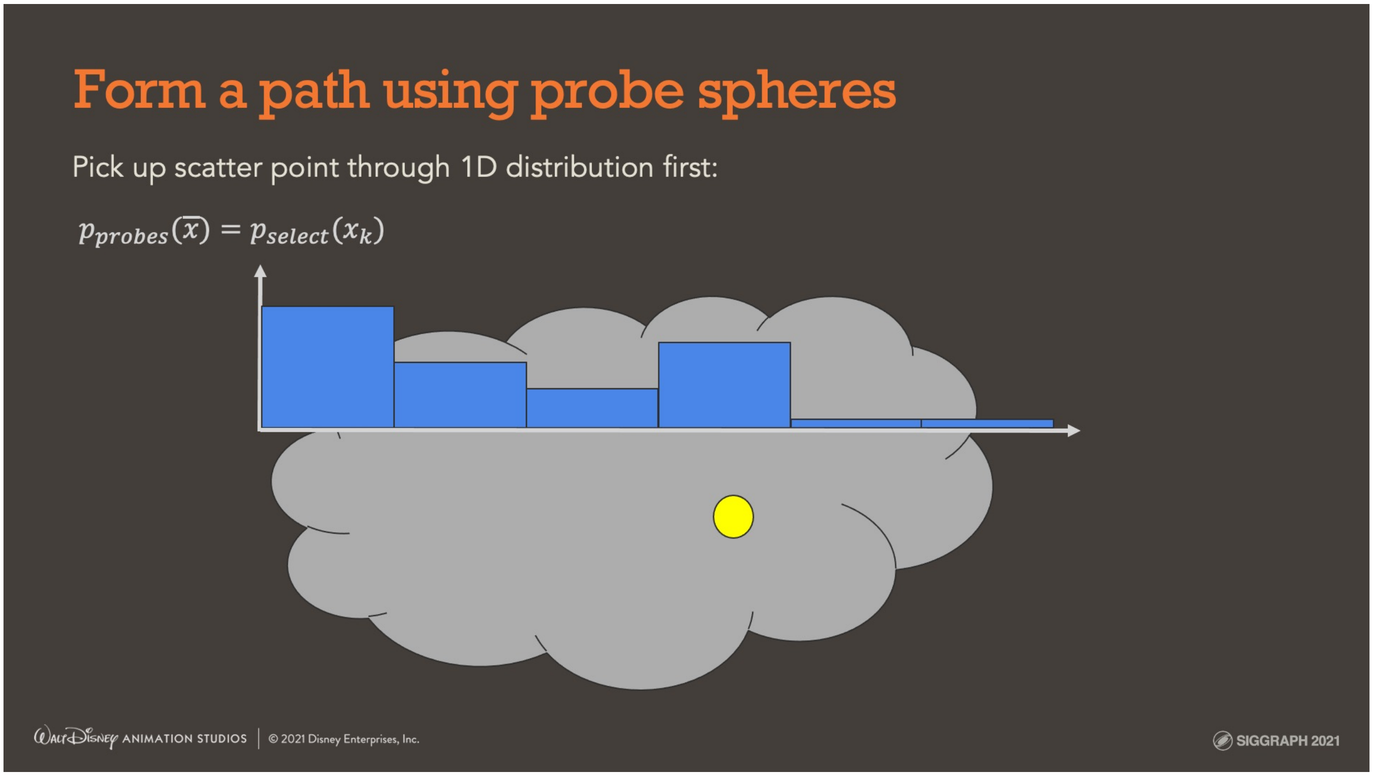

Form a path using probe spheres

Pick up scatter point through 1D distribution first:

$p_{probes}(\overline{x}) = p_{select}(x_k)$

To form the full path that contains the null-collision path vertices for MIS purpose, we use the 1D distribution first to pick up the scattering point

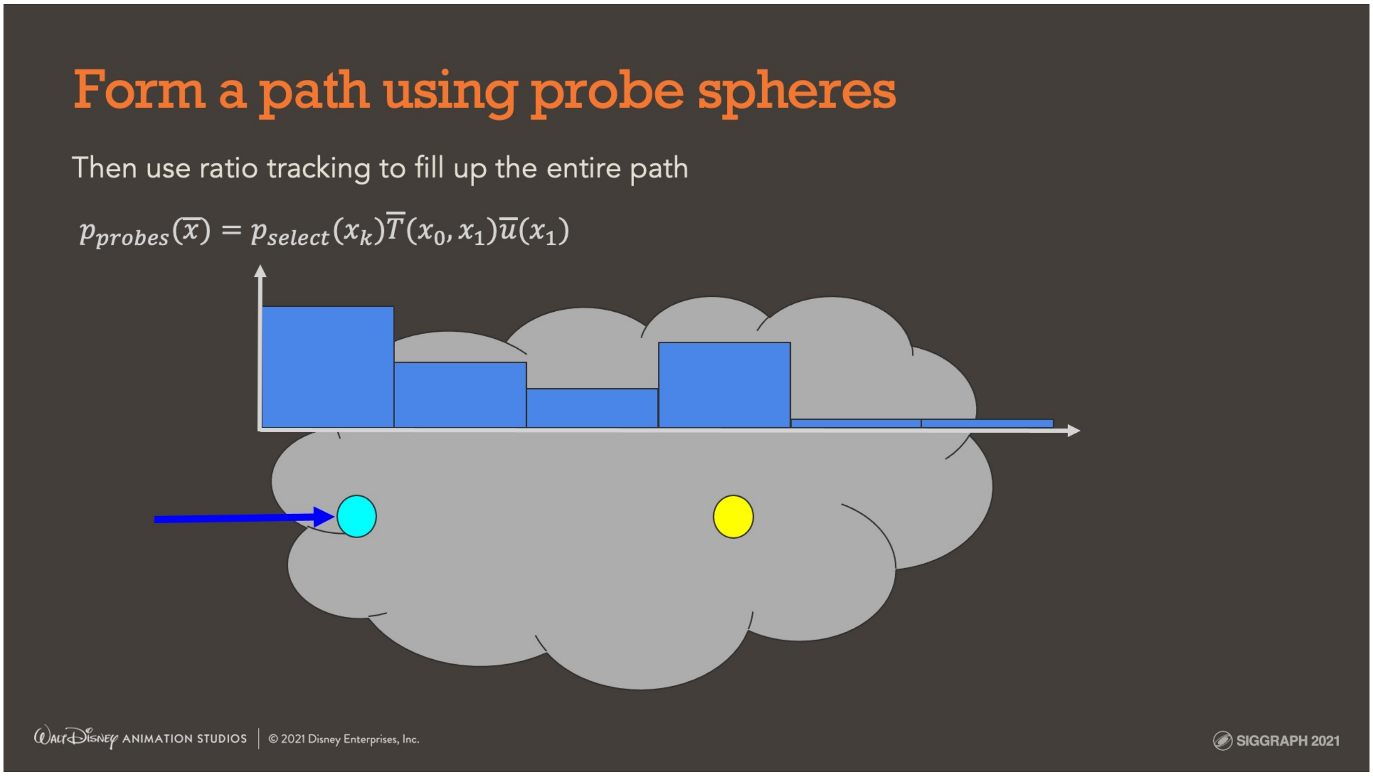

Form a path using probe spheres

Then use ratio tracking to fill up the entire path

$$p_{probes}(\overline{x}) = p_{select}(x_k)\overline{T}(x_0, x_1)\overline{u}(x_1)$$

We then use ratio tracking moving towards the scattering point, and update the path pdf

