Virtual Blue Noise Lighting

Supplemental Document

TIANYU LI*, WENYOU WANG*, DAQI LIN, and CEM YUKSEL, University of Utah, USA

(a) *VPLs [Keller 1997]*
4.17M virtual lights (572 MB)
30 mintues

(b) *VSLs [Hašan et al. 2009]*
4.17M virtual lights (572 MB)
30 minutes

(c) ***VBNL (Ours)***
72,008 virtual lights (84 MB)
1 minute

Fig. S.1. *Images generated using 1M light paths with (a) VPLs [Keller 1997], (b) VSLs [Hašan et al. 2009], and (c) our VSLs with emission profiles. Note that 1M light paths result in 4.17M virtual lights with VPLs and VSLs in this example. The virtual light count for our VSLs are controlled independent of light path count used for computing the emission profiles, resulting in fewer virtual lights with less total storage (10 MB for lights and 74 MB for emission profiles), but improved image quality.*

This supplemental document includes additional examples and comparisons of our virtual blue noise lighting (VBNL) method to prior methods, including virtual point lights (VPLs) [Keller 1997], virtual spherical lights (VSLs) [Hašan et al. 2009], Rich-VSLs [Simon et al. 2015], and photon mapping with final gathering [Jensen 1996]. We also provide comparisons using parts of our method and combining our virtual sampling technique with prior virtual light generation methods.

ACM Reference Format:
Tianyu Li, Wenyou Wang, Daqi Lin, and Cem Yuksel. 2022. Virtual Blue Noise Lighting: Supplemental Document. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 3 (July 2022), 5 pages. https://doi.org/10.1145/3543872

S.1 COMPARISON TO REGULAR VPL AND VSL

Regular VPLs and VSLs do not store directional emission profiles, which gives them a relatively compact size, allowing a larger number of virtual lights to be stored. Our VSLs, however, store directional emission profiles, which makes the storage cost of each virtual light significantly more expensive. On the other hand, storing emission profiles allows reduction in the overall storage (using the same number of light paths) and improvement in light sampling quality and performance.

*Joint first authors: both authors contributed equally to this work.

2022. 2577-6193/2022/7-ART
https://doi.org/10.1145/3543872

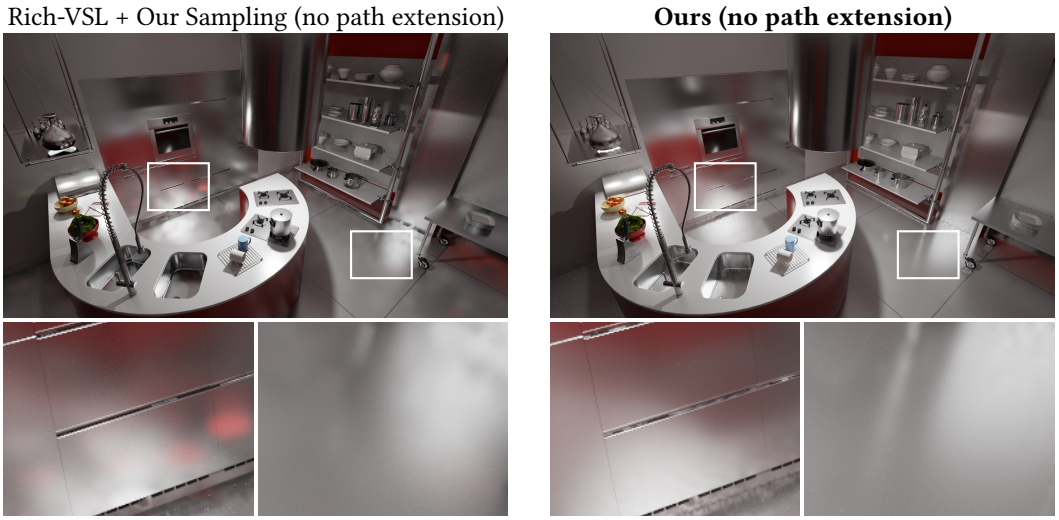

Fig. S.2. *Comparison of applying our light sampling without adaptive path extension using virtual lights generated by Rich-VSLs and our method. Both methods use the same number of virtual lights. Notice that Rich-VSLs produce visible splotches in this example, which do not appear with our method. The differences are purely due to the difference in virtual light distribution and emission profile generation methods.*

