Project-Team FRVSense

Fast Rendering and Visualization Sense

Rennes

| *Activity Report*

2013

1 Team

Head of the team
Kadi Bouatouch, Professor, University of Rennes 1

Administrative assistant
Maryse Fouché, INRIA

University of Rennes 1 personnel
Rémi Cozot, Assistant Professor
Mickael Ribardiere, Research Engineer

External collaborators
Christian Bouville, External Collaborator

PhD students
Adrien Gruson, *MENRT* grant since October 2011
Ronan Boitard, *CIFRE* grant since October 2011
Mohamed Djebbar, *Algerian Government* grant, since February 2011
Ricardo Marques, *CORDI* grant, since October 2010, defense in October 2013
Mahmoud Omidvar, *CSTB* grant, since October 2010
Billal Merabti, *Algerian Government* grant, since October 2012
Matis Hudon *CIFRE* grant since November 2013

2 Overall Objectives

2.1 Objectives

The different topics that will be addressed by the team are:

1. Real-time rendering and global illumination algorithms exploiting the performances of the new GPUs (Graphics Processing Units);
2. Color and visual perception;
   - Color appearance models
   - HDR and tone mapping
   - Multi-view rendering on stereoscopic and auto-stereoscopic display devices;

3. Inverse lighting and image aesthetics

Our objectives for the these topics are:

- Real-time rendering and global illumination
  The main objective is real-time global illumination computation. This is a real challenge that want to take up all the researchers involved in 3D movie and game video as well as virtual reality applications. We want to make available to artists a tool for photorealistic real-time rendering. Regarding cinematographic relighting, our goal is to devise an algorithm for relighting images of an animation sequence. Changing the parameters of light sources and playing back the animation could be done interactively. Following the work of Haÿsan et al. [MEFK08] we want to formalize the problem of animation relighting using a transfer tensor. We want to propose an adaptive method for the tensor evaluation and a practical optimization of the clustered principal component analysis (CPCA) used for tensor compression.

  For global illumination we have proposed another integration procedure known as Bayesian Monte Carlo integration (BMC)[O'H87]. As opposed to importance sampling methods, BMC does not require the samples to be drawn from a predefined pdf and does take into account samples location. The price to pay for this is a higher computational cost per samples but in a recent paper[BBL+09], we have shown that this cost can be kept to the same level as for importance sampling with a suboptimal solution. Compared to the usual combination of importance and stratified sampling procedures, our proposed BMC method allows a 23% decrease in the number of samples for the same noise level. However, these results have only been obtained for the computation of the diffuse component in the final gathering step of a photon mapping renderer and cannot be easily extended to the implementation of a full global illumination renderer. Our approach will allow to analyze for which rendering step or for which effect the use of BMC could be beneficial and then experiment solutions. A first objective will be to deal with the rendering of the specular component and experiment this solution for real time environment map rendering. An important aspect of this work will be the integration of prior knowledge. We are also investigating solutions for spatio-temporal sampling and interpolation as mentioned above. We think that BMC can provide a very powerful framework in which both procedures of approximation and integration can be merged, thus leading to very efficient rendering methods.

- Color and visual perception

  – Color appearance models and tone mapping
    Our main objective is to propose approaches allowing to combine global illumination

[MEFK08] H. M., V.-A. E., P. F., B. K., "Tensor Clustering for Rendering Many-Light Animations", *Computer Graphics Forum 27*, 4, 2008.

[O'H87] A. O'Hagan, "Monte-Carlo is fundamentally unsound", *The Statistician 36*, 2/3, 1987, p. 247–249.

[BBL+09] J. Brouillat, C. Bouville, B. Loos, C. Hansen, K. Bouatouch, "A Bayesian Monte Carlo Approach to Global Illumination", *Computer Graphics Forum 28*, 8, 2009, p. 2315–2329, http://dx.doi.org/10.1111/j.1467-8659.2009.01537.x.

engines and visual perception based algorithms in order to compute perceptually realistic images. We have already proposed a robust and real-time model of visual attention which mixes top-down and botom-up approaches in the context virtual reality[HBO+10,HLC+10]. We have also proposed and evaluated methods to enhance visual feedback by adding a depth of field blur in VR applications[HLCC08]. In the short-term, our work will focus onto three topics:

* the creation of a robust and real-time chromatic adaptation algorithm (estimate of white balance) in the context of global illumination;
* the enhancement of tone mapping algorithm using our visual attention model;
* the proposal of new tone mapping techniques for dynamic scenes;
* the study of new methods for improving visual comfort and perception in the specific context of auto-stereoscopic devices.

