Automatic Parameter Control for Metropolis Light Transport

Károly Zsolnai and László Szirmay-Kalos

Budapest University of Technology and Economics, Hungary



Figure 1: *Images rendered with PSSMLT in 15 minutes with* $p_l = 0.75$ *(image above) and* $p_l = 0.25$ *(below).*

Abstract

Sophisticated global illumination algorithms usually have several control parameters that need to be set appropriately in order to obtain high performance and accuracy. Unfortunately, the optimal values of these parameters are scene dependent, thus their setting is a cumbersome process that requires significant care and is usually based on trial and error. To address this problem, this paper presents a method to automatically control the large step probability parameter of Primary Sample Space Metropolis Light Transport (PSSMLT). The method does not require extra computation time or pre-processing, and runs in parallel with the initial phase of the rendering method. During this phase, it gathers statistics from the process and computes the parameters for the remaining part of the sample generation. We show that the theoretically proposed values are close to the manually found optimum for several complex scenes.

1. Introduction

The performance of global illumination rendering algorithms [Szirm08] has been dramatically increased recently, thus these methods have become viable alternatives in production rendering and also in real-time applications [RDGK12]. However, high performance rendering algorithms often come with many control parameters that need to be set by the artist before starting the rendering process. Unfortunately, the optimal values of the control parameters are scene dependent and a careless setting would significantly slow down the convergence or could even result in images of severe artifacts in case of biased methods. In practice, parameter setting is based on previous experience and a trial and error approach, but this is unacceptable in production rendering where a large number of images are rendered, and also in virtual reality systems when the scene may change in time in a way that is not anticipated by the designer. In these cases robust rendering algorithms are needed that require no manual parameter setting and can deliver high quality images in rendering times comparable to manually controlled methods.

Metropolis Light Transport (MLT) [VG97] is known to be a robust approach that can efficiently handle a large variety of lighting effects and scenes. The power of MLT comes from the Metropolis-Hasting sampling that — unlike other importance sampling methods generating samples independently from a prescribed density — explores important light path regions in the scene by mutating previous light paths and evaluating the importance of the mutations. When a particular MLT algorithm is designed, we have to specify the tentative transition function that mutates a path, taking into account the following criteria:

• The mutation strategy should guarantee that the Markov chain is *ergodic*, i.e. it has an asymptotic distribution that is independent of the initial state. To satisfy this requirement, all light paths of non-zero contribution should be given a chance to be generated as a tentative sample

[©] The Eurographics Association 2013.

sooner or later. By giving chance to the full re-generation of a path, the ergodicity condition can be met.

- Unlike other Monte Carlo algorithms, the Metropolis algorithm generates not statistically independent, but correlated samples, which can increase the error [SKDP99, APSS01]. Thus, a secondary requirement for mutation strategies is to keep the correlation low, which means that the possibility of rejected samples should be reduced.
- Metropolis sampling converges to the desired probability distribution, but at the beginning of the process the samples are not selected with the required probability, which introduces some error in the estimation. This error is called as the *start-up bias* [SKDP99].

The large search space of control parameters is effectively limited and the reduction of rejection rate is also successfully addressed by the method of Primary Sample Space MLT (PSSMLT) [KSKAC02], which executes the perturbations in the space of uniformly distributed random numbers, called the primary sample space. PSSMLT is based on the recognition that any light path generator algorithm is a mapping between the primary sample space and the path space, and if importance sampling is involved in the path generation then more important regions are represented by larger volume in primary sample space. Thus, making uniform mutations in primary sample space automatically reduces the path perturbation size at important regions and increases it in unimportant regions. We can also say that the involved importance sampling in path generation does one part of the job of optimal importance sampling, which is further improved by the Metropolis-Hasting scheme. In this sense PSSMLT is a special type of random number generator that can be plugged into an arbitrary Monte Carlo method. In PSSMLT a simple way of ensuring ergodicity is to generate a primary space sample point from scratch, independently of the current paths with some positive probability. Such independent samples are often called *large steps*.

