

Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes – Supplemental Material

PETER KUTZ, Walt Disney Animation Studios

RALF HABEL, Walt Disney Animation Studios

YINING KARL LI, Walt Disney Animation Studios

JAN NOVÁK, Disney Research

In this document, we present further material for our paper, “Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes”:

- in Section 1.1, we present an extended proof of the main mathematical mechanism our analog decomposition tracking is built upon,
- in Section 1.2, we prove that the maximum-based probabilities cap the local collision weight,
- in Section 1.3, we prove our claim that geometric growth of the throughput can be avoided when using spectral tracking,
- in Section 2.1, we provide a number of plots and empirical results generated during our investigation of and experiments with decomposition tracking, and we discuss the meaning of these results,
- in Section 2.2, we provide an additional spectral tracking example using a special absorbing volume,
- in Section 2.3, we show the effects of non-bounding free-path-sampling coefficients on renders, and
- in Section 2.4 we compare the two history-aware spectral-tracking probability schemes on a volume with an albedo of 1.
- in Section 2.5 we provide additional comparisons of delta and decomposition tracking performance for various octree depths.

1 PROOFS AND DERIVATIONS

1.1 Extended Proof of Analog Decomposition Tracking

We present a proof that the minimum of the distance samples drawn from two components of a volume is distributed identically to the distance samples drawn directly from their sum, the original volume. This proof is essentially an expanded version of the proof of Theorem 1 in Section 4.1 of the main paper, and it uses the same notation. Specifically, this proof shows that a random variable C representing the distance samples from the original volume has the same CDF as the minimum D of two random variables A and B representing the distance samples from the component volumes.

PROOF. We begin by deriving the CDF of the minimum of two independent random variables:

$$\begin{aligned}
 F_D(t) &= P(D \leq t) \\
 &= P(\min(A, B) \leq t) \\
 &= P(A \leq t \vee B \leq t) \\
 &= 1 - P(A > t \wedge B > t) \\
 &= 1 - P(A > t)P(B > t) \\
 &= 1 - (1 - P(A \leq t))(1 - P(B \leq t)) \\
 &= 1 - (1 - P(A \leq t) - P(B \leq t) + P(A \leq t)P(B \leq t)) \\
 &= P(A \leq t) + P(B \leq t) - P(A \leq t)P(B \leq t) \\
 &= F_A(t) + F_B(t) - F_A(t)F_B(t). \tag{1}
 \end{aligned}$$

Next, we incorporate the distance CDFs and simplify:

$$\begin{aligned}
 F_D(t) &= [1 - T_A(t)] + [1 - T_B(t)] - [1 - T_A(t)][1 - T_B(t)] \\
 &= 2 - T_A(t) - T_B(t) - [1 - T_A(t) - T_B(t) + T_A(t)T_B(t)] \\
 &= 2 - T_A(t) - T_B(t) - 1 + T_A(t) + T_B(t) - T_A(t)T_B(t) \\
 &= 1 - T_A(t)T_B(t). \tag{2}
 \end{aligned}$$

Finally, we use the definition of $T_X(t)$ and simplify:

$$\begin{aligned}
 F_D(t) &= 1 - \exp\left(-\int_0^t \mu_A(\mathbf{x}_s)ds\right) \exp\left(-\int_0^t \mu_B(\mathbf{x}_s)ds\right) \\
 &= 1 - \exp\left(-\int_0^t \mu_A(\mathbf{x}_s)ds + -\int_0^t \mu_B(\mathbf{x}_s)ds\right) \\
 &= 1 - \exp\left(-\int_0^t \mu_A(\mathbf{x}_s) + \mu_B(\mathbf{x}_s)ds\right) \\
 &= 1 - \exp\left(-\int_0^t \mu_C(\mathbf{x}_s)ds\right) \\
 &= F_C(t). \tag{3}
 \end{aligned}$$

From the above, we can see that $F_D(t) = F_C(t)$. $\square$

1.2 Maximum Collision Weight of Spectral Tracking

In this section we show that, when using the maximum-based probability scheme for spectral tracking introduced in Section 5.1.1 of the main paper, the maximum value of the local collision weight over all of the channels never exceeds the number of collision types, i.e. 3 for media with absorptive, scattering, and null-collision events. It could be shown analogously that the maximum value never exceeds 2 if the medium is non-emissive.