Once we will be able to evaluate the perceptual realism of the rendered images, then in the long-term, one challenge will be to inverse the perceptual models and let the artist/designer specify some constraints on image appearance and perception. The rendering parameters will be tuned to satisfy these constraints. Note that a high level model of constraints can be used by an artist to enhance the visual perception of 3D scenes.

– Multi-view rendering
Our goal in this project is to focus our research on the problem of real time multiview rendering of 3D scenes. The real time constraint is particularly challenging when it comes to rendering 256 views with a global illumination model. Such performance cannot be obtained without some pre-computation, which impacts visualization interactivity because some scene model characteristics (object attributes, light source position, etc.) will not be allowed to change and the navigation will be restricted.

2.2 Key Issues

3 Scientific Foundations

3.1 Global Illumination

Keywords: Global illumination, relighting, bayesian Monte Carlo, multi-view rendering, image-based rendering, color appearance model, high dynamic range.

To achieve realistic rendering, one has to compute global illumination accounting for: multiple reflections and refractions, sub-surface scattering (within skin, marble), multiple scattering

[HBO+10] S. Hillaire, G. Breton, N. Ouarti, R. Cozot, A. Lecuyer, "Using a Visual Attention Model to Improve Gaze Tracking Systems in Interactive 3D Applications", *Computer Graphics Forum*, 2010.

[HLC+10] S. Hillaire, A. Lecuyer, R. Corte, T., R. Cozot, G. Breton, "A Real-Time Visual Attention Model for Predicting Gaze Point During First-Person Exploration of Virtual Environments", *in : VRST 2010*, 2010.

[HLCC08] S. Hillaire, A. Lecuyer, R. Cozot, G. Casiez, "Depth-of-Field Blur Effects for First-Person Navigation in Virtual Environments", *Computer Graphics and Applications*, 2008.

in participating media (smoke, fire, dust, clouds), chromatic adaptation, tone mapping, etc. It is well known that global illumination computation is demanding in terms of computing and memory resources. This is why the team members have already proposed new solutions based on irradiance/radiance caching methods implemented on the GPU to speed up the different computations involved in global illumination and rendering algorithms[GBP07a,KGPB05].

Radiance caching has also been used to simulate multiple scattering within participating media[RCB11]. Indeed, we proposed a technique for efficiently rendering participating media with irradiance caching scheme using Monte Carlo ray tracing. Our approach extends the original algorithm for surfaces to volumes rendering. Our method allows us to adjust density of cached records according to illumination changes and to store both direct and indirect contributions in the records but also multiple scattering due to medium. To achieve this, we proposed an adaptive shape for records depending on geometrical features and irradiance variations. In order to avoid a too high density of cached records where it is not necessary (for example away from the observer inside a medium with high absorption and scattering properties) a new method is proposed to control the density of the cache during the addition of new records. This volume density control is depending of the interpolation quality with the existing records and the photometric characteristics of the participating media. Limiting the number of records in the volume can accelerate both the computation pass (computation of irradiance and gradients may be time-consuming for a record) and the final rendering pass by limiting the number of searches in the cache. Moreover, storing all contributions in records is of high importance, especially in the case of scenes having many light sources with complex geometry as well as wide surfaces exposed to daylight or again for animation rendering. Finally, instead of using an expensive ray marching to find records that cover the ray, we gather all the records along ray which contribute to the in-scattered radiance. With these different techniques, pre-computing and rendering passes are significantly speeded-up.