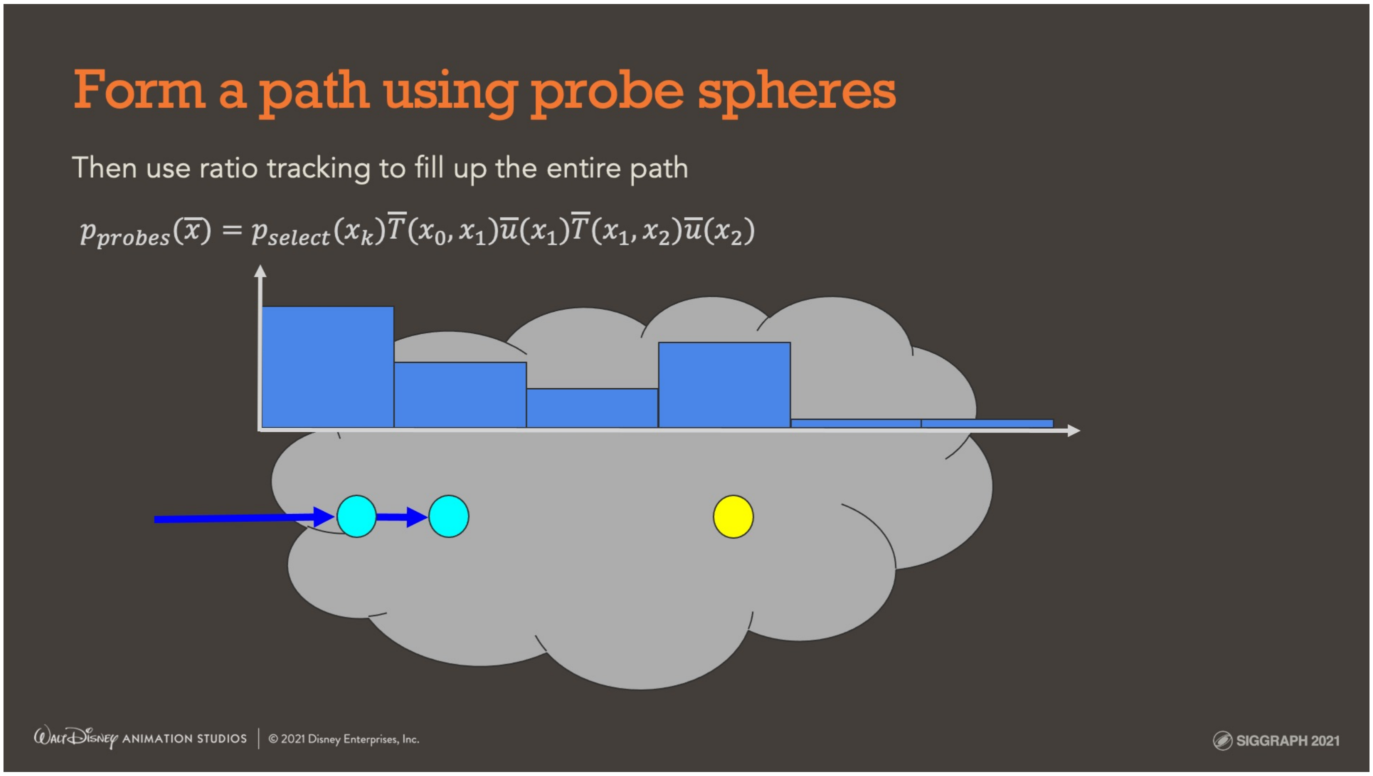

Form a path using probe spheres

Then use ratio tracking to fill up the entire path

$$p_{probes}(\overline{x}) = p_{select}(x_k)\overline{T}(x_0, x_1)\overline{u}(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)$$

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

repeat the distance sampling process...

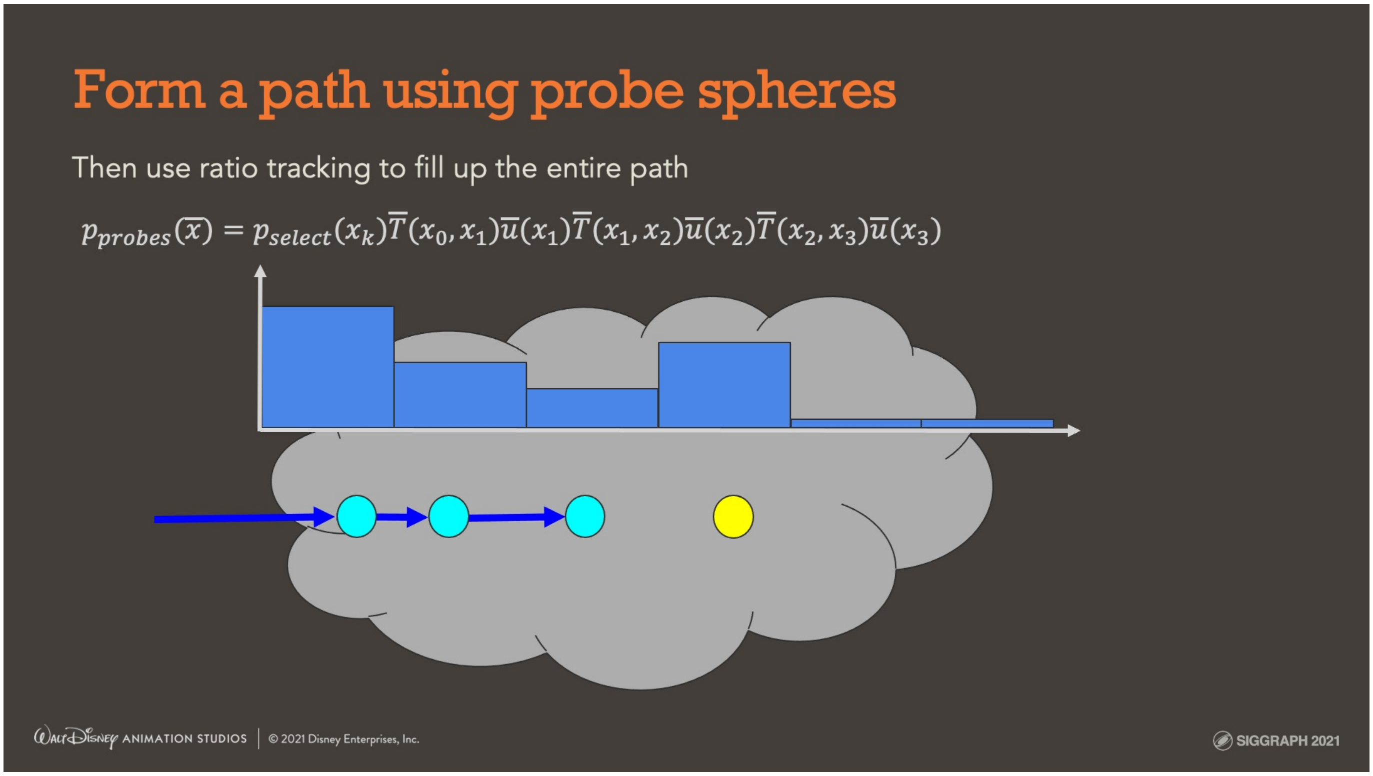

Form a path using probe spheres

Then use ratio tracking to fill up the entire path

$$p_{probes}(\overline{x}) = p_{select}(x_k)\overline{T}(x_0,x_1)\overline{u}(x_1)\overline{T}(x_1,x_2)\overline{u}(x_2)\overline{T}(x_2,x_3)\overline{u}(x_3)$$

WALT DISNEY ANIMATION STUDIOS

SIGGRAPH 2021

...and update the corresponding pdf...

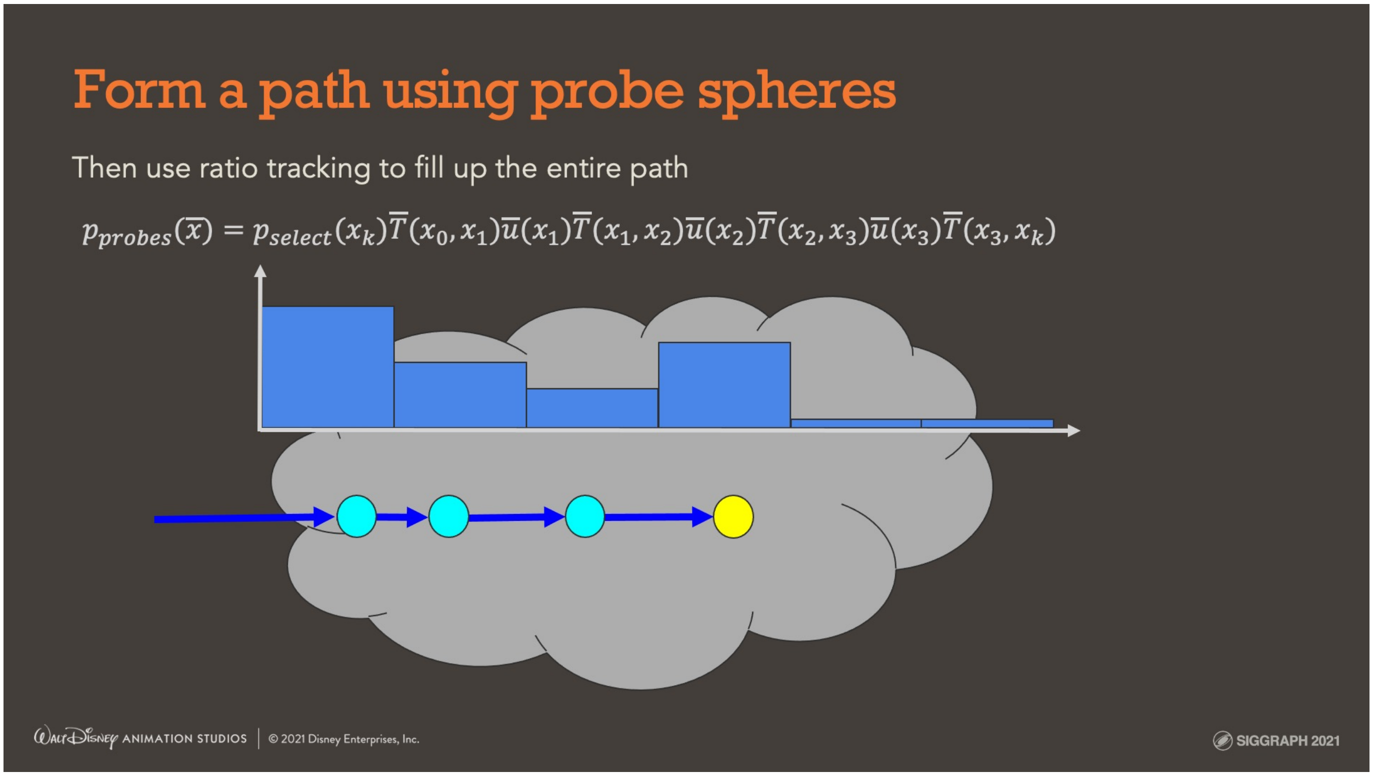

Form a path using probe spheres

Then use ratio tracking to fill up the entire path

$$p_{probes}(\overline{x}) = p_{select}(x_k)\overline{T}(x_0, x_1)\overline{u}(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)\overline{T}(x_2, x_3)\overline{u}(x_3)\overline{T}(x_3, x_k)$$