Theoretically, any positive large step probability makes the Markov process ergodic, but its selection affects convergence and thus rendering performance, thus it must be carefully set. This paper addresses this issue and proposes an adaptive method to control the large step probability. The model must be based on parameters that are easy to estimate while starting the algorithm.

2. The proposed mutation control method

If the path building strategy is efficient for the particular scene, the importance sampling is close to optimal, which means that the path contribution divided by the density is close to constant where it is non-zero. Note that zero contribution regions keep their zero contribution, so the path building strategy can only flatten the contribution of non-zero regions. Metropolis sampling should handle the residual variation of the integrand after transforming the domain according to the path building strategy. If the non-zero contribution regions have low variation in the primary sample space, the efficiency of Metropolis depends on how quickly it explores this space. However, if there is a significant residual variation after transformation, Metropolis should still focus on the high contribution regions, according to the principles of importance sampling. Thus, in the first case the efficiency can be characterized by the size of the explored domain, but in the second case, the integrand values should also be taken into account.

The efficiency of the Metropolis method in obtaining samples that are different from the previous ones can be characterized by the following three measures:

- Small step efficiency η_s that is the average probability that a small perturbation is accepted.
- Large step efficiency η_l that is the average probability that a large perturbation is accepted.
- The large step non-zero tentative sample probability η_0 is the average probability that a large perturbation proposes a non-zero contribution sample. As large steps are independent and the primary sample space has unit size, this probability is the volume of the path domain where paths have non-zero contribution.

Note that these parameters can be easily estimated in parallel to the running of the algorithm and can be obtained from a relatively few number of samples.

These values are not independent of each other when measured on a particular scene. As Metropolis sampling drives samples towards high importance regions where small mutations may cause smaller importance decrease, the relation $\eta_s \ge \eta_l$ generally holds. On the other hand, the rejection of a large mutation can be due to a mutation to a zero contribution point or to a point having positive but smaller contribution than the current state, thus $\eta_l \le \eta_0$.

The ratio η_l/η_0 characterizes the efficiency of the importance sampling of the path building strategy. If the integrand is flat in the nonzero contribution regions, then this ratio is close to 1. If it is not close to one, the integrand still has a significant variation which needs to be compensated by Metropolis sampling.

To explore the path space efficiently, both small and large perturbations must be applied to the path space walk. Small mutations propose new primary sample points that are in a neighborhood around the point defining the current state. Large mutations lead to an arbitrary point of the unit volume primary sample space. The type for the next mutation is given by the *large step probability* p_l . In the stationary case, all positive large step probabilities are equivalent. However, the convergence rate is affected, so, in order to find an optimal large step probability, the dynamic behavior of the Metropolis algorithm should be examined.

2.1. Flattened integrand

First we assume that the path building strategy was good in importance sampling, thus the non-zero contribution regions of the primary sample space have a roughly constant integrand. The convergence is fast if the Markov process explores the integration domain as quickly as possible. Random walks can be decomposed into *local explorations* when only accepted small mutations modify the sample locations. A local exploration phase is terminated when a large step is accepted, and the system starts exploring another part of the integration domain from the seed given by the large step.

The average probability that Metropolis sampling takes an accepted large step is then the product of the large step probability and its success ratio, i.e. $p_l\eta_l$. As the number of tried mutations before an accepted large step follows a geometric distribution, the expected length of local exploration is

$$E[N_{local}] = 1 + \frac{1 - p_l \eta_l}{p_l \eta_l}$$

where 1 is the large step that gets the process to start this local exploration and $1 - p_l \eta_l$ is the probability that the next sample will also belong to this local exploration.