We have brought several contributions to global illumination: radiance cache, participating media, Metropolis Light Transport, etc. Moreover, we proposed a novel approach to progressive photon mapping for scenes containing surface and volume objects. Our method allows to handle scenes with specular and refractive objects as well as homogeneous and heterogeneous participating media.

3.2 Bayesian and Quasi Monte-Carlo methods for image rendering

Keywords: Sampling, likelihood principle, Monte Carlo.

Most global illumination rendering methods are based Monte-Carlo importance sampling but, as pointed out by OÕHagan in 1987[O'H87], this procedure suffers from two inconsistencies

[GBP07a] P. Gautron, K. Bouatouch, S. Pattanaik, "Temporal Radiance Caching", *IEEE Transactions on Visualization and Computer Graphics 12*, 5, 2007.

[KGPB05] J. KrivÃ¡nek, P. Gautron, S. Pattanaik, K. Bouatouch, "Radiance Caching for Efficient Global Illumination Computation", *IEEE Transactions on Visualization and Computer Graphics 11*, 5, September-October 2005, p. 550–561.

[RCB11] M. Ribardiere, S. CarrÃ©, K. Bouatouch, "Adaptive Records for Volume Irradiance Caching", *in : Pacific Graphics'2011*, 2011.

[O'H87] A. O'Hagan, "Monte-Carlo is fundamentally unsound", *The Statistician 36*, 2/3, 1987, p. 247–249.

that we will briefly present in the following. Firstly, the estimator depends on some arbitrary choice of the sampling density which violates the Likehood principle. Secondly, this procedure ignores sample location and thus wastes important information. Although solution exists to alleviate these problems, it still remains true that choosing the right sample distribution is a difficult problem and that Monte Carlo importance sampling does not allow to easily take into account prior knowledge on the data. This leads to a waste of samples which is further increased by the fact that samples location is not exploited. Furthermore, samples reuse from one integral computation to another is difficult since samples must be drawn from a predefined probability density function (pdf).

QMC integration is now extensively used in computer graphics (see e.g. [SEB08]. Keller has shown in [Kel13] that QMC techniques can be applied in a consistent way to deal with a wide range of problems (anti-aliasing, depth of field, motion blur, spectral rendering, etc.). However, few applications have been reported in the literature specifically addressing hemispherical sampling with a view of computing the illumination integral. Unlike the unit square sampling case, no explicit construction of optimal point sets for spherical sampling is known and generally the spherical point sets are generated by lifting point sets from the unit square to the unit sphere through an equal-area transform. Although such point constructions are not proved to be optimal, recent results from the numerical analysis literature suggest that both (0, 2)-sequences and Fibonacci lattices lifted to the sphere are quite close to optimality in terms of discrepancy [ABD12].

Nevertheless their performance is not exactly equivalent: several authors have shown that spherical Fibonacci (SF) lattices are particularly well-suited to sphere sampling compared to other low-discrepancy point sets. Furthermore, similar point structures arise spontaneously in nature so as to implement a best packing strategy on the sphere (e.g. packing of seeds in the sunflowers head [Vog79]), a clear indication that these structures have intrinsically good spherical uniformity properties. We introduced theoretical aspects on QMC spherical integration that have never been used in the graphics community. In concrete terms, we define worst case integration error (w.c.e.), spherical cap discrepancy (s.c.d.) and an inter-samples distance-based energy metrics, which allows to assess the quality of a spherical samples set for spherical integration.

[SEB08] P. Shirley, D. Edwards, S. Boulos, "Monte Carlo and Quasi-Monte Carlo Methods for Computer Graphics", *in: Monte Carlo and Quasi-Monte Carlo Methods 2006*, A. Keller, S. Heinrich, and H. Niederreiter (editors), Springer Berlin Heidelberg, 2008, p. 167–177, http://dx.doi.org/10.1007/978-3-540-74496-2_8.

[Kel13] A. Keller, "Quasi-Monte Carlo Image Synthesis in a Nutshell", *in: Monte Carlo and Quasi-Monte Carlo Methods 2012*, J. Dick, F. Y. Kuo, G. W. Peters, and I. H. Sloan (editors), *Springer Proceedings in Mathematics and Statistics*, *65*, Springer Berlin Heidelberg, 2013, p. 213–249, http://dx.doi.org/10.1007/978-3-642-41095-6_8.

