Adaptive Environment Sampling on CPU and GPU Supplementary Document

Asen Atanasov
Chaos Group

Vladimir Koylazov
Chaos Group

Blagovest Taskov
Chaos Group

Alexander Soklev
Chaos Group

Vassillen Chizhov
IMPRS, Saarland University

Jaroslav Křivánek
Render Legion
Charles University, Prague

ACM Reference Format:
Asen Atanasov, Vladimir Koylazov, Blagovest Taskov, Alexander Soklev, Vassillen Chizhov, and Jaroslav Křivánek. 2018. Adaptive Environment Sampling on CPU and GPU Supplementary Document. In *Proceedings of SIGGRAPH '18 Talks*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3214745.3214808

1 INTRODUCTION

In this supplementary document, we provide details on the integer Summed-Area Table (SAT) for environment map sampling. Table 1 shows how the performance of our algorithm is affected by coarser and finer map tiling and light grid resolution. We also provide additional results in Figures 1 and 3. Figure 4 showcases two important scenarios: rendering an exterior and a participating medium when distant shading points map to the same light grid cell.

Table 1: Ratios of render times for different T_u and G_y parameters compared to the render time using the default parameters $T_u = 16$ and $G_y = 50$ for the "Office" and "Living room" scenes. Since T_u and G_y span the angle of π, and T_v and G_x span 2π, we always set $T_v = 2T_u$ and $G_x = 2G_y$.

		Office		Living room	
T_u	G_y	CPU	GPU	CPU	GPU
8	25	1.27	1.21	1.16	1.05
8	50	1.12	1.12	1.06	1.00
16	25	1.08	1.07	1.10	1.00
16	100	1.16	1.10	0.94	0.95
32	50	1.05	1.05	0.97	0.97
32	100	1.79	1.37	0.97	0.97

2 SAT FOR IMAGE REGION SAMPLING

In the adaptive environment sampler we use a summed-area table (SAT) [Crow 1984] to be able to efficiently sample from an arbitrary rectangular subregion [Bitterli et al. 2015]. However, due to precision issues one cannot implement a high-resolution single-precision SAT directly as described in Bitterli et al. [2015]. Building a single-precision floating-point-valued SAT from an image with a high resolution and a large contrast ratio leads to substantial precision issues, see Table 2.

This is especially problematic for a GPU implementation, since the performance difference between using single-precision and

SIGGRAPH '18 Talks, August 12-16, 2018, Vancouver, BC, Canada
2018. ACM ISBN 978-1-4503-5820-0/18/08…$15.00
https://doi.org/10.1145/3214745.3214808

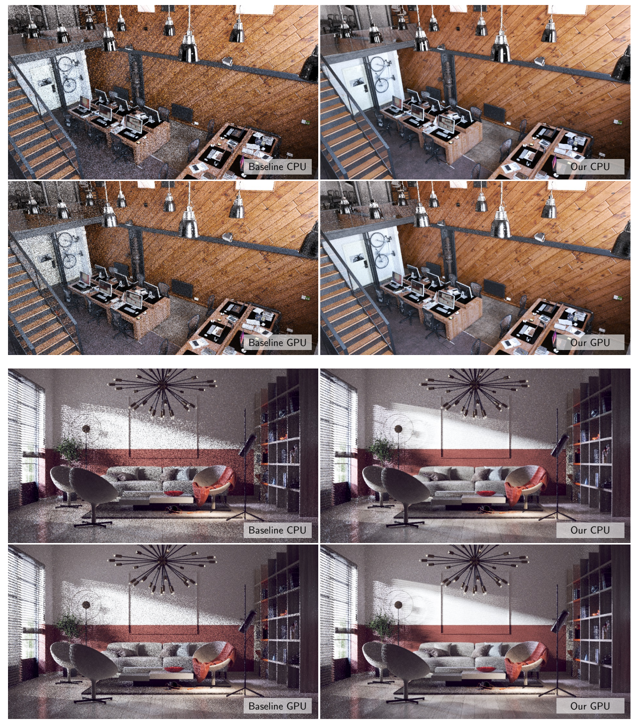

Figure 1: "Office" and "Living room" scenes rendered with classical environment sampling (*Baseline*) and our adaptive strategy. We present both CPU and GPU implementation results and show that our algorithm produces much cleaner images for the same time. Insets in Fig. 1 of the extended abstract are taken from these images. "Office" scene courtesy of Evermotion.

Table 2: Maximum norm error (MNE) and mean squared error (MSE) for sampling with int-valued and float-valued SAT.

