High-Performance Graphics 2025
A. Knoll and C. Peters
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 8

Supplemental material for Triangle Rejection Sampling for Density-Equipped Meshes on GPU

Jérémie Schertzer [1,2] Theo Thonat [1] Tamy Boubekeur [1]

[1] Adobe , [2] Mongolia International University

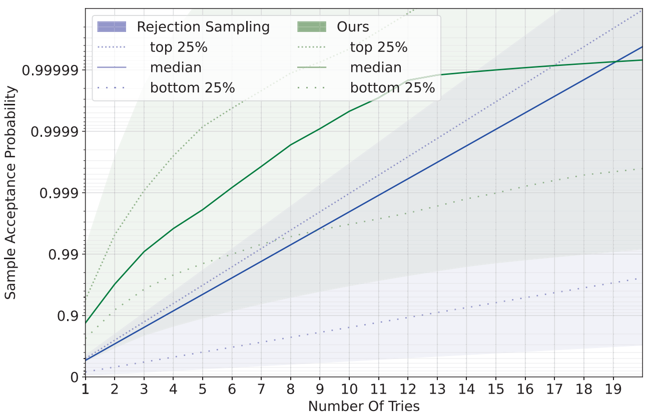

Figure 1: *Samples acceptance probability (**logarithmic scale**) with respect to the number of tries for our method and plain rejection sampling. Median and top/bottom 25% quantile over our benchmark are shown as lines, while filled areas show the range between best and worst cases.*

1. Sample acceptance analysis

Rejection sampling iteratively proposes new samples when a sample is rejected. The number of tries it takes to accept a sample should be as low as possible for good performance, especially in a SIMT context where a long run of rejected samples affects a whole set of threads. We describe in this section how our triangle stratification greatly improves in practice the probability of accepting a sample with respect to plain rejection sampling.

Single try. Rejection sampling first selects a triangle with a probability proportional to its area, and then accepts a random point on this triangle proportionally to the ratio between the point density provided by $\mathcal{D}$ and the maximum density over the mesh. So the probability $\mathbb{P}_{\text{RS}}$ of accepting a sample with a single try is:

$$\mathbb{P}_{\text{RS}} = \frac{1}{\sum \text{Area}(\mathcal{T})} \sum_{\mathcal{T}} \text{Area}(\mathcal{T}) \frac{\overline{\mathcal{D}}_{\mathcal{T}}}{\mathcal{D}^{\max}_{\mathcal{M}}} = \frac{\overline{N}}{\text{Area}(\mathcal{M})\mathcal{D}^{\max}_{\mathcal{M}}} \tag{1}$$

Our method follows a similar scheme, but instead assigns samples to triangles proportionally to their expected sample count $\overline{N}_{\mathcal{T}}$ and relies on the maximum density per triangle instead of the maximum density over the whole mesh. Therefore our probability $\mathbb{P}_{\text{TRS}}$ of accepting a sample with a single try is:

$$\mathbb{P}_{\text{TRS}} = \frac{1}{\sum \overline{N}_{\mathcal{T}}} \sum_{\mathcal{T}} \overline{N}_{\mathcal{T}} \frac{\overline{\mathcal{D}}_{\mathcal{T}}}{\mathcal{D}^{\max}_{\mathcal{T}}} \tag{2}$$

We now compare the two probabilities, using $\lambda_{\mathcal{T}} = \frac{\mathcal{D}^{\max}_{\mathcal{T}}}{\mathcal{D}^{\max}_{\mathcal{M}}} \leq 1$ for convenience, and noting that $\overline{N}_{\mathcal{T}} = \text{Area}(\mathcal{T}) \cdot \overline{\mathcal{D}}_{\mathcal{T}}$:

$$\begin{aligned} \frac{\mathbb{P}_{\text{RS}}}{\mathbb{P}_{\text{TRS}}} &= \frac{\left(\sum \overline{N}_{\mathcal{T}}\right)^2}{\sum \text{Area}(\mathcal{T}) \sum \frac{1}{\lambda_{\mathcal{T}}} \overline{N}_{\mathcal{T}} \overline{\mathcal{D}}_{\mathcal{T}}} \leq \frac{\sum \overline{N}_{\mathcal{T}} \overline{\mathcal{D}}_{\mathcal{T}}}{\sum \frac{1}{\lambda_{\mathcal{T}}} \overline{N}_{\mathcal{T}} \overline{\mathcal{D}}_{\mathcal{T}}} \\ &\leq \frac{\sum \overline{N}_{\mathcal{T}} \overline{\mathcal{D}}_{\mathcal{T}} \lambda_{\mathcal{T}}}{\sum \overline{N}_{\mathcal{T}} \overline{\mathcal{D}}_{\mathcal{T}}} \leq \max_{\mathcal{T}} \lambda_{\mathcal{T}} \leq 1 \end{aligned} \tag{3}$$