THEOREM 1. *Let $\mu_\star : \star \in \{a, s, n\}$ be non-negative collision coefficients, where a , s , and n stand for absorption, scattering, and null-collision, respectively. If $P_\star$ is defined using the (non-history-aware) maximum-based scheme described in Section 5.1 of the main paper, then*

$$\max_{\lambda} (w_\star(\mathbf{x}, \lambda)) \leq 3. \tag{4}$$

PROOF. To reason about the left side of Equation 4, we use the shorthand

$$\overline{w} = \max_{\lambda} (w_\star(\mathbf{x}, \lambda)). \tag{5}$$

First, we expand $w_\star$ using its definition:

$$\overline{w} = \max_{\lambda} \left(\frac{\mu_\star(\mathbf{x}, \lambda)}{\bar{\mu}(\mathbf{x})P_\star(\mathbf{x})} \right), \tag{6}$$

and similarly for $P_\star$:

$$\begin{aligned}\bar{w} &= \max_{\lambda} \left(\frac{\mu_\star(\mathbf{x}, \lambda)}{\frac{\max_{\lambda_g}(\mu_\star(\mathbf{x}, \lambda_g))}{\sum_{i \in \{a, s, n\}} \max_{\lambda_g}(\mu_i(\mathbf{x}, \lambda_g))}} \right) \\ &= \max_{\lambda} \left(\frac{\mu_\star(\mathbf{x}, \lambda)}{\bar{\mu}(\mathbf{x})} \frac{\sum_{i \in \{a, s, n\}} \max_{\lambda_g}(\mu_i(\mathbf{x}, \lambda_g))}{\max_{\lambda_g}(\mu_\star(\mathbf{x}, \lambda_g))} \right).\end{aligned}\quad (7)$$

Rearranging terms yields

$$\bar{w} = \max_{\lambda} \left(\frac{\mu_\star(\mathbf{x}, \lambda)}{\max_{\lambda_g}(\mu_\star(\mathbf{x}, \lambda_g))} \frac{\sum_{i \in \{a, s, n\}} \max_{\lambda_g}(\mu_i(\mathbf{x}, \lambda_g))}{\bar{\mu}(\mathbf{x})} \right).\quad (8)$$

Our assumption that all coefficients are non-negative implies that the free-path-sampling coefficient is a majorant. This means that the maximum value of any wavelength of any collision coefficient is lower or equal to the free-path-sampling coefficient. Hence

$$\bar{w} \leq \max_{\lambda} \left(\frac{\mu_\star(\mathbf{x}, \lambda)}{\max_{\lambda_g}(\mu_\star(\mathbf{x}, \lambda_g))} \frac{\sum_{i \in \{a, s, n\}} \bar{\mu}(\mathbf{x})}{\bar{\mu}(\mathbf{x})} \right),\quad (9)$$

which, given that all coefficients are non-negative, simplifies to

$$\bar{w} \leq 3 \max_{\lambda} \left(\frac{\mu_\star(\mathbf{x}, \lambda)}{\max_{\lambda_g}(\mu_\star(\mathbf{x}, \lambda_g))} \right).\quad (10)$$

The maximum of a vector of values divided by the maximum value of that vector equals to 1, hence

$$\bar{w} \leq 3\quad (11)$$

concluding the proof. $\square$

1.3 Path Throughput of Spectral Tracking

In this section, we prove by mathematical induction that, when using the history-aware average-based probability scheme for spectral tracking introduced in Section 5.1.2 of the main paper, the sum of the throughputs over all of the channels is always equal to the number of channels N_λ . This property also implies that the average of the throughputs over all of the channels is always equal to 1, and that the maximum throughput for any one channel is N_λ .

THEOREM 2. *Let $\mu_\star : \star \in \{a, s, n\}$ be non-negative collision coefficients, where a , s , and n stand for absorption, scattering, and null-collision, respectively. If $P_\star$ is defined using the history-aware average-based scheme described in Section 5.1.2 of the main paper, then the following statement $S(n)$:*

$$\sum_{h=1}^{N_\lambda} w(\mathbf{X}_h, \lambda_h) = N_\lambda,\quad (12)$$

holds for all natural numbers $n \in \mathbb{N}^0$.

PROOF. We prove the above theorem by mathematical induction.