[ABD12] C. Aistleitner, J. Brauchart, J. Dick, "Point Sets on the Sphere with Small Spherical Cap Discrepancy", *Discrete and Computational Geometry 48*, 4, 2012, p. 990–1024, http://dx.doi.org/10.1007/s00454-012-9451-3.

[Vog79] H. Vogel, "A better way to construct the sunflower head", *Mathematical Biosciences 44*, 3-4, 1979, p. 179–189, http://www.sciencedirect.com/science/article/pii/0025556479900804.

3.3 Color, Visual Perception and Inverse Lighting

Keywords: Color appearence model, visual perception, tone mapping, white balancing, high dynamic range.

Color appearance models and visual perception Global illumination algorithms provide, as results, physically based values that need to be transformed into displayable RGB images. This transformation could seem straightforward. But the human visual system does not perform a direct transformation, rather it involves some local and global adaptation operations, takes into account the viewing conditions and some cognitive tasks. To give the same visual perception as that of a user immersed in a 3D virtual world, this transformation into images of these physical values must take into account the features of the human visual system (HVS) as well as the characteristics of the display device. The main key features of the human visual system to be accounted for are the luminance adaptation, the chromatic adaptation, the focus point (at which the user is looking) and the depth of field. On the display device side, we have to take into account the luminance range, the color range, the size of the screen and if necessary the stereo features.

Chromatic adaptation consists in recovering the color appearance of objects as if the user was within a 3D scene (as if he sees directly the scene). But color appearance for the human visual system does not depend only on the spectral radiance of the stimulus, but it also takes into account the viewing conditions. Color appearance models allow to model the perception of color stimulus according to the viewing conditions through colorimetric and photometric quantities such as: brightness, lightness, chroma, saturation, colorfulness, and hue of the stimulus [KJF07,Fai05]. Many woks also deal with chromatic adaptation. They belong to different research fields: computer graphics, color science, computer vision, etc [Fai91b,Fai91a,VK70,WW09]. Colors are perceived as constant even though the illuminant color changes. Indeed, the perceived color of a diffuse white sheet of paper is still white even though it is illuminated by a single orange tungsten light, whereas it is orange from a physical point of view. Unfortunately global illumination algorithms only focus on the physics aspects of light transport. The output of a global illumination engine is an image which has to undergo chromatic adaptation to recover the color as perceived by the HVS. We proposed a new color adaptation method well suited to global illumination. This method estimates the adaptation color by averaging the irradiance color arriving at the eye. Unlike other existing methods, our approach is not limited to the view frustum, as it considers the illumination from all the scene. Experiments

[KJF07] J. Kuang, G. M. Johnson, M. D. Fairchild, "iCAM06: A refined image appearance model for HDR image rendering", *Journal of Visual Communication*, 2007.

[Fai05] M. D. Fairchild, *Color Appearance Model, Second edition*, John Wiley and sons, Ltd, 2005.

[Fai91b] M. D. Fairchild, "A model of incomplete chromatic adaptation", *in : the 22nd Session of the CIE*, CIE, p. 33–34, 1991.

[Fai91a] M. D. Fairchild, "Formulation and testing of an incomplete-chromatic-adaptation model", *Color Research Application 16*, 1991, p. 243–250.

[VK70] J. Von Kries, *Chromatic adaptation*, MIT Press, [1905] 1970, p. 109–119.

[WW09] A. Wilkie, A. Weidlich, "A Robust Illumination Estimate for Chromatic Adaptation in Rendered Images", *in : Eurographics Symposium on Rendering 2009*, 2009.

have shown that our method outperforms the state of the art methods.

