Report on Dagstuhl-Seminar 01242:

Stochastic Methods in Rendering

June 11 - 15, 2001

Stochastic methods have become indispensable tools in computer graphics, and more specifically in rendering, since the mid 1980s. Now these techniques are used in all subfields of rendering, and are part of every major commercial rendering package.

As with most computer graphics research, work from related fields such as radiative heat transfer and neutron transport have been modified and applied in rendering. However, once the basic technique is introduced to computer graphics, its researchers often improve on it in such a way that sometimes it is transferred back to the field from where it was borrowed, as it has happened with some radiative heat transfer methods. A second closely related topic is quasi-Monte Carlo integration and randomized quasi-Monte Carlo integration.

Although specific conferences exist separately on Monte Carlo techniques and rendering, they are both too general to give an opportunity to Monte Carlo researchers in rendering for meeting and evaluating the specific impact of these techniques on their field. The purpose of this first seminar focused on ”Stochastic Methods in Rendering” was thus to give the opportunity to evaluate the past, present and future perspectives of the use of stochastic and quasi-Monte Carlo techniques in rendering.

The workshop had to be scheduled for the same week as another seminar at Schloss Dagstuhl, and therefore the number of participants was restricted to be low (32 people from 10 countries). However this reduced number turned out ideal to stimulate real discussions. Many well established specialists came together with several young researchers, mostly PhD students. The overall highlight of the seminar was the presence of professor John Halton, a pioneer of the quasi-Monte Carlo method. It was very interesting to follow the discussions between people, once they had been made aware of tricks and hints used by others in completely different contexts. Most participants profited a lot from this exchange of normally not mentioned details in Monte Carlo and quasi-Monte Carlo algorithms.

”Stochastic Methods in Rendering” are concerned with numerically computing the integrals underlying the generation of synthetic images from digital data. These integrals very often are high-dimensional and usually the integrands are discontinuous, too. So only methods based on stochastic tools provide appropriate algorithms for the simulation of light transport. 21 talks of high quality were given on that subject. This document contains the abstracts of these talks, enabling the reader to search for more information on those topics she/he is interested in. Many talks described the application of stochastic methods in diverse algorithms, but some also described algorithms that were developed for stochastic use. A few others were rather on principles of Monte Carlo and quasi-Monte Carlo methods. Among others the following topics were discussed during and after the talks:

- theoretical comparisons of radiosity algorithms,
- utilization of random walk methods for global illumination,
- importance sampling and optimal ray distribution techniques,
- rendering animations,
- stochastic methods for real-time rendering, and
- complex natural objects.

The permanent discussion, whether random or quasi-random methods are better, resulted in a tutorial on correlated sampling by Alexander Keller, which stimulated a very interesting discussion. The main facts were that the lower bound of Monte Carlo integration is obtained using correlated sampling with the separation of the main part and the multi-level method of dependent tests, that even multiple importance sampling can arbitrarily fail due to the problem of insufficient techniques, and that correlated stratification is intrinsic to (randomized) quasi-Monte Carlo techniques. A then obvious but very important conclusion was that correlated sampling in rendering is worth further profound investigations.

As always in Dagstuhl, several activities outside the scientific field helped to recreate a little bit and encouraged people to get to know each other more. Among others we should mention the daily volleyball sessions, the

Wednesday afternoon hike through surrounding hills, several table tennis matches, and jogging and biking activities. And the organization in Dagstuhl was perfect as always, thanks to Angelika Müller and her colleagues. So at the end of the week we could look back at a very successful and relaxing seminar.

Mateu Sbert, Werner Purgathofer, Pete Shirley

Abstracts in alphabetical order by speaker's surname:

Multiple Strategy Stochastic Iteration
György Antal
joint work with László Szirmay-Kalos **5**

Stochastic Integration of Vector Fields
David C. Banks
joint work with Anuj Srivastava **5**

A Theoretical Comparison of Monte Carlo Radiosity Algorithms
Philippe Bekaert **6**

Efficient Monte Carlo Algorithms for Animated Radiosity Environments
Gonzalo Besuievsky **6**

The Fuzzy Random Walk
Francesc Castro **7**

Flame-lit Environments
Alan Chalmers **7**

Adjoints and Importance in Rendering
Per H. Christensen **7**

Scene Complexity Measures
Miquel Feixas
joint work with Jaume Rigau, Mateu Sbert, and Philippe Bekaert **7**

