Eurographics Symposium on Rendering 2021
A. Bousseau and M. McGuire
(Guest Editors)

Volume 40 (2021), Number 4

Supplemental Material : Optimised Path Space Regularisation

paper1026

1. Multi-layered Materials

We can also apply our regularisation technique in the presence of multi-layered materials without further optimisation as long as a conversion to Beckmann roughness is available for every considered layer. When evaluating a regularised path of length k, we compute the accumulated roughness at vertex $k-1$ independently for every layer. We can then sample the regularised lobes based on their associated weights, as would be done for non-regularised BSDFs and compute the regularised contribution. For every other vertex before vertex $k-1$ we can simply assume the non-regularised sampled lobe's roughness for the path roughness. The WATER DROPLETS scene in Figure 1 applies our regularisation using the strategy described above.

2. Extended Comparison against Specular Manifold Sampling

In this section, we provide further comparison between path tracing with regularised NEE against different variant of Specular Manifold Sampling. In Figure 2 we demonstrate how the water's depth influences both our regularisation and Specular Manifold Sampling. Figure 3 shows that Specular Manifold Sampling's *biased* method, which uses a trial set size M to fix the maximal number of attempts to find deterministic connections to a light source, can loose a consequent amount of energy in presence of complex geometry.

3. Supplemental Video : Optimisation Process

During the optimisation process shown in our supplemental video, we can observe the effect of using our variance-aware objective function. The optimiser slowly reduces bias until the variance threshold defined by our hyperparameter β is reached, at which point the loss heavily penalises any further bias reduction and quickly readjusts the attenuation factors until reaching a local minimum.

4. Supplemental Video : Temporal Stability

In our supplemental video, we show the animation of the beach scene from our teaser to demonstrate the temporal stability of our method for $\beta = 0.001$ and $\beta = 0.05$.

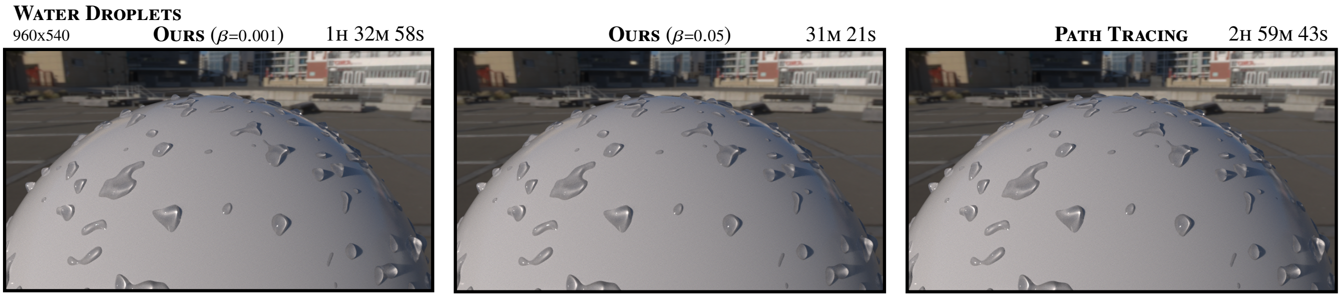

Figure 1: Equal noise comparison between a path traced reference and path tracing with regularised NEE for two sets of learnt parameters. *Our regularisation can help resolve the water droplets that introduce complex SDS paths, even on a multi-layered diffuse-specular material. Furthermore, it does not introduce visible bias in other regions that do not require regularisation.*