	MNE		MSE	
HDR Image	int	float	int	float
Hallway	5.6×10^{-3}	4.2×10^{-0}	1.0×10^{-5}	3.8×10^{-1}
Night	1.2×10^{-3}	1.0×10^{-0}	4.9×10^{-7}	8.6×10^{-3}
Day	2.1×10^{-4}	1.2×10^{-1}	1.4×10^{-8}	4.1×10^{-4}
Sunset	1.8×10^{-4}	1.2×10^{-1}	1.1×10^{-8}	3.6×10^{-4}

Figure 2: Reconstruction of Hallway HDR image in resolution 10000×5000 using 32-bit integer (top image) and 32-bit float (middle image). The bottom image is the difference between the original image and its float reconstruction. The difference between the original image and its integer reconstruction is a nearly black image and it is omitted here. For an accurate error measurement see Table 2.

double-precision values is considerable. The problem is the rounding error, as this error increases with the image resolution. Additionally, the error grows along both the x and y axes producing a conspicuous pattern, see Figure 2. The goal of this section is to present our findings and solutions concerning this issue.

In the work of Bitterli et al. [2015], while the radiance is evaluated on the full-resolution image, sampling is done on a rescaled 512×512 SAT. This resolution constraint may help to avoid large errors but it also does not permit to importance sample fine structures in higher resolution images. A method to mitigate the error was also proposed by Hensley et al. [2005], where an offset is applied to the initial image, after which the SAT is built from the modified image. This results in a non-monotonic table that has sums of lower magnitude in each entry. While this technique does improve the error moderately, it is usually insufficient for higher-resolution images. Finally, Crow [1984] has proposed a partitioning of the SAT into regions, which works very well but is complex and has additional computational and memory requirements.

Our solution to the problem is building integer-valued SATs from the floating-point images. To the best of our knowledge, this has not been done previously. The improvements we get for the error compared to floating-point SATs are substantial, see Table 2 and Figure 2. In our implementation we perform a simple linear remapping such that the floating-point range $[0, sum([0, width) \times [0, height))]$ is mapped onto the integer range $[0, 2^{32})$. In practice this works well even for HDR images with resolutions of up to 15000×7500.

2.1 Rounding error

While floating-point-valued images enable a nonlinearly distributed range of values, this does not carry over to floating-point-valued SATs. In a floating-point SAT each entry is the sum of multiple floating-point values, which results in the SAT representing pixels values correctly only as long as the sum fits in the mantissa of each entry, effectively transforming the advantages of floating-point representation into rounding error. Therefore, integer fixed-point representation should be preferred in order to save the bits from the exponent and avoid the error introduced by it.

2.2 Integer-valued SAT

The main advantage of an integer SAT over a floating-point one is that the introduced error is a lot smaller, see Table 2, and that it is uniformly distributed, see Figure 2. It is important to mention, that neither integer-valued SATs nor floating-point-valued SATs have the capacity to deal with all possible HDR images correctly since RGBE supports a contrast ratio over 2^{256}. However, we have confirmed that our 32-bit integer SAT works well with images used in practice.

2.3 SAT building

We originally build the SAT in double-precision by performing two passes: one to sum the rows and another to sum the columns. Then it is remapped to an integer SAT. This approach is fairly straightforward to implement and parallelize. Note that faster approaches have been proposed based on recursive doubling [Hensley et al. 2005] and balanced-trees [Slomp et al. 2012].

ACKNOWLEDGEMENTS

The HDR image "Hallway" and the HDR images "Day", "Sunset" and "Night" were kindly provided by Wouter Wynen (Aversis 3D), and Peter Sinitra and Marek Denko (NoEmotion), respectively.

REFERENCES

Benedikt Bitterli, Jan Novák, and Wojciech Jarosz. 2015. Portal-Masked Environment Map Sampling. *Comput. Graph. Forum* 34, 4 (July 2015), 13–19. https://doi.org/10.1111/cgf.12674

Franklin C. Crow. 1984. Summed-area Tables for Texture Mapping. *SIGGRAPH Comput. Graph.* 18, 3 (Jan. 1984), 207–212. https://doi.org/10.1145/964965.808600

Justin Hensley, Thorsten Scheuermann, Greg Coombe, Montek Singh, and Anselmo Lastra. 2005. Fast Summed-Area Table Generation and its Applications. *Computer Graphics Forum* (2005). https://doi.org/10.1111/j.1467-8659.2005.00880.x

Marcos Slomp, Michihiro Mikamo, and Kazufumi Kaneda. 2012. Fast Local Tone Mapping, Summed-Area Tables and Mesopic Vision Simulation.