Base case: Because

$$w(\mathbf{X}_0, \lambda) \equiv 1,\quad (13)$$

$S(0)$ holds:

$$\sum_{h=1}^{N_\lambda} w(\mathbf{X}_0, \lambda_h) = \sum_{h=1}^{N_\lambda} 1 = N_\lambda.\quad (14)$$

Inductive step: Assuming that $S(j)$ holds, we show that $S(j+1)$ also holds, that is:

$$\sum_{h=1}^{N_\lambda} w(\mathbf{X}_{j+1}, \lambda_h) = N_\lambda.\quad (15)$$

We first express the path throughput $w(\mathbf{X}_{j+1}, \lambda_h)$ as a product of throughput of path prefix $\mathbf{X}_j$ and the local collision weight at $\mathbf{x}_{j+1}$. We use $\tilde{w}$ as a shorthand for the left-hand side of Equation (12) and $\mathbf{x}_+$ as a shorthand for $\mathbf{x}_{j+1}$ for clarity:

$$\tilde{w} = \sum_{h=1}^{N_\lambda} [w(\mathbf{X}_j, \lambda_h) w_\star(\mathbf{x}_+, \lambda_h)].\quad (16)$$

The definition of $w_\star$ can be substituted in to obtain

$$\tilde{w} = \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+) P_\star(\mathbf{x}_+)} \right],\quad (17)$$

and then the definition of $P_\star$ can be substituted in to obtain

$$\begin{aligned}\tilde{w} &= \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+) \frac{\text{avg}(\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g))}{\sum_{i \in \{a, s, n\}} \left[\frac{\text{avg}(\mu_i(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g))}{\lambda_g} \right]}} \right] \\ &= \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+)} \frac{\sum_{i \in \{a, s, n\}} \left[\frac{\text{avg}(\mu_i(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g))}{\lambda_g} \right]}{\text{avg}(\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g))} \right].\end{aligned}\quad (18)$$

Next, the avg function can be replaced with its definition, yielding

$$\tilde{w} = \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+)} \frac{\sum_{i \in \{a, s, n\}} \left[\frac{1}{N_\lambda} \sum_{g=1}^{N_\lambda} [\mu_i(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)] \right]}{\frac{1}{N_\lambda} \sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \right].\quad (19)$$

We then extract $1/N_\lambda$, which does not depend on i or g . The $1/N_\lambda$ terms can then be canceled. Next, we can switch the nested summations in preparation for later steps:

$$\tilde{w} = \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+)} \frac{\sum_{g=1}^{N_\lambda} \left[\sum_{i \in \{a, s, n\}} [\mu_i(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)] \right]}{\sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \right]\quad (20)$$

Because we use one $\bar{\mu}$ for all λ , we know that

$$\sum_{i \in \{a, s, n\}} \mu_i(\mathbf{x}, \lambda) \equiv \mu_a(\mathbf{x}, \lambda) + \mu_s(\mathbf{x}, \lambda) + \mu_n(\mathbf{x}, \lambda) \equiv \bar{\mu}(\mathbf{x}),\quad (21)$$

and we know that the path throughput $w(\mathbf{X}_j, \lambda_g)$ does not depend on i , so the inner sum can be eliminated. The $\bar{\mu}$ can then be extracted from the remaining sum; then we cancel the $\bar{\mu}$ terms:

$$\begin{aligned}\tilde{w} &= \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+)} \frac{\sum_{g=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_g) \left(\sum_{i \in \{a, s, n\}} \mu_i(\mathbf{x}_+, \lambda_g) \right) \right]}{\sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \right] \\ &= \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \frac{\mu_\star(\mathbf{x}_+, \lambda_h)}{\bar{\mu}(\mathbf{x}_+)} \frac{\bar{\mu}(\mathbf{x}_+) \sum_{g=1}^{N_\lambda} w(\mathbf{X}_j, \lambda_g)}{\sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \right].\end{aligned}\quad (22)$$

Using the induction hypothesis that $S(j)$ holds, the remaining sum in the numerator can be replaced with N_λ , yielding

$$\tilde{w} = \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \mu_\star(\mathbf{x}_+, \lambda_h) \frac{N_\lambda}{\sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \right].\quad (23)$$