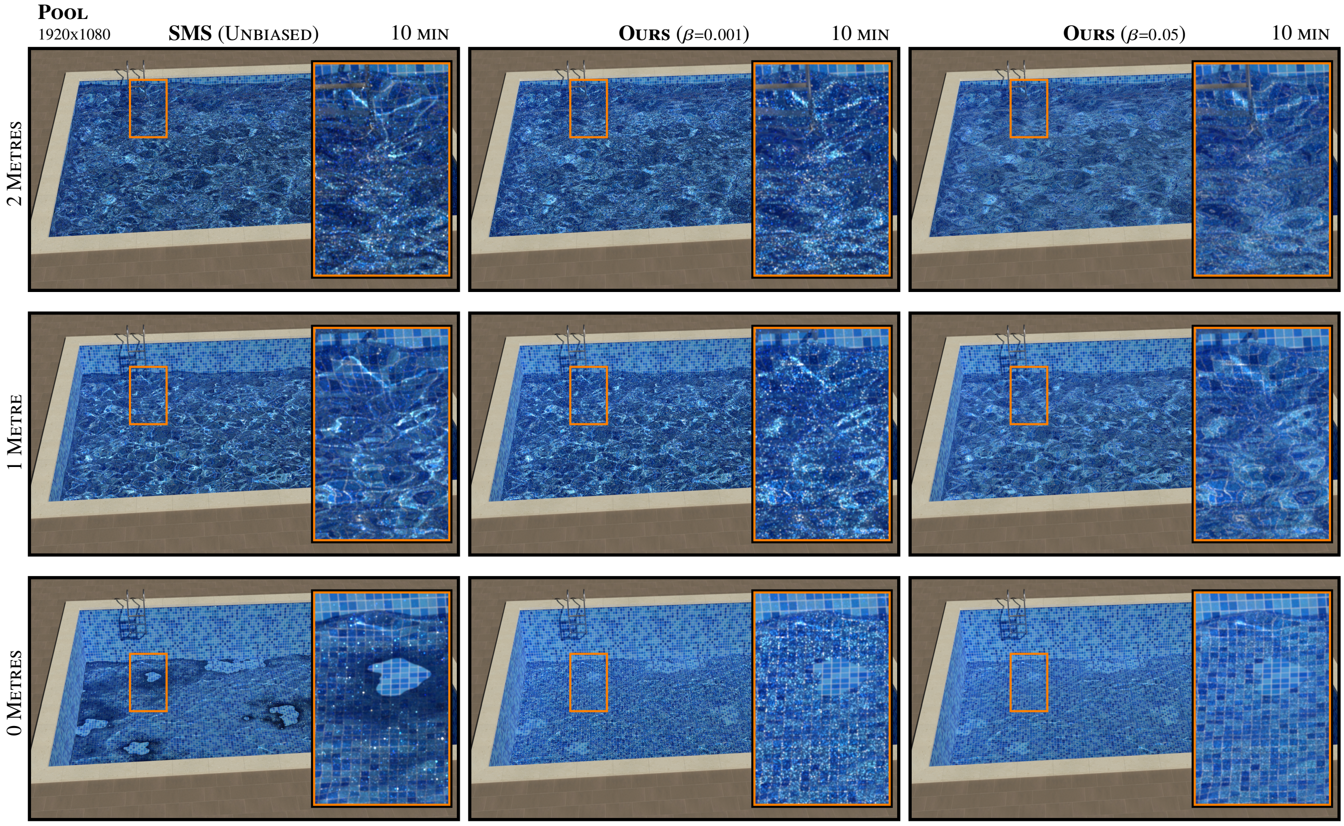

Figure 2: Equal time comparison between Unbiased Specular Manifold Sampling (SMS) and path tracing with regularised NEE for different depth of water. *As the depth of the water decreases, SMS performs better and finds caustic paths more consistently. However, at zero metres, when the water surface is very close to the pool's floor, SMS has difficulties converging. On the other hand, our method achieves very consistent results across the whole depth range. Moreover, we can notice that the bias introduced by our regularisation with* $\beta = 0.05$ *decreases as we reduce the depth, a direct consequence of our model not being geometry-aware.*

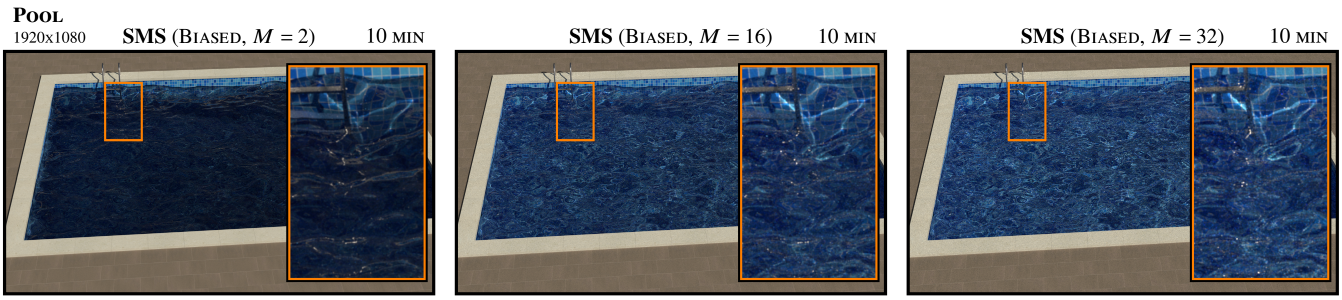

Figure 3: Equal time comparison between different trial set size M for the biased variant of Specular Manifold Sampling. *In the presence of complex geometry, the trial set size needs to be set relatively high to avoid significant energy loss (when compared to Figure 2). In this scene, the biased variant of SMS with* $M = 32$ *could only compute 11 samples per pixels in the time frame given.*