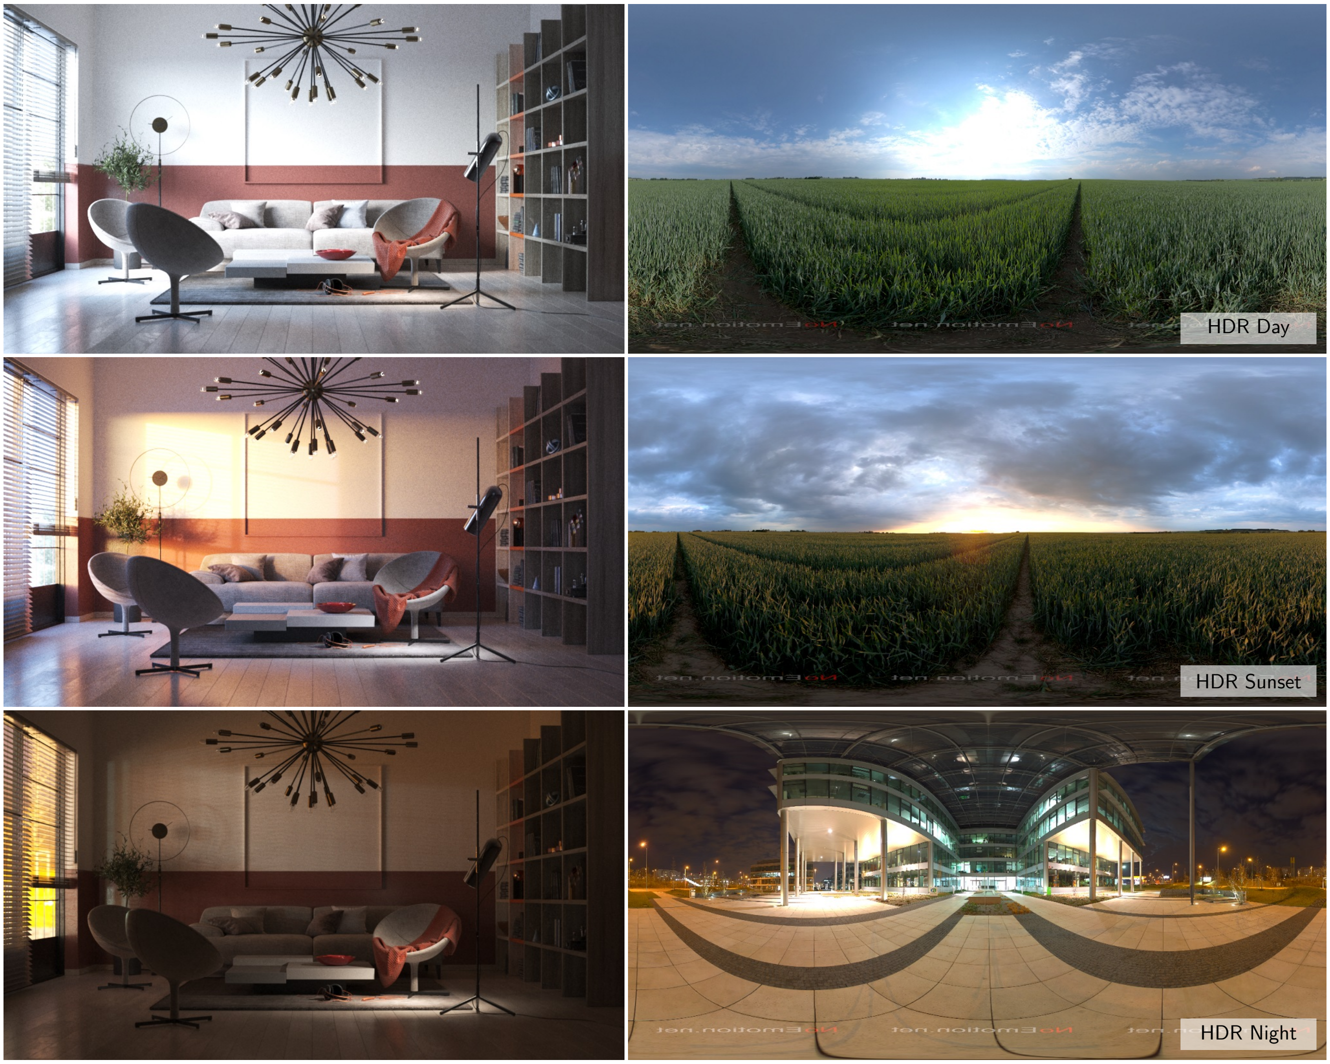

Figure 3: Living room scene lit with four different HDR images: "Day", "Sunset" and "Night". The effective speedup of our algorithm, measured for the same noise level on CPU/GPU for this scene is: Day - 2.2/1.6, Sunset - 1.9/1.6 and Night - 3.8/3.

Figure 4: Garden scene without (left) and with (right) participating medium. The effective speedup of our algorithm for the same noise level is 2.3 (left) and 3.4 (right) with CPU and 1.8 (left) and 2.6 (right) with GPU. The scene with the fog showcases the important scenario when many distant points are projected to the same light grid cell.