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

...until distance sampling passes through the scattering point

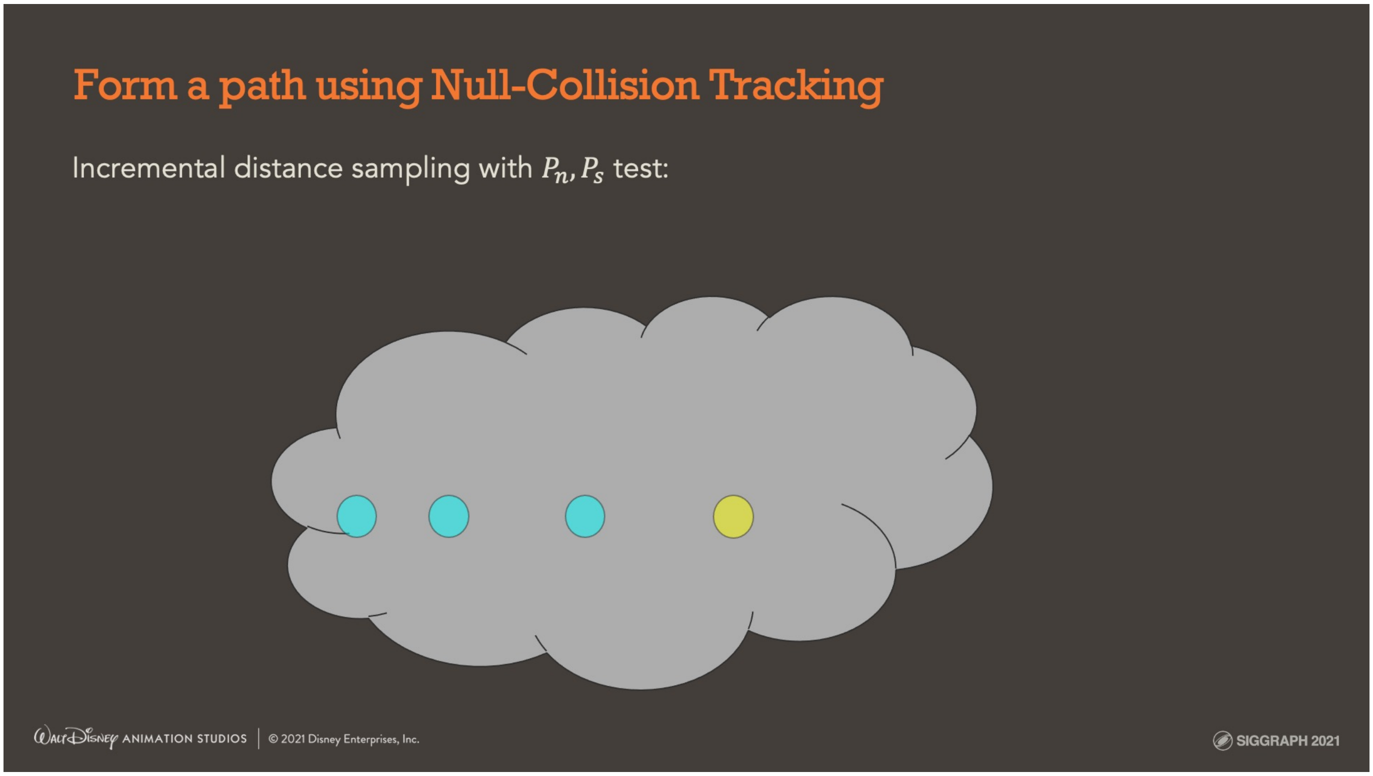

Form a path using Null-Collision Tracking

Incremental distance sampling with P_n, P_s test:

To formulate the same path using null-collision tracking to get the pdf...

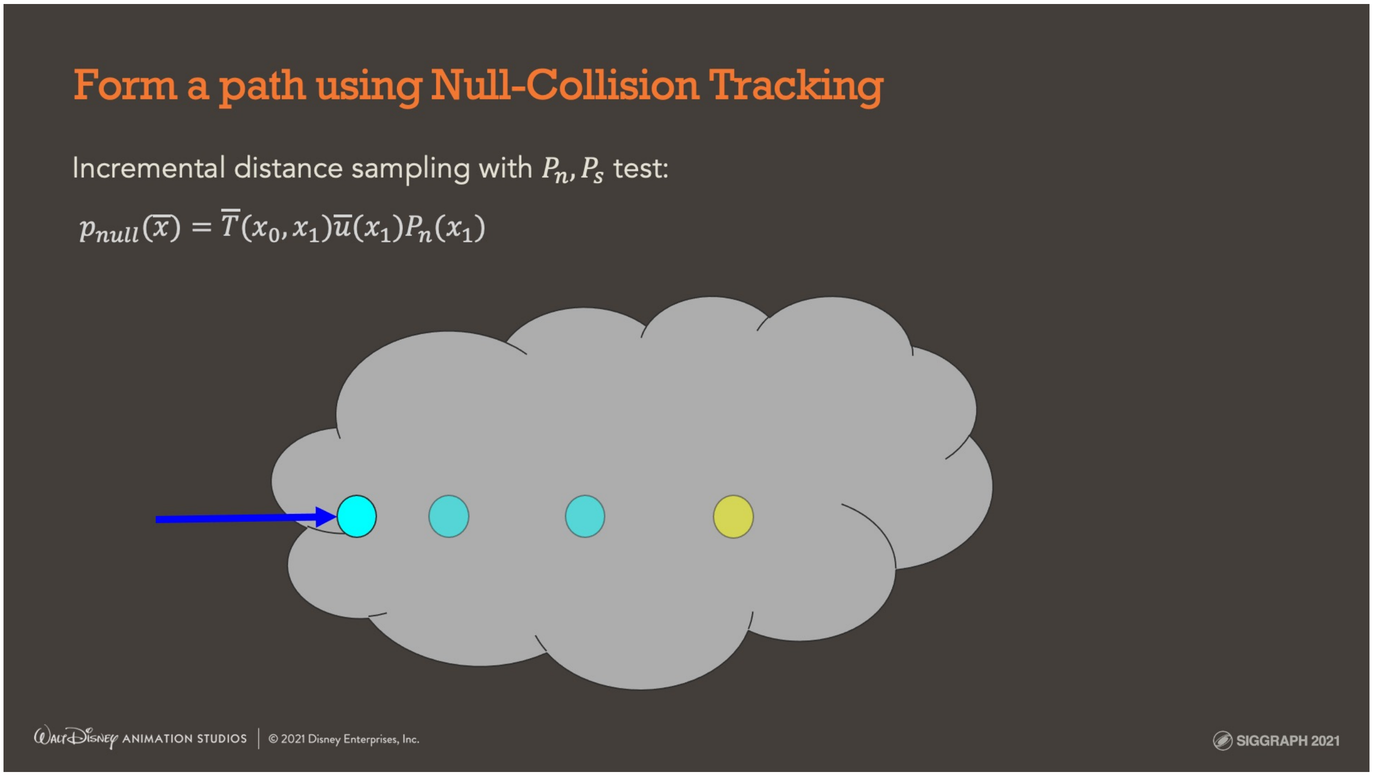

Form a path using Null-Collision Tracking

Incremental distance sampling with P_n, P_s test:

$$p_{null}(\overline{x}) = \overline{T}(x_0, x_1)\overline{u}(x_1)P_n(x_1)$$

...we have the distance and mu_bar to compute T_bar, and we know Pn and Ps based on the tracking algorithm we use.

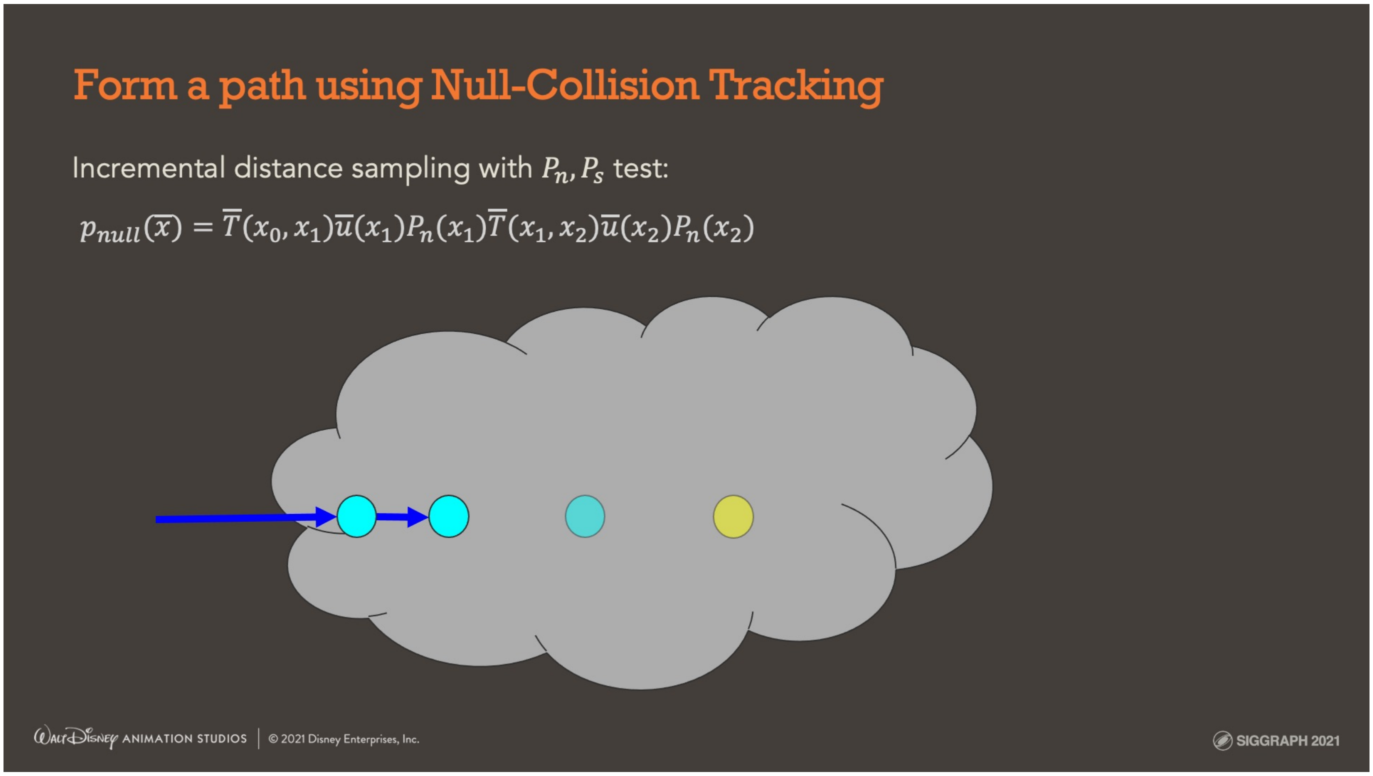

Form a path using Null-Collision Tracking

Incremental distance sampling with P_n, P_s test:

$$p_{null}(\overline{x}) = \overline{T}(x_0, x_1)\overline{u}(x_1)P_n(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)P_n(x_2)$$

For all the vertices before the scattering, we apply pdf Pn...