This is demonstrated in Figure S.1, comparing our method to regular VPLs and regular VSLs generated using the same number of light paths. In the case of VPLs and VSLs, each light vertex along a path is recorded as a virtual light, because each one of these regular VPLs and VSLs store a single incident light direction. In comparison, our method uses about 58× fewer VSLs, but with emission profiles computed using the same number of light vertices. The number of VSLs with our method is specified independently. Notice that, even though all three virtual lighting representations are computed using the same amount of light path information, virtual lights with emission profiles provide improved quality.

In addition, using much fewer lights for lighting estimation, our method renders significantly faster. Our result in Figure S.1 is rendered in 1 minute, while we used 30 minutes for rendering VPLs and VSLs to produce a reasonably converged image using an efficient implementation of BVH-based light sampling [Estevez and Kulla 2018].

S.2 USING OUR LIGHT SAMPLING FOR RICH-VSLS

Our light sampling method is tightly coupled with our virtual light generation and distribution approaches. In Figure S.2, we show that it can be unsuitable for Rich-VSLs. In this example both our VSLs and Rich-VSLs are sampled using our combined power-based and BSDF sampling technique without adaptive camera path extension. Notice that the results with Rich-VSLs include strong visible artifacts (in the form of colored splotches), which do not appear with our method. This is mainly because our method is able to produce a better virtual light distribution.

Introducing adaptive camera path extension helps with both methods, as shown in Figure S.3. In this case, however, Rich-VSLs must heavily rely on camera path extension, as the virtual light density is determined to be too low to avoid camera path extension. This reduces the efficiency of light sampling and results in substantial amount of noise in comparison to ours.