Why Quasi-Monte-Carlo Methods are Statistically Valid and How Their Errors Can Be Estimated Statistically
John H. Halton **8**

Beyond Monte Carlo
Alexander Keller **8**

Efficient Bidirectional Path Tracing by Randomized Quasi-Monte Carlo Integration
Thomas Kollig
joint work with Alexander Keller **8**

Perception-Guided Indirect Lighting Computation for High-Quality Animation
Karol Myszkowski **9**

Reflectance Models for Monte Carlo Ray-tracing
Attila Neumann
joint work with László Neumann **9**

Optimal Ray Shooting in Monte Carlo Radiosity
Mateu Sbert
joint work with Alex Brusi, Philippe Bekaert, Robert Tobler and Werner Purgathofer **9**

Interactive Ray-Tracing
Phillip Slusallek **9**

Path Differentials for Monte Carlo Rendering
Frank Suykens **10**

Simple and Robust Mutation Strategy for Metropolis Light Transport Algorithm
László Szirmay-Kalos
joint work with Csaba Kelemen **10**

Density Estimation on the Tangent Plane for Radiosity
Carlos Ureña-Almagro
joint work with Miguel Lastra-Leidinger **11**

Towards Interactive Global Illumination
Ingo Wald **11**

Photon Tracing for Complex Environments
Alexander Wilkie **11**

Using Monte Carlo in a Real-Time Visualization Application
Michael Wimmer **12**

Abstracts of Talks
(in alphabetical order by speaker's surname)

Multiple Strategy Stochastic Iteration

György Antal
joint work with László Szirmay-Kalos

This talk introduces a global illumination method that combines several strategies to meet the contradicting criteria of quality and speed. The methods include parallel and perspective ray-bundle shooting and ray shooting. Each method is designed to randomly approximate the effect of the light transport operator. Parallel ray-bundle tracing transfers the radiance of all points parallel to a randomly selected global direction, with perspective ray-bundles we can shoot the radiance of a single patch in all directions, and ray shooting transfers the radiance of a randomly selected point at a randomly selected direction. These strategies are of complementary character since each of them is effective in different illumination conditions. The proposed algorithm is iterative and the steps realized by different methods randomly follow each other. In each step, the applied strategy is selected randomly according to the properties of the current radiance distribution, thus we can exploit that the used strategies are good in different conditions. The formal framework of their combination is the stochastic iteration. Although the final result is the image, i.e. the algorithm is view dependent, a rough approximation of the radiance function is stored in object space, that can allow fast movements at reasonable storage requirements and also speed up Monte Carlo simulations which result in the final image. The method is also suited for interactive walkthrough animation in glossy scenes since when the viewpoint changes, only the object space radiance values remain valid and the algorithm quickly adapts to the new values.

Stochastic Integration of Vector Fields

David C. Banks
joint work with Anuj Srivastava

A streamline is a curve $c(t)$ in a vector field $V(x)$ satisfying

$$\frac{d}{dt}c(t) = V(c(t)).$$

A vector field with error or noise or diffusion can be modeled as $V = \mu + \sigma\xi$, where μ is the mean field and ξ is white noise.

We demonstrate how to integrate a stochastic vector field containing a mean component and a stochastic component, modeled as a Wiener process. We develop a Taylor expansion for the integration and demonstrate the result on different vector fields.

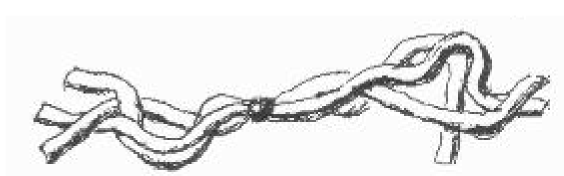

A Theoretical Comparison of Monte Carlo Radiosity Algorithms

Philippe Bekaert

The purpose of this talk is to share some of my experience and insights concerning Monte Carlo radiosity algorithms. Although the focus is on diffuse environments, much of what is said can also be applied to general light emission and scattering.