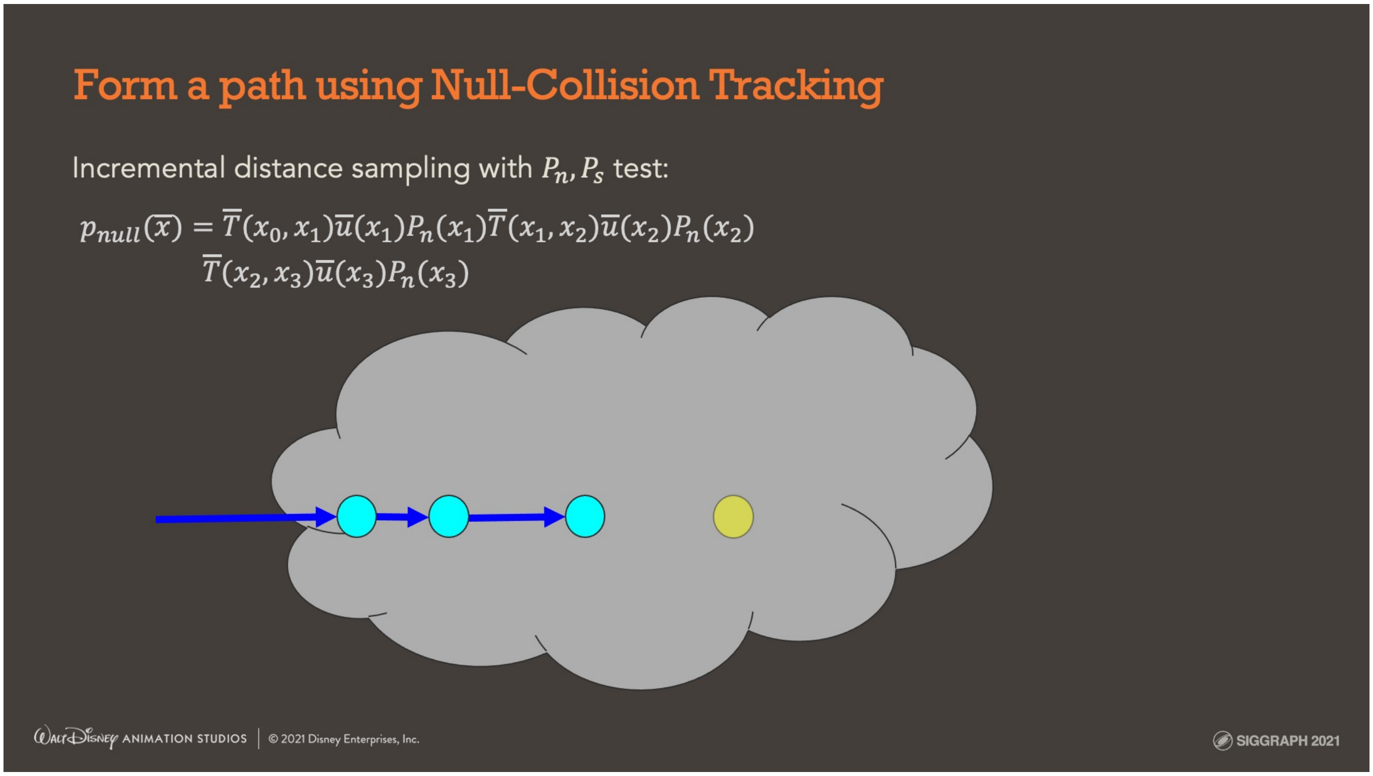

Form a path using Null-Collision Tracking

Incremental distance sampling with P_n, P_s test:

$$p_{null}(\overline{x}) = \overline{T}(x_0, x_1)\overline{u}(x_1)P_n(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)P_n(x_2)$$
$$\overline{T}(x_2, x_3)\overline{u}(x_3)P_n(x_3)$$

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

...repeat the distance sampling process and update the corresponding pdf...

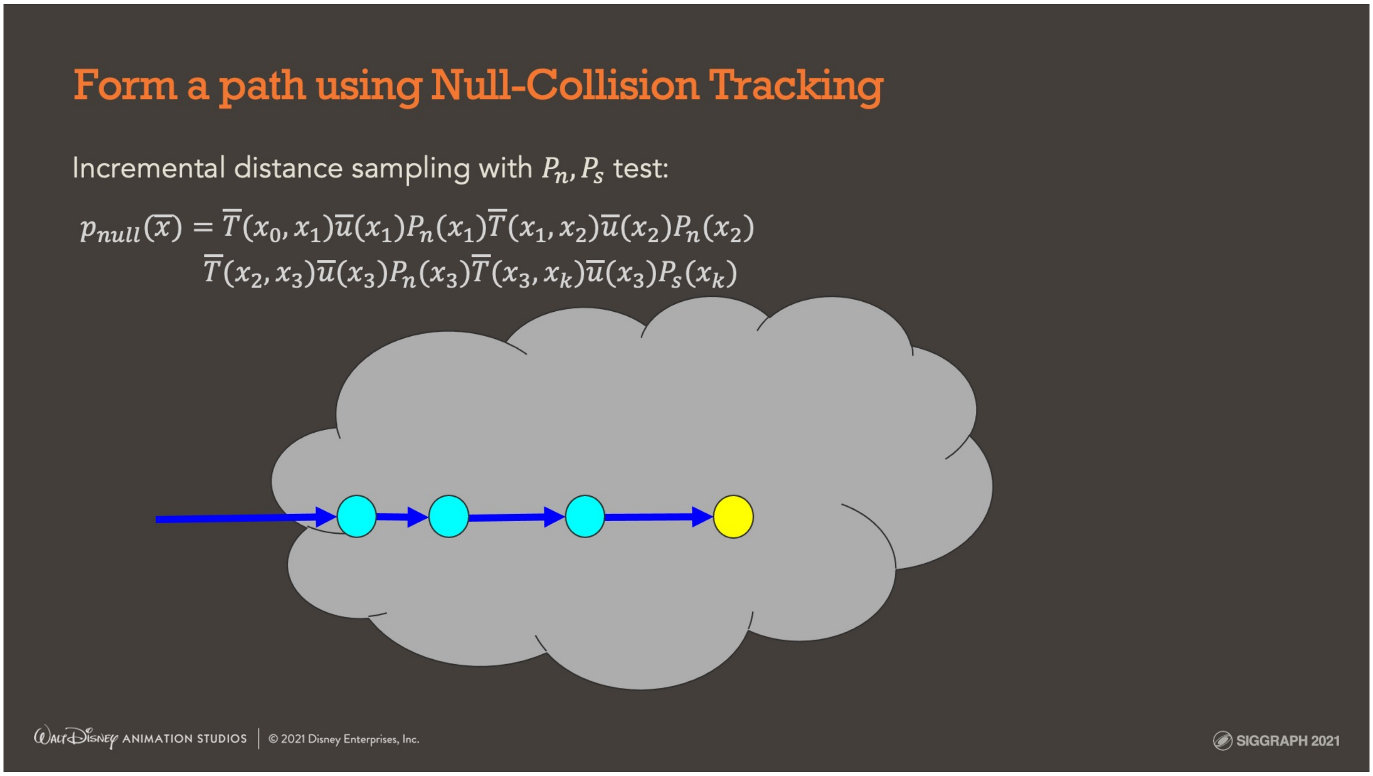

Form a path using Null-Collision Tracking

Incremental distance sampling with P_n, P_s test:

$$p_{null}(\overline{x}) = \overline{T}(x_0, x_1)\overline{u}(x_1)P_n(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)P_n(x_2)$$
$$\overline{T}(x_2, x_3)\overline{u}(x_3)P_n(x_3)\overline{T}(x_3, x_k)\overline{u}(x_3)P_s(x_k)$$