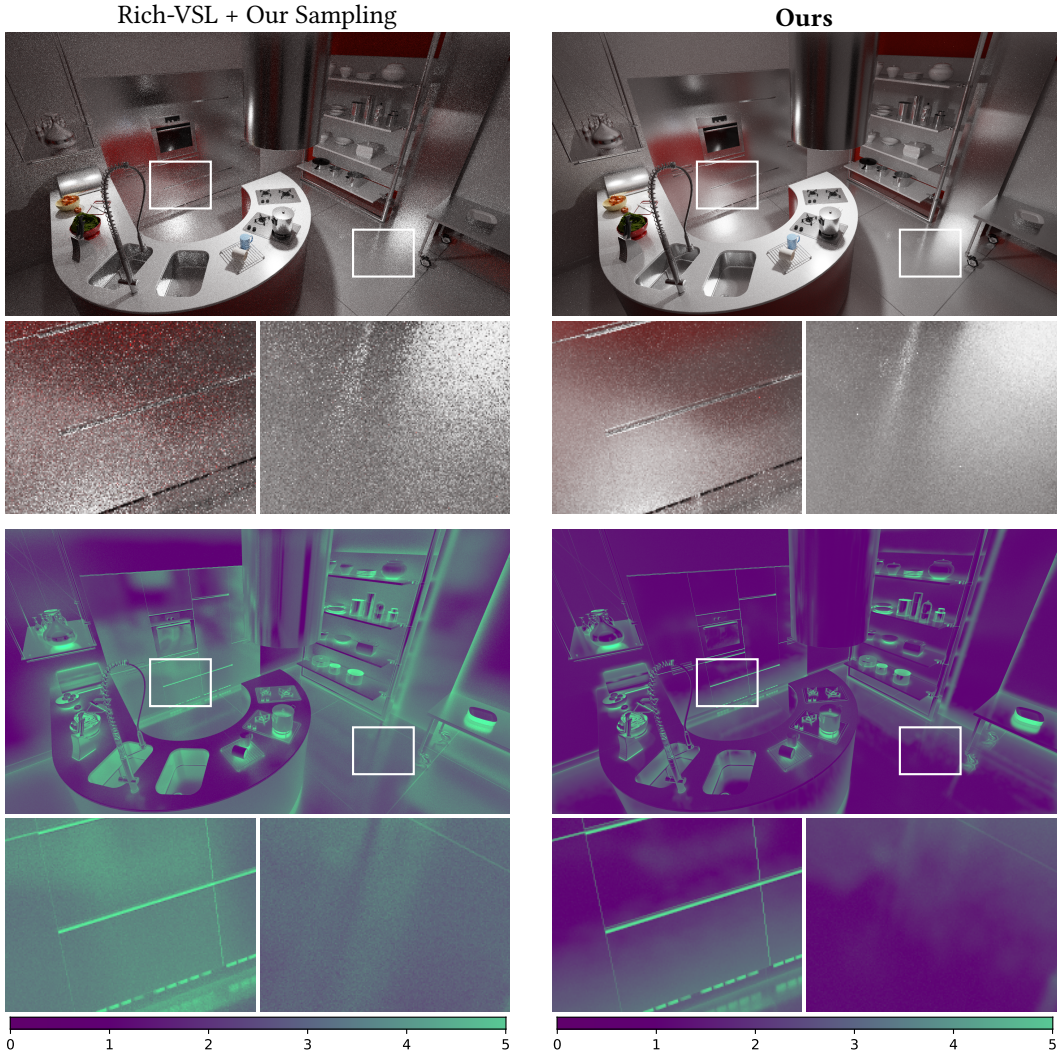

Fig. S.3. *Images rendered using our full light sampling method (including adaptive path extension) in 500ms (not including virtual light preparation time) with virtual lights generated using Rich-VSLs and our method. The bottom row shows visualizations of the average path length. Notice that Rich-VSLs use longer light paths for the majority of the image, reducing the effective performance of light sampling and resulting in higher noise with the same render time.*

S.3 COMPARISON TO PHOTON MAPPING

Our BSDF sampling method with adaptive path extension has similarities to photon mapping with final gathering [Jensen 1996]. In Figure S.4 we provide a direct comparison between the two methods, showing the improvement of the virtual light representation, as compared to directly using the photons for illumination estimation. As can be seen in the images, photon mapping with final gathering leads to significantly increased noise, even though it uses more than 3× longer render time. In comparison, our method uses the extra time for virtual light preparation and

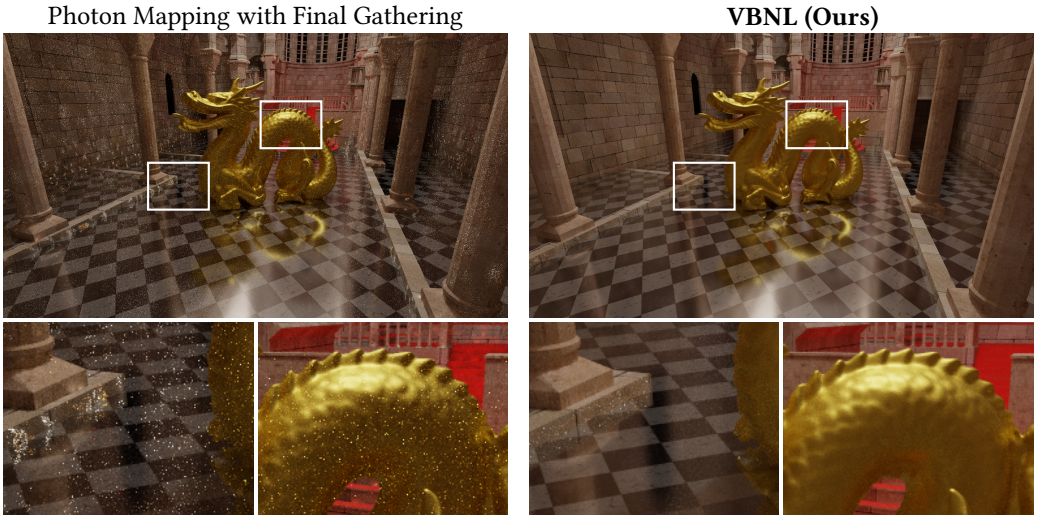

Fig. S.4. *Comparison of our method to photon mapping with final gathering [Jensen 1996]. Both images are generated using the same total time with 0.8 million photon paths.*

needs shorter time for rendering to produce an image with lower noise. More significantly, photon mapping produces clearly-visible splotches, using a relatively large density estimation radius. Using a progressive variant of photon mapping [Hachisuka et al. 2008] would help with eliminating such artifacts at the cost of longer render times.