The radiosity problem can be described mathematically in two ways, leading to three basic solution strategies:

- The radiosity integral equation, an instance of the general rendering integral equation, but for diffuse light emission and scattering. This equation can be solved using so call continuous collision shooting random walks: these are called "continuous" because the states of the random walk form a continuous set, they are "collision" random walks because every collision of the random walk with surfaces in the scene is counted, and they are "shooting" random walks because they simulate the emission of particles from light sources that are propagated through the environment until they arrive in some region of interest. I show that algorithms such as the photon bucketing (histogram method), random walk algorithms for higher order approximations, the kernel method for density estimation, and a kind of final gathering algorithm differ basically in the measurement function used;

- The radiosity system of linear equations. Monte Carlo solution strategies include discrete random walks, which are very similar to continuous random walks. The states correspond to the patches in the scene and thus form a discrete set in this case. Alternatively, the radiosity system can be solved using stochastic adaptations of the Jacobi iterative method.

A rough outline is given how the variance of the corresponding Monte Carlo estimators for radiosity can be calculated. Such variance expressions can be used in order to compare algorithms, in order to derive good default parameters and heuristics for algorithms, and in order to design and study the effectiveness of variance reduction techniques for Monte Carlo radiosity.

A comparison of the variance of continuous random walks, discrete random walks and the stochastic Jacobi method learns that for an equal cost, these three basic estimators will yield to good approximation an equal accuracy in practice. The cost is measured by the number of rays shot. Minor differences are explained.

The calculation of the variance of the random walk estimators is however much more involved than the calculation of the variance of the stochastic Jacobi estimator. In my experience, not only the analysis of the stochastic Jacobi method is easier, but also the design and implementation of variance reduction techniques. Variance reduction techniques moreover often appear to be more effective as well, indicating that the stochastic Jacobi iterative method is a good alternative to random walks for further optimizing Monte Carlo rendering algorithms.

The talk is concluded with a general recipe to derive stochastic iteration algorithms corresponding with other measurement methods than bucketing. This recipe may allow new and more effective algorithms for general light emission and scattering to be developed.

Efficient Monte Carlo Algorithms for Animated Radiosity Environments

Gonzalo Besuievsky

Photorealistic based animation is known to be an expensive task when physically based global illumination is used in order to improve illumination details. In this talk we present a Monte Carlo approach to compute efficiently animated sequences in radiosity environments. We develop a method, called the Multi-Frame Lighting Method, that computes all frames of a given animated sequence in a single simulation. The method can be used both to compute the sequence or to pre-process it in order to detect, and further exploit, temporal illumination coherence. Results are completely view-independent and present significant high speed-ups showing that the technique could be an interesting alternative to deterministic methods for computing non-interactive radiosity animations for moderately complex scenarios.

The Fuzzy Random Walk

Francesc Castro

Stochastic methods are widely used in the context of radiosity. In several algorithms, random rays simulate the exchange of power between the patches in which the scene is discretized. We present a quite simple idea to reduce the aliasing in the resulting images with a very small increase of computational cost. It consists in splattering part of the power that is received by a patch to its neighboring patches. The amount of power that is distributed depends on the area of the receiver patch and also on the length of the ray segment, and the distribution of this power among the neighboring patches is related with the distance between the intersection point and every neighboring patch. Reductions of the mean square error to less than a half have been obtained in our tests.

Flame-lit Environments

Alan Chalmers

Archaeology provides an excellent application for computer graphics to explore scientific methods and problems at a human scale and to introduce a cultural dimension which opens up avenues for new and alternative interpretations. Much research has been undertaken into accurately modeling archaeological sites. Unfortunately, the luminaires used in standard modeling packages for the final renderings are based on parameters for daylight, filament bulbs or cold fluorescent tubes, rather than lamp or candle light. Flame was the predominant method for illuminating ancient sites and thus it is essential that these are modeled accurately in order to recreate the ancient lighting conditions. This is especially true as the light cast by different flame-fuel types can be significantly different which may be crucial to the perception of a reconstructed scene.

This paper describes techniques which combine experimental archaeological reconstruction of ancient lamps and candles with computer graphics and vision in order to simulate accurately flame-lit ancient environments.

Adjoints and Importance in Rendering

Per H. Christensen

This survey gives an overview of the use of importance, an adjoint of light, to speed up rendering. The importance of a light distribution indicates its contribution to an image. Importance can be used to focus global illumination and ray tracing calculations where they matter most for image accuracy. In this talk, we attempt to clarify the various uses of adjoints and importance in rendering by bringing them into a unified framework.