...until we reach the scattering point and apply pdf P_s

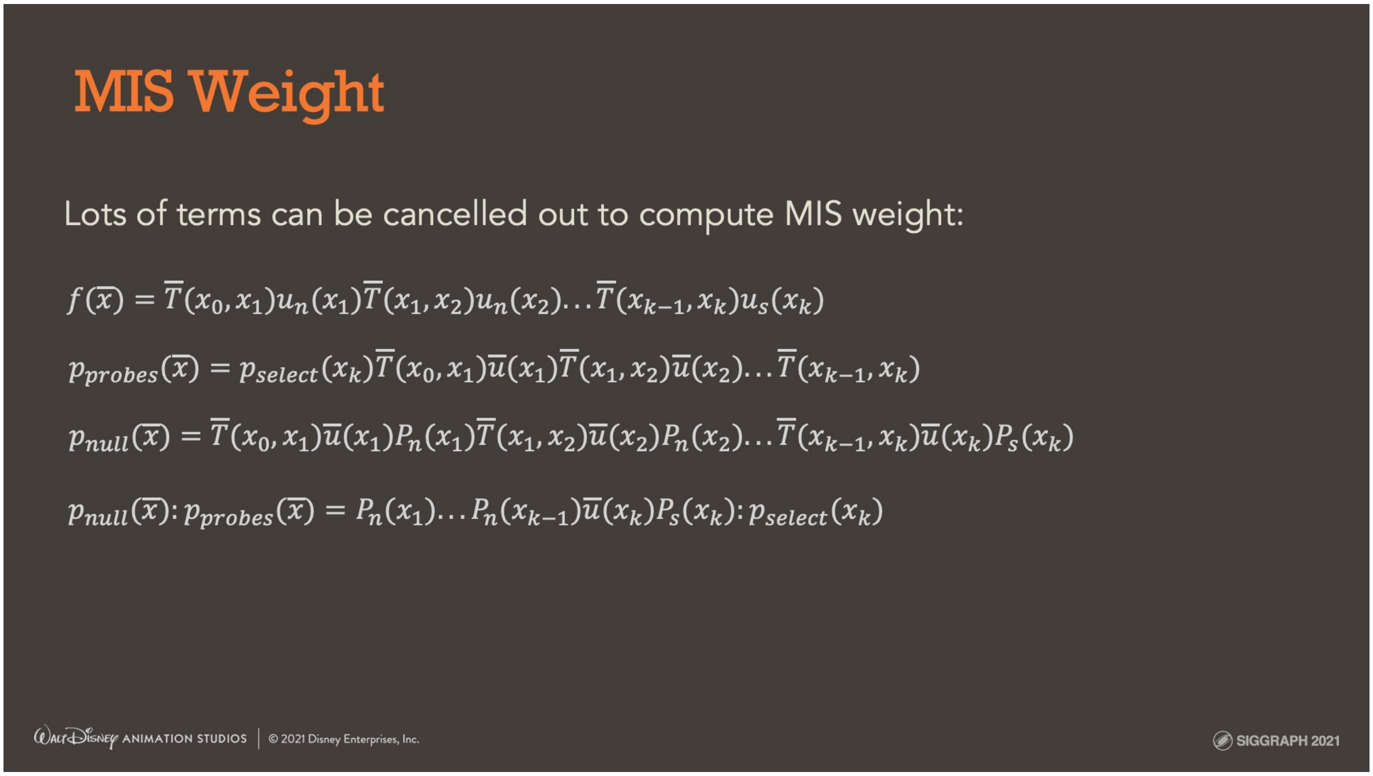

MIS Weight

Lots of terms can be cancelled out to compute MIS weight:

$$f(\overline{x}) = \overline{T}(x_0, x_1)u_n(x_1)\overline{T}(x_1, x_2)u_n(x_2)\ldots\overline{T}(x_{k-1}, x_k)u_s(x_k)$$

$$p_{probes}(\overline{x}) = p_{select}(x_k)\overline{T}(x_0, x_1)\overline{u}(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)\ldots\overline{T}(x_{k-1}, x_k)$$

$$p_{null}(\overline{x}) = \overline{T}(x_0, x_1)\overline{u}(x_1)P_n(x_1)\overline{T}(x_1, x_2)\overline{u}(x_2)P_n(x_2)\ldots\overline{T}(x_{k-1}, x_k)\overline{u}(x_k)P_s(x_k)$$

$$p_{null}(\overline{x}) : p_{probes}(\overline{x}) = P_n(x_1)\ldots P_n(x_{k-1})\overline{u}(x_k)P_s(x_k) : p_{select}(x_k)$$

The contribution function and the path pdf looks quite verbose, but lots of the terms can actually be cancelled. We can then use the final simplified terms to form the mis weight.

Null-Collision Tracking

Heterogeneous volume with low (top) and high (down) extinction coefficients

This is a test render of forward-scattering heterogeneous volume with strong lights embedded within. Null-collision tracking doesn't do well on thinner volumes here.

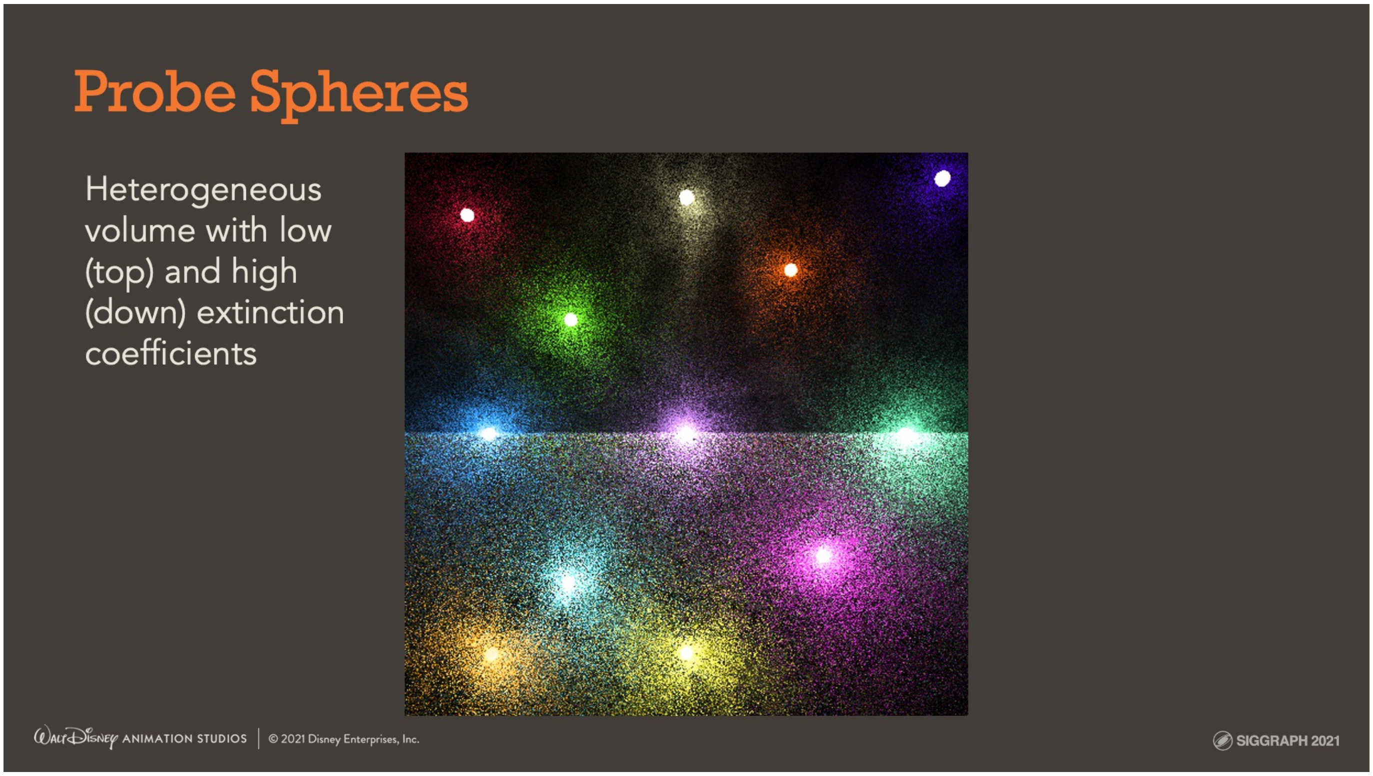

Probe Spheres

Heterogeneous volume with low (top) and high (down) extinction coefficients

...and this is an equal sample count test render using our probe sphere strategy. it does well in the thin volume but behaves worse in the thicker volume.