The Cauchy–Schwarz inequality was applied on the numerator with $\overline{N}_{\mathcal{T}} = \text{Area}(\mathcal{T})^{\frac{1}{2}} \cdot \text{Area}(\mathcal{T})^{\frac{1}{2}} \overline{\mathcal{D}}_{\mathcal{T}}$, followed by the Jensen inequality using the convex function $x \mapsto x^{-1}$. Our probability of accepting a sample after a single try is therefore always greater or equal than the one from plain rejection sampling. Equation (3) shows that the benefit over plain rejection sampling depends on how much the per-triangle max density $\mathcal{D}^{\max}_{\mathcal{T}}$ is a tighter bound than the mesh maximum density $\mathcal{D}^{\max}_{\mathcal{M}}$, especially for triangles with a large expected sample count and a large average density.

Many tries. For plain rejection sampling, the sampling process is independent for each new try, so the probability of accepting a sample after n tries is:

$$\mathbb{P}_{\text{RS}}(n) = 1 - (1 - \mathbb{P}_{\text{RS}})^n = 1 - \left(\frac{\sum \text{Area}(\mathcal{T}) \left(1 - \frac{\overline{\mathcal{D}}_{\mathcal{T}}}{\mathcal{D}^{\max}_{\mathcal{M}}}\right)}{\sum \text{Area}(\mathcal{T})} \right)^n \tag{4}$$

On the contrary, our sampling process only proposes new locations on the same triangle. So our probability of accepting a sample after n tries is:

$$\mathbb{P}_{\text{TRS}}(n) = 1 - \frac{1}{\sum \overline{N}_{\mathcal{T}}} \sum_{\mathcal{T}} \overline{N}_{\mathcal{T}} \left(1 - \frac{\overline{\mathcal{D}}_{\mathcal{T}}}{\mathcal{D}^{\max}_{\mathcal{T}}}\right)^n \tag{5}$$

$\mathbb{P}_{\text{RS}}(n)$ is a single exponential whose base is the average rejection probability $\overline{q} = \frac{1}{\sum \text{Area}(\mathcal{T})} \sum \text{Area}(\mathcal{T}) \left(1 - \frac{\overline{\mathcal{D}}_{\mathcal{T}}}{\mathcal{D}^{\max}_{\mathcal{M}}}\right)$ while

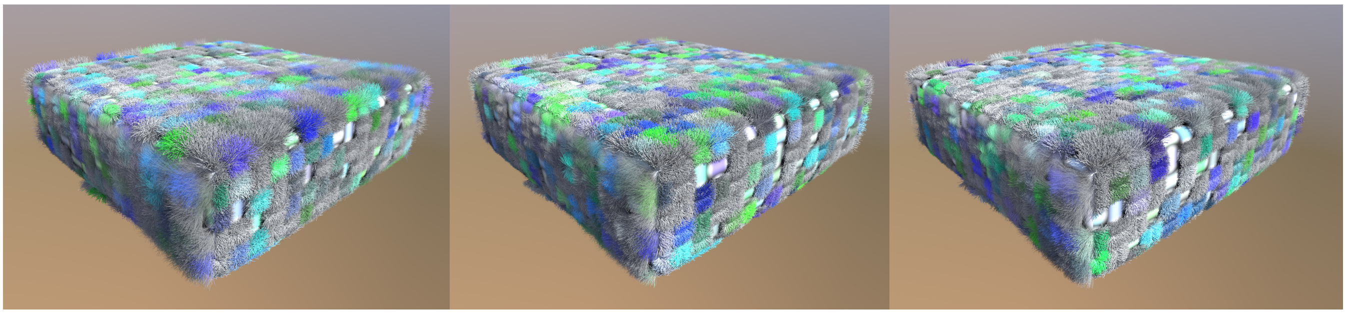

Figure 2: *Real-time procedural variations on the material of this pillow. The procedural model is edited interactively and includes a density map among its output channels which feeds the GPU path tracer on-the-fly.*

$\mathbb{P}_{\text{TRS}}(n)$ is the weighted average of multiple exponentials whose bases are the rejection probabilities $q_{\mathcal{T}} = 1 - \frac{\overline{\mathcal{D}}_{\mathcal{T}}}{\mathcal{D}_{\mathcal{T}}^{\max}}$. We show in Figure 1 that in practice, plain rejection sampling can have a better acceptance probability than ours, but only for numbers of tries large enough that the probability is extremely close to 1, so it matters little.

2. Additional application results

Figure 2 and Figure 3 provide additional examples of downstream application results, specifically for interactive fiber distribution design.