Scene Complexity Measures

Miquel Feixas
joint work with Jaume Rigau, Mateu Sbert, and Philippe Bekaert

In our previous work we introduced a new framework based on Information Theory for the analysis of scene visibility and radiosity complexity. Mutual information was presented as a scene complexity measure and scene entropy as the scene randomness. In this talk, these measures are reviewed emphasizing their relationship with Monte Carlo. First, we present how uniformly distributed global lines are used in order to compute scene complexity (continuous mutual information) and also their application to scene discretization. In particular, we show the Monte Carlo integration of the complexity integral. Second, we present a discussion about the relationship between the scene entropy, the Fano's inequality, and the computational Monte Carlo error of the form factors. A first step is given in the prediction of the number of lines to be sampled in order to compute the form factors to given accuracy.

Why Quasi-Monte-Carlo Methods are Statistically Valid and How Their Errors Can Be Estimated Statistically

John H. Halton

Essentially, the Monte Carlo method estimates a Lebesgue integral by sampling values of the integrand at random points in the domain of integration.

Actual use of random numbers is rare, because of the difficulty of generating reliable, consistent, truly random numbers. Consequently, we use, instead, either ”pseudo-random numbers” (PRN), which satisfy a variety of statistical tests of randomness, or ”quasi-random numbers” (QRN), which are clearly not fully ”random-seeming”, but which have strong properties of uniform distribution, which are what is required to accurately and efficiently estimate our integrals.

While QRN are empirically found to be very effective, the mathematically rigorous basis of their use has hitherto been rather mysterious. We now present a mathematical basis for understanding why their use is statistically valid.

In addition, while there are results giving mathematical upper bounds for the errors made in quasi-Monte-Carlo, and these bounds predict asymptotic behavior with large N (the sample size), the coefficients of these bounds are so hugely over-estimated as to be of little use in estimating actual errors. On the other hand, the use of the customary statistical estimators of error (sample standard deviations) is not justified when we use QRN, and also yields over-estimates of error. However, we present a new approach which has been found to lead to much more accurate error-estimates of the statistical type.

Beyond Monte Carlo

Alexander Keller

While the Monte Carlo method of integration (MC) is independent of the dimension and smoothness of the integrand, its convergence is rather slow. Abandoning the independence of the random samples and using correlated deterministic low-discrepancy sequences for sampling, the quasi-Monte Carlo method (QMC) provides much faster convergence. However then the possibility of a cheap an efficient error estimate is lost. Randomized quasi-Monte Carlo integration (RQMC) allows to almost preserve the fast convergence and at the same time have an efficient error estimate.

After surveying Monte Carlo extensions of quasi-Monte Carlo integration, we present new and very simple algorithms for randomly scrambling $(0,m,s)$-nets. These randomized low-discrepancy point sets are stratified and Latin hypercube samples at the same time, including even more stratification properties imposed by the $(0,m,s)$-net structure. Examples of the efficient application of such sampling patterns to core sampling problems of computer graphics are given. Due to their stratification properties these patterns cannot perform worse than standard stratified sampling patterns, and in practice perform much better due to their low-discrepancy. Finally we indicate the benefits of derandomized randomized quasi-Monte Carlo methods (DRQMC).

Efficient Bidirectional Path Tracing by Randomized Quasi-Monte Carlo Integration

Thomas Kollig
joint work with Alexander Keller

As opposed to Monte Carlo integration the quasi-Monte Carlo method does not allow for an error estimate from the samples used for the integral approximation. In addition the deterministic error bound of quasi-Monte Carlo integration is not accessible in the setting of computer graphics, since usually the integrands are of unbounded variation. The application of randomized quasi-Monte Carlo integration allows to exploit low-discrepancy sampling and at the same time we can estimate the variance. The structure of the high dimensional functionals computed for photorealistic image synthesis indicates different approaches to apply randomized quasi-Monte Carlo integration to bidirectional path tracing. The resulting techniques are simple to implement and much more efficient than previous bidirectional path tracing algorithms.