Since N_λ and the inner summation in the denominator are both independent of h , we can extract them, producing a fraction where we can cancel the numerator and denominator to produce N_λ :

$$\begin{aligned}\tilde{w} &= \sum_{h=1}^{N_\lambda} \left[w(\mathbf{X}_j, \lambda_h) \mu_\star(\mathbf{x}_+, \lambda_h) \right] \frac{N_\lambda}{\sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \\ &= N_\lambda \frac{\sum_{h=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_h) w(\mathbf{X}_j, \lambda_h)]}{\sum_{g=1}^{N_\lambda} [\mu_\star(\mathbf{x}_+, \lambda_g) w(\mathbf{X}_j, \lambda_g)]} \\ &= N_\lambda,\end{aligned}\quad (24)$$

thereby showing that $S(j+1)$ also holds. $\square$

2 ADDITIONAL ANALYSIS AND DETAILS

2.1 Control Extinction in Decomposition Tracking

In this section, we provide more results from our experiments with and analysis of decomposition tracking described in Section 4.3 of the main paper. Figure 1 is an extended version of Figure 7 in the main paper. Figures 2 and 3 together are an extended version of Figure 8 in the main paper.

2.2 Additional Spectral Tracker Comparison

Figure 4 shows a special case and how it can be handled with spectral and decomposition tracking. In particular, the figure shows a comparison of five multi-channel trackers on a volume *with* absorption for which the non-history-aware maximum-based probabilities with the reduced-termination-rates adjustment work well. The volume has spectrally varying absorption and scattering coefficients that are computed as the product of a spatially varying density and the vector multipliers (0.2, 0.3, 0.6) and (1.1, 1.2, 1.0), respectively. Figure 5 shows all of the collision coefficients for a part of the volume with full density and a part of the volume with half density. The amount of absorption in this test is set such that the scattering and null probabilities can always be set to bound the maximum proportions of scattering and null particles, respectively. This allows the maximum-based probability scheme to keep local collision weights, and thus path throughputs, at or below 1 at all times. To ensure fair and straightforward comparisons, we do not employ sample splitting or Russian roulette in these tests.

This configuration also allows decomposition to be implemented in a way that does not change results apart from decreasing the number of lookups. Variance can increase somewhat when decomposition splits a volume into achromatic and chromatic parts, which is the straightforward solution described in Section 5.2 of the main paper; however, if something is known about the properties of the volume in advance, the decomposition can be more optimally tailored to the situation. Because the albedo is constant in the volume of Figure 4, the albedo and collision probabilities can be set analogously for the control and residual components when decomposition is applied, ensuring that the decomposition does not affect the variance. Of the trackers compared in this figure, the decomposed spectral tracker using the maximum-based probability scheme has the lowest overall variance, the lowest maximum path throughput, and the lowest number of spatially varying density lookups.

2.3 Effects of Non-Bounding Free-path-sampling Coefficients on Renderers

Figure 6 shows a selection of the renders that were used to generate the data for the plots in Figure 13 of the main paper, which illustrates how the variance explodes with non-bounding free-path-sampling coefficients. The sequences of renders for this analysis use the history-aware average-based probabilities for maximum robustness to high albedos. While absorption stifles the explosion of the variance slightly, it does not prevent it. When computing the RMSEs and LTUVs, the maximum absolute pixel value was clamped to 65504 as if the images were encoded using standard 16-bit floating-point numbers; without this clamping, the RMSEs and LTUVs increase even more quickly due to occasional extreme outliers.

2.4 Comparison of History-aware Probability Schemes

Figure 7 compares the history-based and average-based probabilities on the non-absorbing color explosion, illustrating that the latter is robust to high albedos. This is an expanded version of Figure 14 in the main paper.

2.5 Additional Octree Depth Comparisons

In this section, we provide additional comparisons of delta and decomposition tracking performance for various octree depths. In Table 2 of the main paper, we provide a comparison using the optically thicker cloud from Figure 11 of the main paper. In addition to reproducing that table here as Table 2, we also provide in Table 1 the same comparison using the optically thinner cloud from Figure 11 of the main paper. Just like in the main paper, the timings were measured on a dual 12-core 2.50 GHz Intel Xeon E5-2680 v3 but we report them in single-core equivalents.




Figure 1. For the media in columns (a) and (c) we visualize the average relative cost ρ (purple curve) of weighted decomposition tracking w.r.t. delta tracking, the variance of path throughput until the first real collision (orange curve), and their product (black curve).