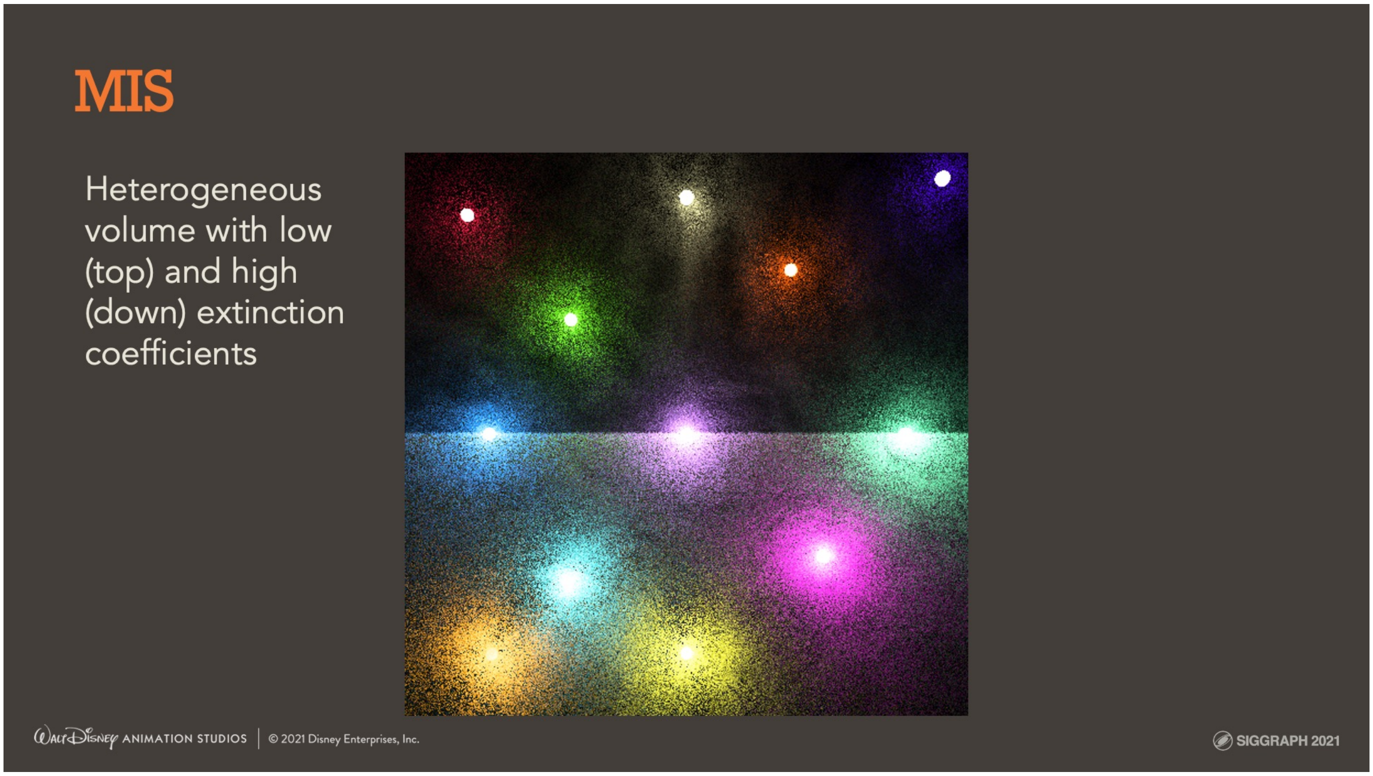

MIS

Heterogeneous volume with low (top) and high (down) extinction coefficients

And this is a test render using MIS to combine the two strategies, and now we have both cases covered in one unified volume integrator.

Optimized emission sampling
Optimized emission+scattering sampling

This is an equal-sample-count production shot test. The upside one is a render using our emission sampling improvements. The downside one is a render where we further layer in our scattering sampling improvements. You can see that the noise is further reduced around the atmospheric area illuminated by the emissive sources.

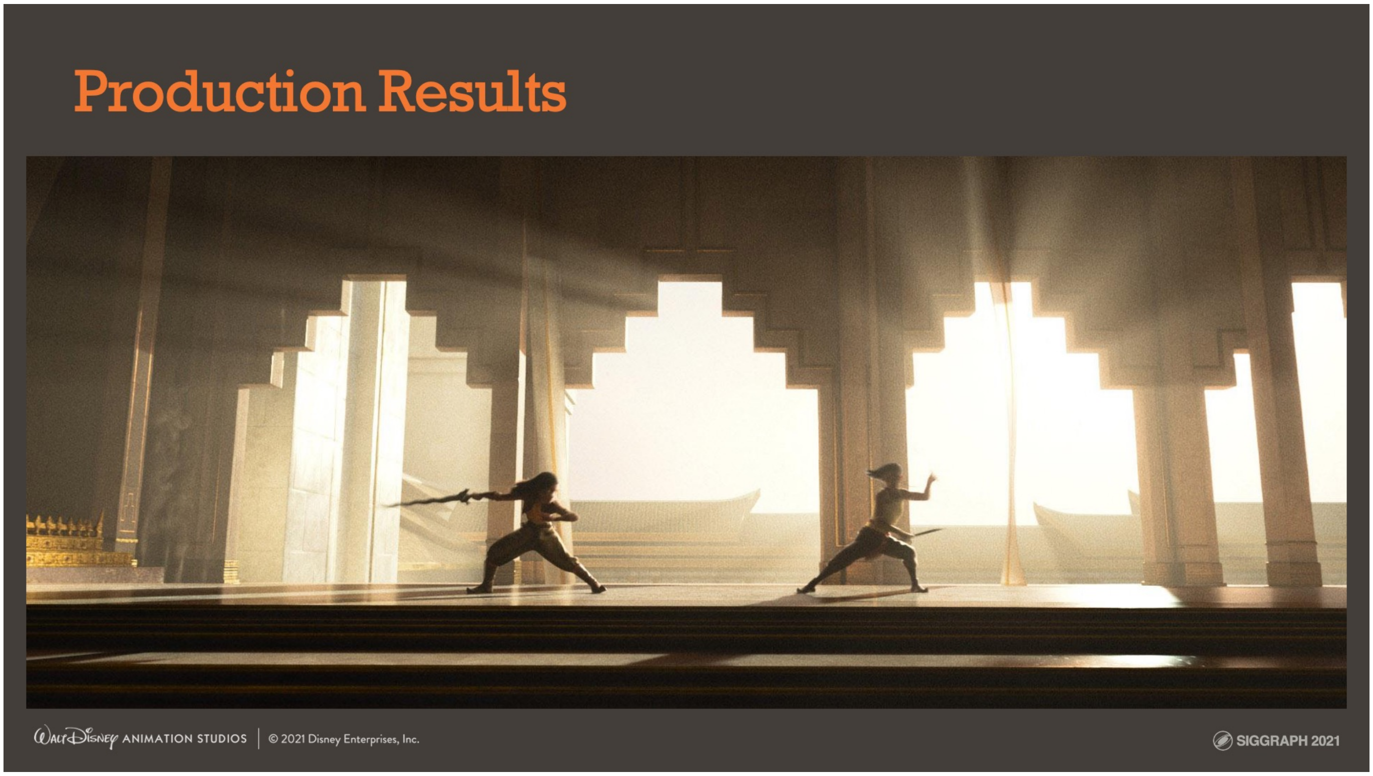

Production Results

Walt Disney ANIMATION STUDIOS

And here are some still frames from the film Raya and the Last Dragon, showcasing the beautiful work our artists put together using our improved volume integrator.

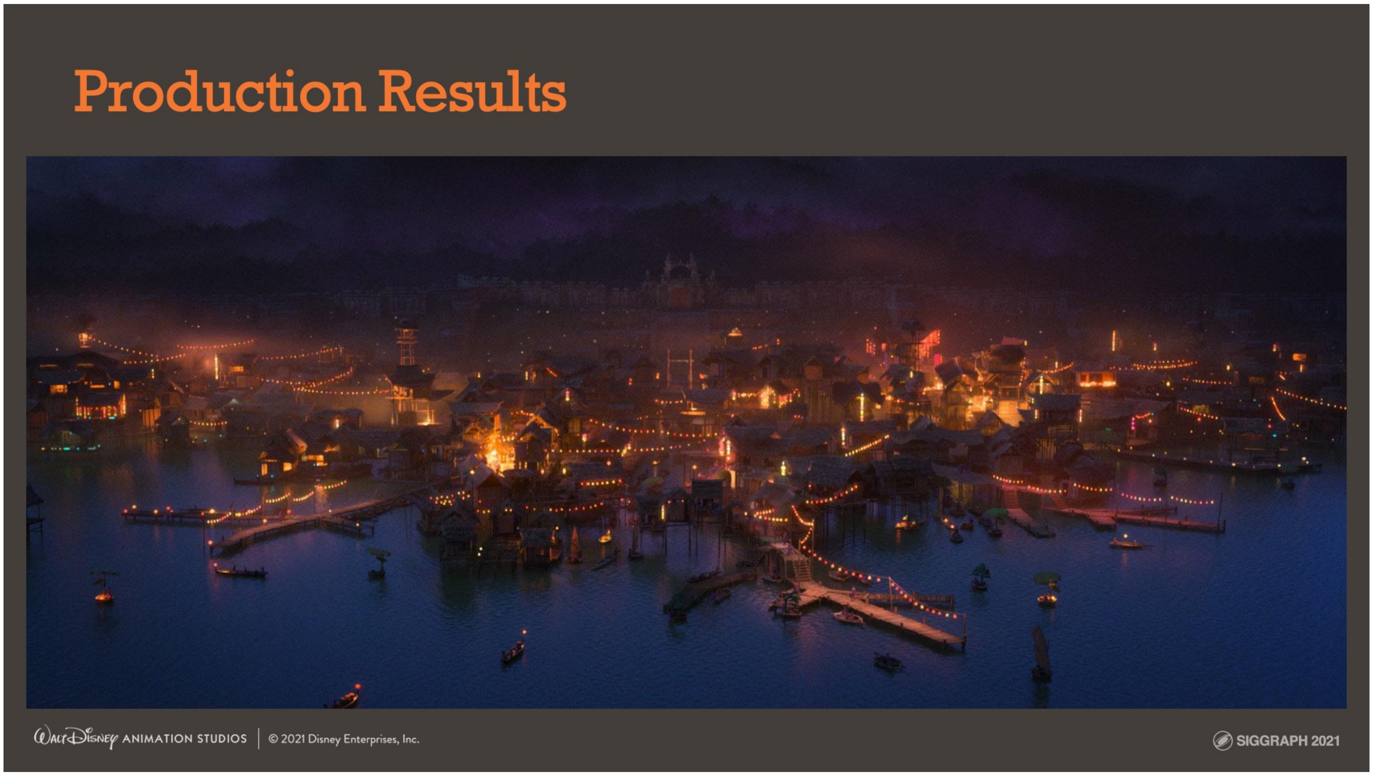

Production Results

WALT DISNEY ANIMATION STUDIOS |

SIGGRAPH 2021

The renderer can handle more combinations of complex lighting and volumetric scenarios now...