Perception-Guided Indirect Lighting Computation for High-Quality Animation

Karol Myszkowski

We present a method for efficient global illumination computation in dynamic environments by taking advantage of temporal coherence of lighting distribution. The method is embedded in the framework of stochastic photon tracing and density estimation techniques. A locally operating energy-based error metric is used to prevent photon processing in the temporal domain for the scene regions in which lighting distribution changes rapidly. A perception-based error metric suitable for animation is used to keep noise inherent in stochastic methods below the sensitivity level of the human observer. As a result a perceptually-consistent quality across all animation frames is obtained. Furthermore, the computation cost is reduced compared to the traditional approaches operating solely in the spatial domain.

Reflectance Models for Monte Carlo Ray-tracing

Attila Neumann
joint work with László Neumann

The appearance and computational cost of a picture created with distributed or Monte-Carlo ray-tracing depend on the applied reflectance models. The well-known and used BRDF formulas have several problems. Often they do not fulfill the energy conservation criterion like the most of metal models at grazing angles. The appearance of Phong and Blinn models seems non-metallic and artificial. The central problem is to generate refracted rays following the energy distribution in recursive ray-tracing method. The presented reflectance models can be used for importance sampling very easily. However, the formulas are short and simple and the appearance is adjustable from retro-reflective character until metals with off-specular peak.

Optimal Ray Shooting in Monte Carlo Radiosity

Mateu Sbert
joint work with Alex Brusi, Philippe Bekaert, Robert Tobler and Werner Purgathofer

In this talk we deal with the Monte Carlo radiosity problem of how to aim more rays towards where they are needed more: small patches and high reflectivity regions. We derive the optimal ray shooting probabilities when taking the Mean Square Error as the error measure. The standard case of using the Form Factors is shown to be optimal in the particular case of equal relation between reflectivity and square root of patch area throughout the scene. We also present an adaptive algorithm to compute optimal probabilities. Examples showing optimal direct illumination distribution are given, exhibiting concordance with the theoretically predicted results.

Interactive Ray-Tracing

Phillip Slusallek

For almost two decades researchers have argued that ray tracing will eventually become faster than the rasterization technique that completely dominates todays graphics hardware. However, this has not happened yet. Ray tracing is still exclusively being used for off-line rendering of photorealistic images and it is commonly believed that ray tracing is simply too costly to ever challenge rasterization-based algorithms for interactive use. However, there is hardly any scientific analysis that supports either point of view. In particular there is no evidence of where the crossover point might be, at which ray tracing would eventually become faster, or if such a point does exist at all.

This talk provides several contributions to this discussion: We first present a highly optimized implementation of a ray tracer that improves performance by more than an order of magnitude compared to currently available ray tracers. The new algorithm makes better use of computational resources such as caches and SIMD instructions and better exploits image and object space coherence. Secondly, we show that this software implementation can challenge and even outperform high-end graphics hardware in interactive rendering performance for complex environments. We also provide an brief overview of the benefits of ray tracing over rasterization algorithms and point out the potential of interactive ray tracing both in hardware and software.

Path Differentials for Monte Carlo Rendering

Frank Suykens

Photo-realistic rendering algorithms such as Monte Carlo ray tracing sample individual paths to compute images. Noise and aliasing artifacts are usually reduced by supersampling. Knowledge about the neighborhood of the path, such as an estimated footprint, can be used to reduce these artifacts without having to trace additional paths. Ray differentials estimate such a footprint for classical ray tracing, by computing ray derivatives with respect to the image plane. The footprint proves useful for filtering textures locally on surfaces. Path differentials generalize this idea to arbitrary Monte Carlo path sampling, including general reflection and refraction functions. Sampling new directions introduces additional partial derivatives, which can all be combined into a footprint estimate. Additionally the path gradient is introduced; it gives the rate of change of the path contribution. When this change is too steep the size of the footprint is reduced. The resulting footprint can be used in any global illumination algorithm that is based on path sampling. We have developed two applications to show its potential: texture filtering in distributed ray tracing and a new approach to hierarchical refinement in particle tracing radiosity.