During a local exploration, rejected mutations keep the original state while accepted small mutations modify the sample position additively, i.e. the new sample will be in the small neighborhood of the previous sample, where their difference is governed by the probability density of small mutations. As small mutations are independent, the variance of the perturbations caused by accepted small mutations is added, thus the average radius d of the space explored by N_{as} accepted small mutations grows proportional to the square root of the number of new samples. However, we cannot say that only small mutations can explore the full space. The first step of the local exploration, which is an accepted large step, also places a sample. As the perturbation of small steps is set to give the possibility to walk the whole space when the total number of mutations are executed, we can safely assume that the large step starting the local exploration is responsible for the same exploration. Thus, the space visited by a local exploration phase has expected radius

$$E[d] = \sigma \sqrt{N_{as} + 1}$$

where σ is the standard deviation of a single small mutation.

In a local exploration not all N_{local} steps belong to the category of new samples, only those mutations should count that lead to accepted small mutations:

$$E[N_{as}] = \frac{1 - p_l \eta_l}{p_l \eta_l} \cdot \frac{(1 - p_l) \eta_s}{1 - p_l \eta_l} = \frac{(1 - p_l) \eta_s}{p_l \eta_l}$$

since $(1 - p_l)\eta_s$ is the probability that a successful small mutation is tried, and $1 - p_l\eta_l$ is the probability that this sample does not terminate the local exploration. If we generate *M* Metropolis samples, in average $Mp_l\eta_l$ local explorations are established, each spreading over a subspace of radius $\sigma\sqrt{N_{as}+1}$, thus the size *D* of the total explored space,

having projected onto a line is

$$E[D] = M p_l \eta_l \sigma \sqrt{1 + \frac{(1-p_l)\eta_s}{p_l \eta_l}}.$$

The objective of large step probability selection is to maximize the size of the space explored by M samples, thus we obtain:

$$p_l = \underset{p_l}{\operatorname{argmax}} E[D] = \frac{\eta_s}{2(\eta_s - \eta_l)}$$

If large steps are almost as successful as small mutations, this formula can result in values that are larger than 1. This means that the large step probability should be set to 1, i.e. only large steps should be executed.

2.2. High variation integrand in the non-zero contribution regions of the primary sample space

So far, we assumed that the path building strategy already flattens the integrand transformed to primary sample space in the non-zero contribution regions, and thus the efficiency can be characterized by the speed of space exploration. However, when the transformed integrand still has large variation, Metropolis sampling should still focus on the peaks of the integrand, thus the explored size itself is not an appropriate measure for efficiency. Unfortunately, this analysis would also involve the consideration of the integrand, which cannot be robustly estimated with a few samples and in parallel with the sample generation. We can only state that in this case, the small steps must be given higher probability since they will concentrate regions of high contribution while they are poorer in exploring lower contribution regions. The analysis of the previous section would propose $p_l = 0.5$ for this case, which is thus an overestimation. Without robust estimates upon which a theoretical model can be built, we simply propose the application of $p_l = 0.25$ in this case. The cases of flattened and not flattened integrands can be distinguished by considering η_l/η_0 , where we set the threshold to 0.1 above which the integrand is considered as flattened.

3. Results

We have tried a variety of scenes to demonstrate robustness of the new control method and selected bi-directional path tracing as the path building strategy. Table 3 summarizes the statistical parameters obtained for these scenes, the proposed large step probability p_l and the range of optimal probabilities p_l^{opt} determined by an extensive parameter study.

The *LuxTime scene* is a typical indoors setup with multiple area light sources where only slight difficulties are present, such as the dial of the watch which can only be illuminated by light that passes through the glass. This scene is lacking complex specular paths and the illumination is mainly diffuse interreflection, which is successfully

[©] The Eurographics Association 2013.

Zsolnai and Szirmay-Kalos / Automatic Parameter Control for Metropolis Light Transport



Figure 2: Rendering results with three large step probabilities (high resolution images are in the supplemental material).