Production Results

We can use emissive volumes to produce more realistic and accurate shadow movements

Production Results

...and we can now also render more volumetric effects directly using our renderer

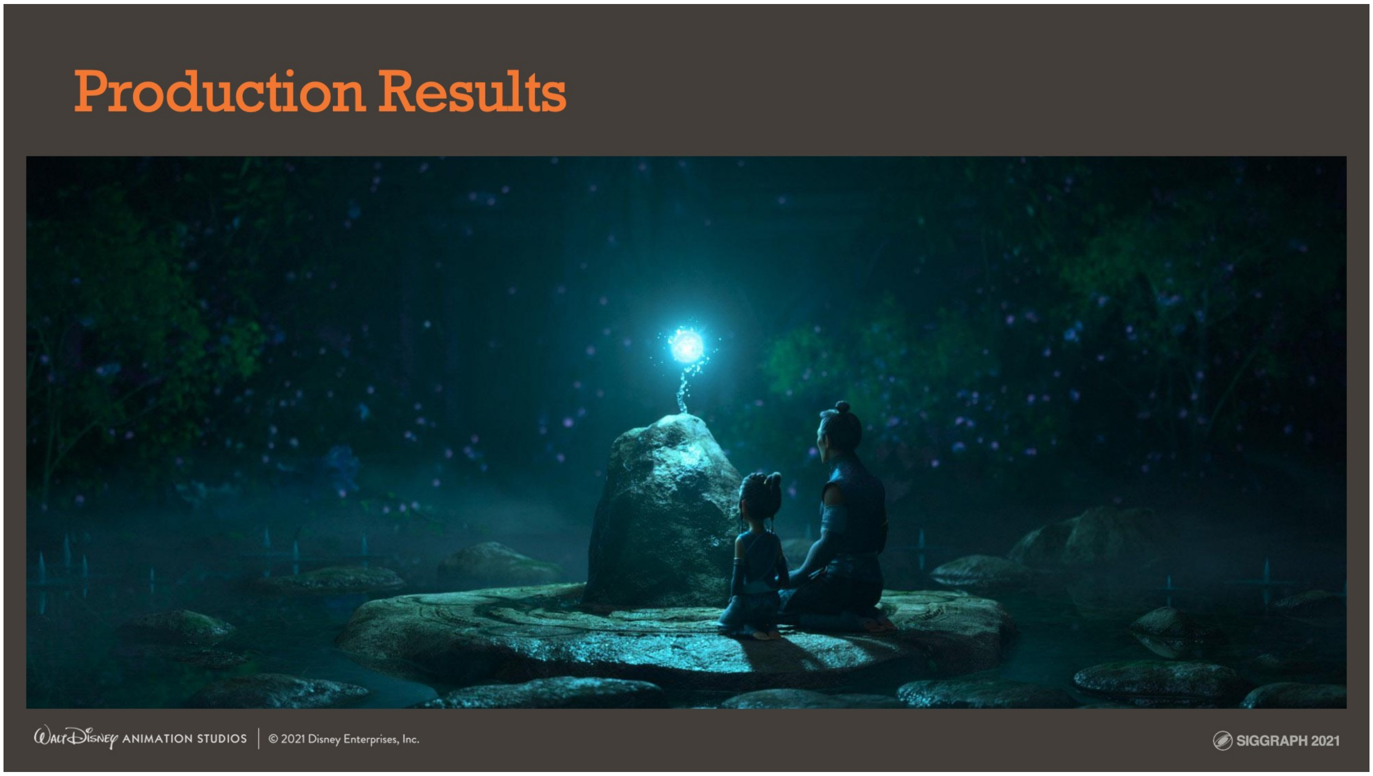

Production Results

...which reduces the complexity of our compositing setups

Production Results

Walt Disney ANIMATION STUDIOS |

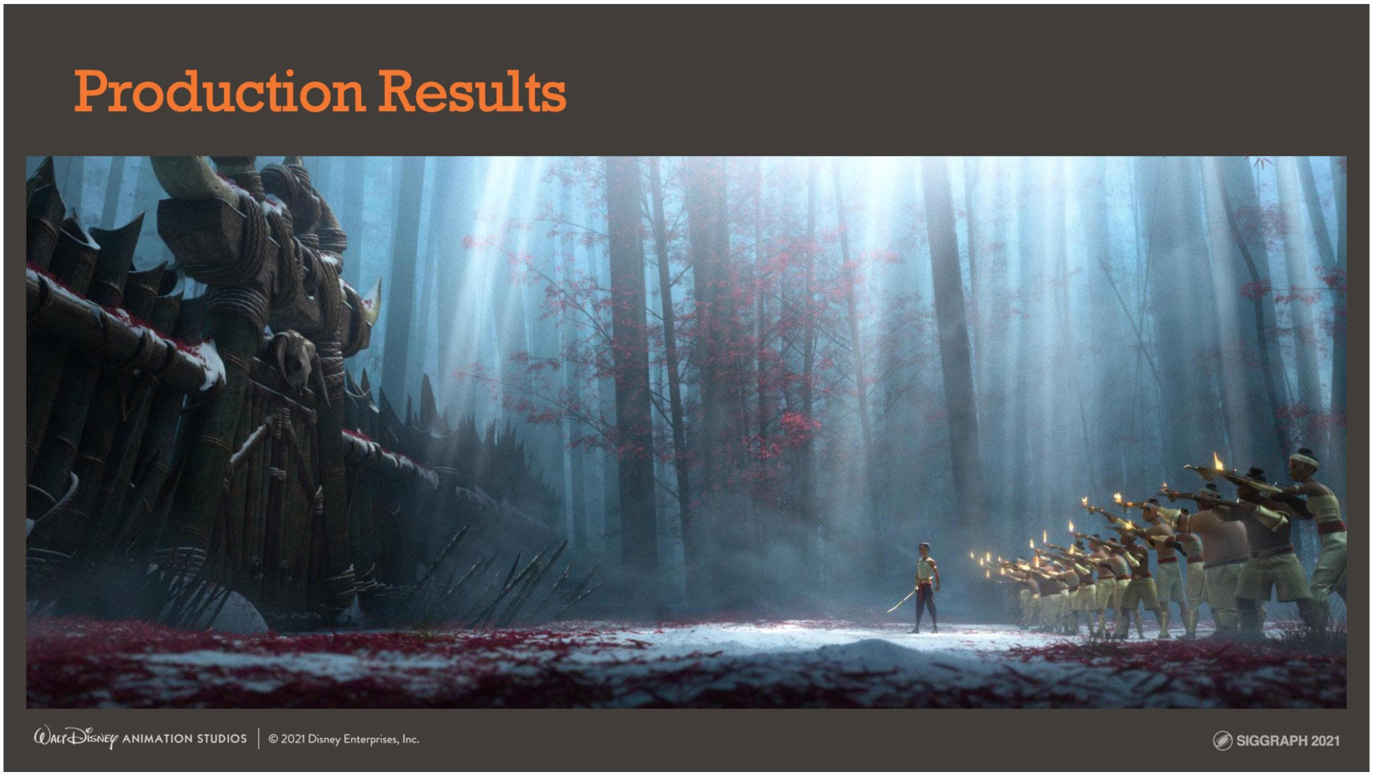

Production Results

Production Results

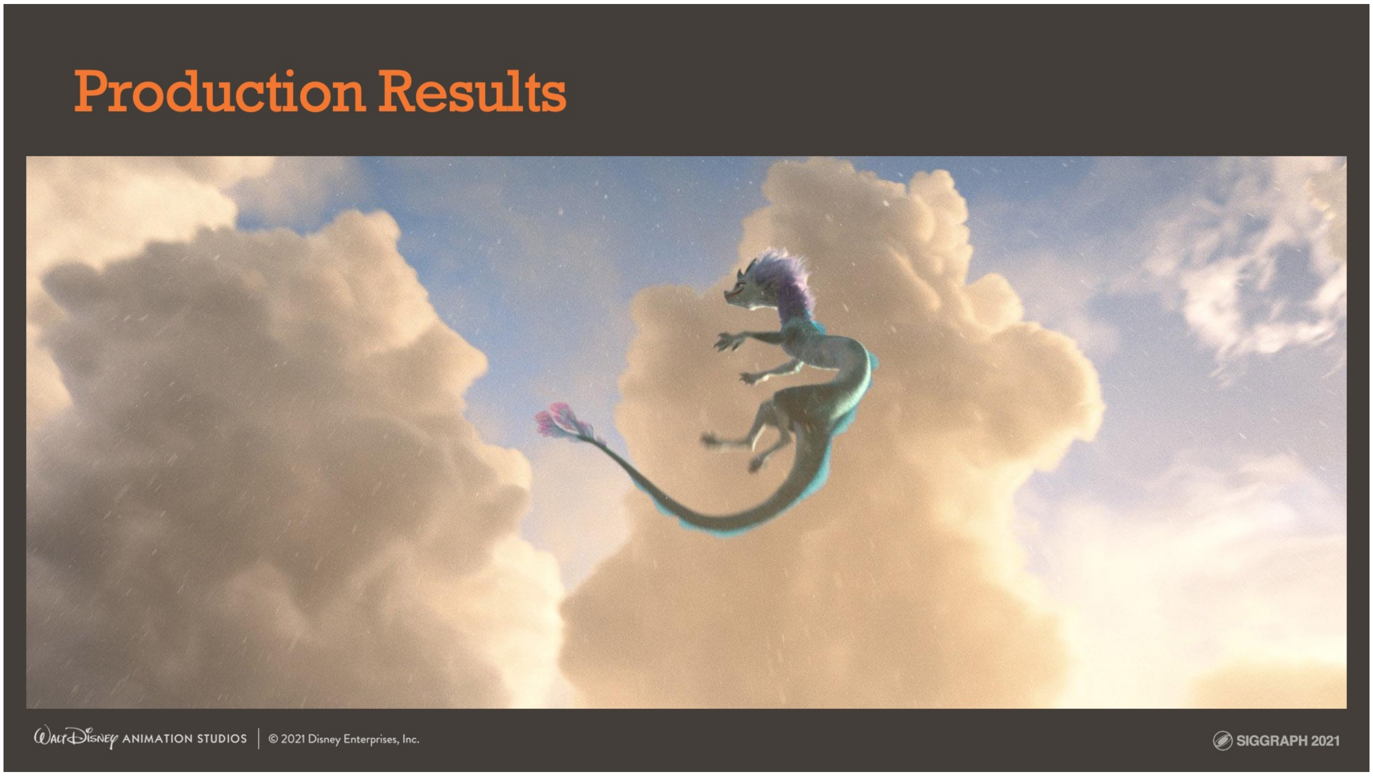

Production Results

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

...and of course we are very happy with our ability to still render white clouds efficiently.