Figure 2. For the media in column (a), we visualize the the absolute cost of our *weighted* decomposition tracking in column (b), the relative cost of our *weighted* decomposition tracking w.r.t. delta tracking ($\rho = N_{dcmp}/N_{delta}$) in column (c), the mean path throughput $w(X)$ at the first real collision (including the local collision weight) in column (d), and the variance of $w(X)$ in column (e).




Figure 3. Continuation of Figure 2. For the media in column (a), we visualize the the absolute cost of our *weighted* decomposition tracking in column (b), the relative cost of our *weighted* decomposition tracking w.r.t. delta tracking ($\rho = N_{dcmp}/N_{delta}$) in column (c), the mean path throughput $w(X)$ at the first real collision (including the local collision weight) in column (d), and the variance of $w(X)$ in column (e).




Figure 4. Comparison of naive delta tracking (ND), free-path-sampling residual ratio tracking (FPSRR), spectral tracking with *average* probability scheme (SA), spectral tracking with *maximum* probability scheme (SM), and spectral tracking with *maximum* probability scheme and *decomposition* (SMD) in a canonical scene, where light paths enter a heterogeneous RGB medium (a). We visualize the mean cost (b), the maximum path throughput (c), and the variance of the RGB path throughput (d); all statistics are w.r.t. the first *real* collision with the medium or transmission through it. The images in column (e) show the noise in the path throughput—each pixel was obtained by averaging 2 instances of the tracker. Note how tracking all channels at once reduces the color noise. This comes at the cost of computing $3\times$ more weights; however, since most heterogeneous-volume renderers are memory bound, arithmetic operations are rarely the bottleneck in practice.




Figure 5. Collision coefficients for two points in the volume of Figure 4, showing the absorption “gap” between the scattering and null coefficients and the probability settings used by the history-aware average-based probability scheme.






Figure 6. A few of the renders that were used in the non-bounding free-path-sampling coefficient analysis. The volume in the first column exhibits absorption, the volume in the second column is non-absorptive. Both volumes have the same colored extinction coefficients. The exacerbated noise in certain regions is caused by a mixture of positive and negative fireflies, all of very high magnitude.

(a) History-aware maximum-based probabilities



(b) History-aware average-based probabilities



(c) Reference




Figure 7. Approximately equal-time comparison of two of the spectral tracking probability schemes presented in the main paper on a volume without absorption, illustrating the ability of the history-aware average-based probabilities to bound the throughput. The images were rendered with 16 spp, and the exposure was reduced to better show the fireflies. The free-path-sampling coefficient does not perfectly bound the extinction in this case, resulting in a few fireflies in the reference image.

Table 1. Comparisons of delta and decomposition tracking performance for various octree depths. The scene is the thin cloud from Figure 11 in the main paper, rendered with 32 spp. The last two rows report the non-lookup tracker time.

Octree depth	1	4	8	16
Octree leaves visited	0.12 G	0.31 G	1.32 G	2.02 G
Lookup num. (Delta)	0.74 G	0.24 G	0.080 G	0.074 G
Lookup num. (Decomp)	0.74 G	0.23 G	0.032 G	0.022 G
Lookup time (Delta)	201 s	88 s	44 s	43 s
Lookup time (Decomp)	201 s	87 s	26 s	20 s
Octree time (Delta)	3.58 s	40 s	216 s	365 s
Octree time (Decomp)	3.58 s	42 s	253 s	414 s
Tracker time (Delta)	49 s	30 s	78 s	115 s
Tracker time (Decomp)	49 s	31 s	98 s	142 s

Table 2. Comparisons of delta and decomposition tracking performance for various octree depths. The scene is the thick cloud from Figure 11 in the main paper, rendered with 32 spp. The last two rows report the non-lookup tracker time.

Octree depth	1	4	8	16
Octree leaves visited	1.86 G	2.06 G	3.11 G	4.05 G
Lookup num. (Delta)	106 G	22.8 G	2.96 G	2.49 G
Lookup num. (Decomp)	106 G	22.6 G	1.71 G	1.03 G
Lookup time (Delta)	20065 s	5785 s	1030 s	876 s
Lookup time (Decomp)	20065 s	5472 s	536 s	336 s
Octree time (Delta)	64 s	315 s	734 s	916 s
Octree time (Decomp)	64 s	319 s	714 s	918 s
Tracker time (Delta)	6108 s	1389 s	299 s	321 s
Tracker time (Decomp)	6108 s	1379 s	280 s	300 s