

Denoising at Scale for Massive Animated Series

Tamy Boubekeur
Malik Boughida
LTCI, Telecom ParisTech,
Paris-Saclay University

Laurent Noël
Jérémie Defaye
Farchad Bidgolirad
Ubisoft Motion Pictures




Figure 1: All-effects denoising of production Monte Carlo rendering (32 spp). Close-up are indicated with colored frames.

ABSTRACT

In the modern era of physically-based shading, removing the substantial amount of high frequency noise produced by Monte Carlo rendering techniques is a key challenge for production renderers. Beyond the recent advances in sample-based and feature-based denoising, production constraints and scale introduce additional mandatory features for candidate denoisers. In this talk, we discuss how denoising is deployed in *Shining*, the production renderer developed by Ubisoft Motion Pictures for the *Rabbids Invasion* animated TV series. The scale of the show, as well as the required control for artists, led us to the integration of a sample-based denoiser, which enables per-AOV denoising control, with a minimum overhead regarding engine integration and production workflow. As a result, *all-effects* denoising is made possible for the new TV series season and proved useful in numerous lighting and material scenarios. At the core of the denoising pipeline, our *BCD* algorithm, recently made open source, provides a robust and fast mechanism to filter out Monte Carlo noise while retaining features, for complex lighting and viewing conditions, with trivial per-AOV setup.

CCS CONCEPTS

• Computing methodologies → Rendering; Ray tracing;

KEYWORDS

path tracing, Monte Carlo rendering, denoising

ACM Reference Format:

Tamy Boubekeur, Malik Boughida, Laurent Noël, Jérémie Defaye, and Farchad Bidgolirad. 2018. Denoising at Scale for Massive Animated Series. In

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH '18 Talks, August 12–16, 2018, Vancouver, BC, Canada

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5820-0/18/08.

<https://doi.org/10.1145/3214745.3214815>

Proceedings of SIGGRAPH '18 Talks. ACM, New York, NY, USA, 2 pages.
<https://doi.org/10.1145/3214745.3214815>

1 INTRODUCTION

Unilateral Monte Carlo path tracing offers a predictable, simple and robust physically-based solution to the rendering equation. For the *Rabbids Invasion* animated TV series - a full 3D animated TV series broadcasted in more than 110 countries and reaching 290 millions people in 2016 - Ubisoft Motion Pictures has developed *Shining*, a new production path tracer adapted to the needs of such a series. As for all path tracing engines, even using hundreds of samples per-pixel (spp), noise often remains present in the final image. In order to remove this noise and control rendering costs, *Shining* integrates *BCD*, a new denoiser which has been developed with these specific production constraints in mind. *BCD* is successfully used in the production of the new season of the *Rabbids Invasion* show and smoothly integrates with the artist-friendly production workflow, where each AOV (i.e., render pass) shall be denoised or not, independently, to later get composited under the artist control.

2 PRODUCTION

The *Rabbids Invasion* TV series features 78 episodes per season, with each episode lasting 6.45min. at 24 frames per season, which amounts to 748k images rendered at 720p resolution. Each image is composited from 5 to 6 AOVs, each of which must be denoised independently, and two try-and-test images are produced on average for the adjustment of one final frame. As such, each season represents more than 13 millions frames to denoise, with each set of AOVs consuming 300 spp initially (maximum sampling rate of the adaptive sampler), reduced to 200 spp thanks to the denoiser.

3 CONSTRAINTS

Considering the massive amount of frames to denoise for each season, five major production constraints were identified, namely:

- **low invasiveness** in the core renderer for easy integration;

- **all-effects** ability to filter e.g., volumetric effects, complex geometry, specular and transparent materials, defocused or motion-blurred regions, for which screen-space features (e.g. normal, BRDF) cannot be used for guiding the filtering;
- **reduced number of parameters**, so that artists can quickly activate denoising on a per-AOV basis;
- **natural temporal stability**, as consecutive image data may not be maintained on a given rendering node, and the engine must be able to denoise independently each image;
- **speed**, to cope with the amount of images to generate within a few months for the entire season.