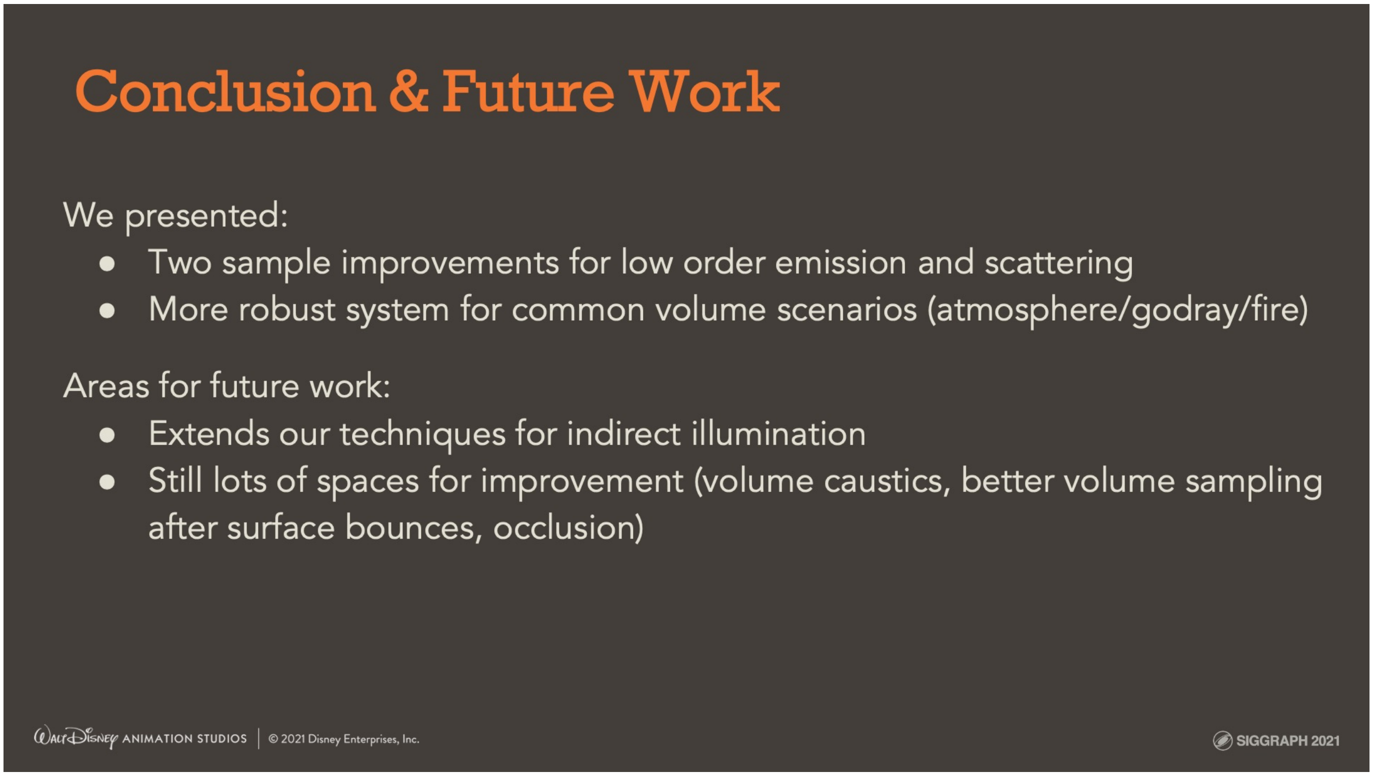

Conclusion & Future Work

We presented:

- Two sample improvements for low order emission and scattering
- More robust system for common volume scenarios (atmosphere/godray/fire)

Areas for future work:

- Extends our techniques for indirect illumination
- Still lots of spaces for improvement (volume caustics, better volume sampling after surface bounces, occlusion)

Walt Disney ANIMATION STUDIOS |

SIGGRAPH 2021

To summarize: We improved our volume emission and scattering sampling through extending the null-collision formulation. These improvements make our rendering system more robust handling various common volume scenarios, without losing its strength for high order scattering.

In the future, we would like to have scattering sample technique that is not limited to primary ray direct lighting. Better sampling for volume caustic and complex occlusion without losing the performance would be always desired.

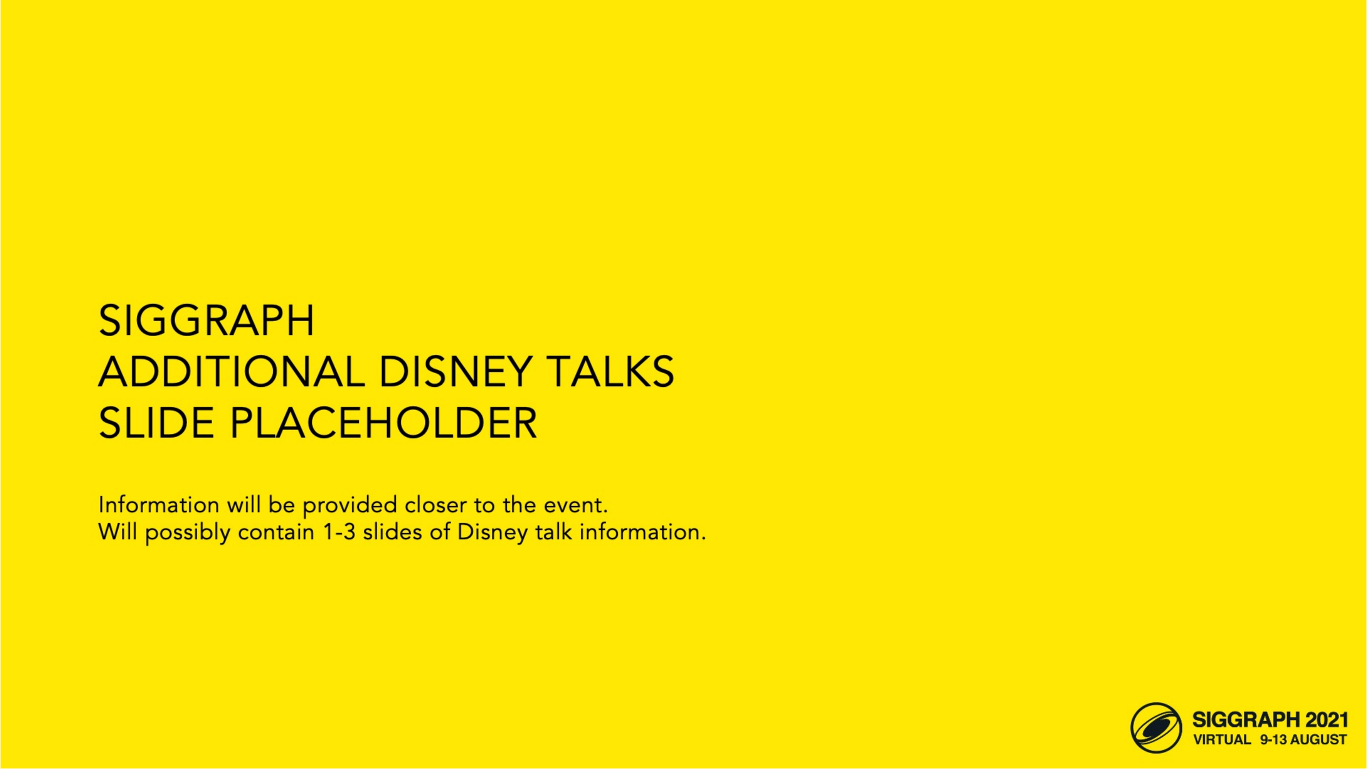

SIGGRAPH ADDITIONAL DISNEY TALKS SLIDE PLACEHOLDER

Information will be provided closer to the event.
Will possibly contain 1-3 slides of Disney talk information.

Placeholder slide where we will place information related to other Walt Disney Animation Studios presentations taking place around Siggraph 2021

SIGGRAPH 2021

THANK YOU

FOR JOINING US!

© 2021 SIGGRAPH. ALL RIGHTS RESERVED.

THE PREMIER **CONFERENCE** & **EXHIBITION** IN
COMPUTER GRAPHICS & INTERACTIVE TECHNIQUES

We thank you for joining the talk and look forward to discuss with you later. See you in the rest of SIGGRAPH