Scene	η_l	η_s	η_0	p_l	p_l^{opt}
LuxTime	0.377	0.783	0.985	0.95	(0.65-1)
Spheres	0.005	0.394	0.487	0.25	(0.2-0.3)
Chess Day	0.061	0.641	0.886	0.25	(0.2-0.4)
Cherry	0.126	0.487	0.911	0.67	(0.4-0.5)
Cornell	0.004	0.438	0.022	0.50	(0.5-0.6)
Glass Ball	0.088	0.489	0.87	0.61	(0.5-0.6)

Table 1: *Scene statistics with the proposed* p_l *values.*

addressed by bi-directional path tracing, indicating by the higher $\eta_{\rm I}/\eta_0\approx 0.4$ factor.

The *Spheres scene* includes dispersion and heavy volume scattering, where the light enters the scene following a path through multiple thick glass-like surfaces making it a very difficult light transport situation. Standard bidirectional path tracing is unable to render this scene efficiently, thus there is a significant residual variation of the integrand even in primary sample space, which is also shown by the very low $\eta_l/\eta_0 \approx 0.01$ factor.

The *Chess Day scene* looks simple, but there are significant glossy interreflections that cannot be mimicked effectively by the deterministic connection rays of bi-directional path tracing. As a result $\eta_l/\eta_0 \approx 0.07$.

In the *Cherry Splash scene* we can also observe complex specular paths, but they have significantly smaller total contribution than in the Chess Day and Spheres scenes.

The *Cornell Box scene* is a pathological case, where illumination comes in a light tube that allows the light to enter the box only after very many specular interreflections. In this scene, the size of the non-zero contribution primary sample space domain is 0.022, i.e. even after emphasizing important paths, non-zero contribution paths occupy only 2% of the space of paths. However, within this small domain, the integrand is relatively constant, thus $\eta_l/\eta_0 \approx 0.18$.

Finally, the *Glass Ball scene* is a typical outdoors setup with glossy interreflections, depth of field and caustics.

4. Conclusions

In this paper we analyzed the convergence properties of PSSMLT, and proposed a method to control its large step probability parameter. Our method is based on a few statistical parameters, including the success ratios of small perturbations, large perturbations, and perturbations leading to non-zero contribution tentative samples. These parameters can be obtained easily and robustly at the beginning of the rendering process. We also showed that these parameters tell us a lot about the scene and its suitability for sampling by the given path building strategy. As Metropolis is responsible to do importance sampling on the function that is already flattened by the path building strategy, we can build our proposed large step value on these parameters.

5. Acknowledgements

This work has been supported by TÁMOP-4.2.2.B-10/1-2010-0009 and OTKA K-104476. We thank Kai Schwebke for providing Lux-Time, Vlad Miller for the Spheres, Giulio Jiang for the Chess, Aaron Hill for the Cornell Box, Andreas Burmberger for the Cherry Splash and Glass Ball scenes.

References

- [APSS01] ASHIKHMIN M., PREMOZE S., SHIRLEY P., SMITS B.: A variance analysis of the Metropolis light transport algorithm. *Computers & Graphics 25*, 2 (2001), 287–294.
- [KSKAC02] KELEMEN C., SZIRMAY-KALOS L., ANTAL G., CSONKA F.: A simple and robust mutation strategy for the Metropolis light transport algorithm. In *Eurographics '02* (2002), pp. 531–540.
- [RDGK12] RITSCHEL T., DACHSBACHER C., GROSCH T., KAUTZ J.: The State of the Art in Interactive Global Illumination. *Computer Graphics Forum* 31, 1 (2012), 160–188.
- [Szirm08] SZIRMAY-KALOS L.: Monte-Carlo Methods in Global Illumination, VDM, Verlag Dr. Müller, 2008.
- [SKDP99] SZIRMAY-KALOS L., DORNBACH P., PURGATH-OFER W.: On the start-up bias problem of Metropolis sampling. In Winter School of Computer Graphics '99 (Plzen, Czech Republic, 1999), pp. 273–280.
- [VG97] VEACH E., GUIBAS L.: Metropolis light transport. SIG-GRAPH '97 Proceedings (1997), 65–76.

© The Eurographics Association 2013.