S.4 EMISSION PROFILE COMPUTATION

Our delayed conversion approach for computing emission profiles improves the computation time without any visible degradation in the final rendering quality. Figure S.5 compares our delayed conversion method to directly computing the emission profiles. Note delayed conversion incurs the overhead of a separate pass, but still provides faster computation. The resulting emission profiles slightly differ, but the resulting difference is not noticeable in the final images, as demonstrated in this example.

REFERENCES

Alejandro Conty Estevez and Christopher Kulla. 2018. Importance sampling of many lights with adaptive tree splitting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 2 (2018), 25.

Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive Photon Mapping. In *ACM SIGGRAPH Asia 2008 Papers* (Singapore) *(SIGGRAPH Asia '08)*. Association for Computing Machinery, New York, NY, USA, Article 130, 8 pages.

Miloš Hašan, Jaroslav Křivánek, Bruce Walter, and Kavita Bala. 2009. Virtual spherical lights for many-light rendering of glossy scenes. In *ACM SIGGRAPH Asia 2009 papers*. 1–6.

Henrik Wann Jensen. 1996. Global illumination using photon maps. In *Eurographics workshop on Rendering techniques*. Springer, 21–30.

Alexander Keller. 1997. Instant radiosity. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 49–56.

Florian Simon, Johannes Hanika, and Carsten Dachsbacher. 2015. Rich-VPLs for improving the versatility of many-light methods. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 575–584.

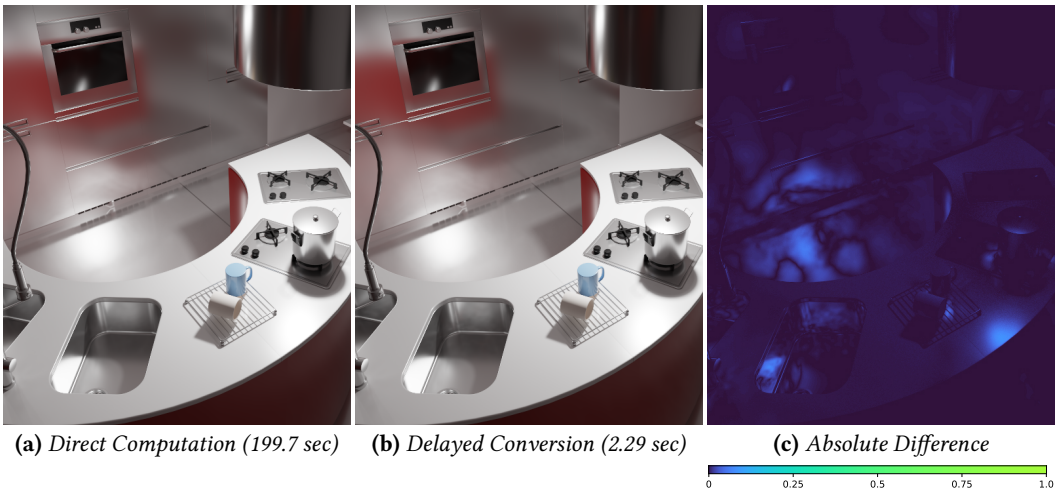

(a) *Direct Computation (199.7 sec)* **(b)** *Delayed Conversion (2.29 sec)* **(c)** *Absolute Difference*

Fig. S.5. *Comparison emission profile computations with photon tracing(20,000,000 light paths): (a) direct conversion of photons to outgoing radiance and (b) delayed conversion using photon splitting. Note that our delayed conversion with photon splitting provides significantly improved computation speed without a noticeable difference in the rendered result. (c) the Absolute Difference Map. Note that the difference is mainly on glossy surfaces where narrow specular lobe causes some aliasing in the resampling process in conversion.*