Simple and Robust Mutation Strategy for Metropolis Light Transport Algorithm

László Szirmay-Kalos
joint work with Csaba Kelemen

The paper presents a new mutation strategy for the Metropolis light transport algorithm, which works in the space of uniform random numbers used to build up paths and transforms the points to the path space according to BRDF sampling, light source sampling and Russian roulette. This transformation makes the integrand and the importance function flatter and thus increases the acceptance of the new samples. Higher acceptance ratio, in turn, reduces the correlation of the samples, which increases the speed of convergence. When mutations are calculated, a new random point is selected usually in the neighborhood of the previous one, but according to our proposition called "large steps", sometimes an independent point is obtained. Large steps greatly reduce the start-up bias and guarantee the ergodicity of the process. Due to the fact that some samples are generated independently of the previous sample, this method can also be considered as a combination of the Metropolis algorithm with a classical random walk. In order to keep the merits of both approaches, we use multiple importance sampling to combine their results, that is, the combined method will be as good at bright regions as Metropolis and at dark regions as random walks. The resulting scheme is robust, efficient, but most importantly, is easy to implement and to combine with an available random walk algorithm.

Density Estimation on the Tangent Plane for Radiosity

Carlos Ureña-Almagro
joint work with Miguel Lastra-Leidinger

In this talk we describe a new technique for computing approximations to radiosity in the context of photon-tracking algorithms. Previous techniques (such as light-maps or photon-maps) show low performance in the presence of either small patches or small objects, because these entities have a low probability to receive any photon hit.

We propose a method which is not based on hits in surfaces but rather on the local density of photons traveling near the point where we wish to estimate radiosity. Concretely we do kernel density estimation by using a plane tangent to surface at target point. In this context, we describe a spatial index structure we use to quickly find which lines go through any given planar circle, which is the essential operation in our algorithm.

This new method can be applied to scenes described by using a wide range of types of geometric models. The only requirement we impose to geometric models is the ability to compute intersections with rays and then obtain the normal at intersection point. Other advantage is in the possibility to select in advance the desired precision, independent of the complexity of the geometric model. We have implemented and tested this method for polygonal meshes, with promising results.

Towards Interactive Global Illumination

Ingo Wald

Many disciplines must handle the creation, visualization and manipulation of huge and complex 3D environments. Examples include large structural and mechanical engineering projects dealing with entire cars, ships, buildings, and processing plants. The complexity of such models is usually far beyond the interactive rendering capabilities of todays 3D graphics hardware.

In this talk we show that using a highly optimized software ray tracer we are able to achieve interactive rendering performance for highly complex models of up to 50 million triangles including reflection and shadow computations. Interactivity is achieved with a novel approach to distributed rendering based on coherent ray tracing.

Due to the drastically reduced time to trace rays, we argue that even interactive global illumination finally comes into reach. However, many of todays commonly used global illumination algorithms do not efficiently map to such an interactive, distributed environment. Therefore, we end this talk with a discussion on requirements for future interactive global illumination algorithms: We have to search for algorithms that are highly efficient, generate rays coherently, work well in parallel, require little communication, and tolerate latency.

Photon Tracing for Complex Environments

Alexander Wilkie

We present a method that makes the use of photon tracing methods feasible for complex scenes when a totally accurate solution is not essential. This is accomplished by using *orientation lightmaps*, which average the illumination of complex objects depending on the surface normal. Through this averaging, they considerably reduce the variance of the stochastic solution. In order to use these specialized lightmaps, which consume comparatively small amounts of memory, no changes have to be made to the basic photon tracing algorithm. Also, they can be freely mixed with normal lightmaps. This gives the user good control over the amount of inaccuracy he introduces by their application. The area computations necessary for their insertion are performed using a stochastic sampling method that performs well for highly complex objects.

Using Monte Carlo in a Real-Time Visualization Application

Michael Wimmer

This talk is intended to show that stochastic methods can be applied to some degree in real-time visualization. Large virtual environments usually face two problems: rendering time because of the overall complexity of the model, and aliasing issues due to the complexity of distant objects projected to the screen. This talk introduces a new primitive, point-based impostors, that remedy the problems described. The new representation depends on a sufficiently dense geometric sampling of a distant object (so as to prevent holes from appearing in any re-projection) and calculates appearance in the form of a small texture map for each sample point. Appearance is calculated via a 6-dimensional Monte Carlo integral that accomodates possible observer movements within a view cell, view-dependent effects, arbitrary camera orientations and the monitor reconstruction function. Point-based impostors provide a fast, antialiased representation for distant objects viewed from a view cell. After the impostor is calculated in a preprocessing step, it can be efficiently rendered with current graphics hardware.