4 BCD DENOISER

These constraints have led to the integration of *BCD* [Boughida and Boubekeur 2018] into the *Shining* engine. *BCD* is an open source implementation of the sample-based denoising algorithm proposed by Boughida and Boubekeur [2017]. This denoiser is agnostic to the way samples (i.e., single path color responses) are generated, handling all kinds of effects from soft shadows on specular materials to (semi-)transparent medium under defocus rendering. The key concept of *BCD* is to maintain per-pixel sample statistics – average (noisy) color value, histogram of the samples color distribution and covariance matrix of this distribution – to later (i) seek for pixels of similar nature using their histograms as a signature [Delbracio et al. 2014] and (ii) build a bayesian model of this distribution [Lebrun et al. 2013], over all such similar pixels, in the form of a single anisotropic gaussian in RGB space. This model is used to collaboratively filter them, by estimating a local empirical model of the noise. As a result, *BCD* is well adapted to high spp scenarios such as production rendering. To integrate *BCD* in *Shining*, we interfaced the sample stream coming from the *Shining* path tracing, with a *sample accumulator* that maintains the mandatory statistics. Once the rendering process is over, the sample accumulator produces the images statistics and feed *BCD* for denoising, which then outputs a denoised image in seconds (CPU execution mode).

5 RESULTS

BCD is made available to artists for all AOVs independently. Its single parameter policy makes possible the individual denoising control for each AOVs in a reasonable amount of time: with 30 to 45 minutes of rendering time per-frame, the 30 sec. denoising appears negligible. The series production takes place on a render farm made of 180 rendering nodes, each nodes featuring 18 Intel Xeon E3-12xx cores @ 2.5GHz and 45Gb RAM. The entire rendering and denoising is happening in main memory, *BCD* being integrated as a dynamic library. Interestingly, we observe that often, the default parameters provide a convincing enough result and therefore, artists just need to decide whether they want to denoise an AOV or not (see Fig. 2), maintaining the fluidity of the post-processing workflow. Regarding performances (Tab. 1), the denoising cost is mainly dictated by the sample accumulation phase, which dominates the overhead when reaching high image resolution/spp. The whole process, however, still represents less than 0.5% of the total rendering time. Currently, our denoiser runs in CPU mode, while a GPU mode is already available in the public source code, and will be activated as soon as the rendering nodes feature GPUs.

Table 1: Denoising performances (in sec.) for HD and 4K rendering at various production level sampling rates (spp).

Res. & spp	Samples Accumulation	Denoising	Total
HD 64spp	28.8	14.8	42.8
HD 128spp	33.8	16.8	50.6
HD 256spp	30.6	18.3	48.9
HD 512spp	44.4	93.0	137.4
4K 64spp	100.0	54.0	154.0
4K 128spp	82.9	40.2	123.1
4K 256spp	93.9	64.8	158.7
4K 512spp	167.0	445.1	612.1




Figure 2: Per-AOV denoising. From top to bottom: diffuse direct, diffuse indirect, glossy direct, glossy indirect and glossy reflection/refraction AOVs are independently denoised using *BCD* before being composited in the finale image.

6 DISCUSSION

Just such as automatic tone mapping has progressively made space for interactive color grading, the design of the denoising post-process is becoming an artist-tuned stage. Currently, the artistic control is mainly located in the binary decision to denoise or not each AOV, before compositing. On top of getting rid from the Monte Carlo artifacts, this workflow reduces by a third the sampling rate (spp) used during production. We foresee in the future that more powerful control primitives may help the artist to better balance simulation and denoising, possibly in a spatially-varying way, to produce larger amounts of artifact-free images even faster. A deeper integration of the *BCD* algorithm, which has proved efficient at controlling adaptive rendering [Boughida and Boubekeur 2017], is also a potential future work. The evolution of the *BCD* denoiser can be followed on its dedicated GitHub repository.

ACKNOWLEDGMENTS

This work is partially supported by the French Agence Nationale pour la Recherche under grant ANR 16-LCV2-0009-01 ALLEGORI and by BPI France, under grant PAPAYA.

REFERENCES

- Malik Boughida and Tamy Boubekeur. 2017. Bayesian Collaborative Denoising for Monte Carlo Rendering. *Computer Graphics Forum (Proc. EGSR 2017)* 36, 4 (2017), 137–153.
- Malik Boughida and Tamy Boubekeur. 2017–2018. *BCD: Bayesian Collaborative Denoiser for Monte-Carlo Rendering*. <https://github.com/superboubek/bcd/>. (2017–2018).
- Mauricio Delbracio, Pablo Musé, Antoni Buades, Julien Chauvier, Nicholas Phelps, and Jean-Michel Morel. 2014. Boosting Monte Carlo Rendering by Ray Histogram Fusion. *ACM Transactions on Graphics* 33, 1, Article 8 (2014), 8:1–8:15 pages.
- M. Lebrun, A. Buades, and J. M. Morel. 2013. A Nonlocal Bayesian Image Denoising Algorithm. *SIAM Journal on Imaging Sciences* 6, 3 (2013), 1665–1688.