The conversion of the HDR (High Dynamic Range) luminance range of the images (computed using a global illumination algorithm) to LDR (Low Dynamic Range) images displayable onto LDR display devices is called Tone Mapping. This topic has been widely addressed by the computer graphics community. A complete state of the art of tone mapping algorithms can be found in [RWD+10]. While many solutions have been designed over the last decade, only few of them can cope with video sequences. Indeed, these TMOs tone map each frame of a video sequence separately, which results in temporal incoherency. Two main types of temporal incoherency are usually considered: flickering artifacts and temporal brightness incoherency. While the reduction of flickering artifacts has been well studied, less work has been performed on brightness incoherency. We proposed a method that aims at preserving spatio-temporal brightness coherency when tone mapping video sequences. Our technique computes HDR video zones which are constant throughout a sequence, based on the luminance of each pixel. Our method aims at preserving the brightness coherency between the brightest zone of the video and each other zone. This technique adapts to any TMO and results show that it preserves well spatio-temporal brightness coherency. We validate our method using a subjective evaluation. In addition, unlike local TMOs, our method, when applied to still images, is capable of ensuring spatial brightness coherency. Finally, it also preserves video fade effects commonly used in post-production.

Inverse Lighting As many phenomena are accounted for by global illumination algorithms, it is difficult to know beforehand what the resulting image will look like. For instance, we would like an object of a scene to be bright, with a lot of contrast, while we would like the background of the scene to be dark and uniform. Answering such specifications can be a time-consuming and tedious effort if it is not done automatically. Thus, a trial-and-error approach is not satisfactory. To obtain a satisfactory result automatically, the user has to express his aesthetic intent by providing a set of target values for properties of the final image. The goal is to change iteratively the parameters of the 3D scene to fit the specification. There are a lot of parameters to take into account. Among them are: position, size, luminance of the light sources, reflectance of the materials, distribution of the objects in the scene. The methods used to compute these lighting parameters to fit a set of constraints are called inverse lighting methods. In these methods, an objective function that accounts for the target parameters is defined to express the intent of the user. The objective function expresses the distance between the current lighting design and the desired result. To find the optimal set of parameters, an objective function has to be minimized. There are many ways to express the aesthetic intent, as there are many possible aesthetics.

We designed a new inverse lighting framework. Given a set of target parameters provided by the user, we set up a platform that can render a scene with global illumination techniques and allows to minimize the objective function to find the desired parameters (light source size, position, luminance) to meet the user's intent.

[RWD+10] E. Reinhard, G. Ward, P. Debevec, S. Pattanaik, W. Heidrich, K. Myszkowsk, *High Dynamic Range Imaging*, Morgan Kaufmann Publishers, San Francisco, USA, 2010.

Multiview Rendering Since the outset of research in global illumination, most of the research works in this domain has focused on the rendering of a single image at a time. But recently, there has been a growing interest in the problem of rendering several views of a scene taken at closed distances in the spatio-temporal domain while efficiently exploiting image coherence. This is mainly due to recent developments of two well-known research domains of computer graphics and virtual reality:

- 3D stereoscopic visualization is booming again with the advent of 3DTV at home. 3D displays are becoming affordable for consumer mass market with an offering ranging from 3-D LCDs requiring glasses to multiview autostereoscopic displays. From the computational point of view, stereoscopic visualization only requires the rendering of two views, which can be sustained by most state-of-the-art graphic boards. However, to avoid the discomfort caused by glasses, autostereoscopic technologies have aroused a considerable research effort and the commercial offering is constantly expanding. These technologies allow glass-free stereoscopic visualization in an angular range of up to 360¡ but they require the rendering of a much greater number of views than basic two-view stereo displays[BWS+07]. Presently commercialized autostereoscopic displays use either a parallax barrier or a lenticular technology with a 9-view limit in either case; however some research prototypes[DML+01] can display as much as 256 views[TN10], which requires specific rendering techniques.

- Network-based 3D graphics and especially Web-based applications have regained interest these last years for various reasons: ever increasing 3D data complexity that measures in Gigabytes today, better protection of property rights on secure servers, exponential growth of services for light mobile clients with limited resources, emergence of Cloud Computing services, etc. Network-based 3D applications may have very different requirements in terms of real time, navigation, interactivity, rendering quality and 3D data complexity. However, the used rendering techniques must make the most of spatio-temporal coherence through appropriate pre-computed data structures so as to reduce network latency effects and to cope with limited computational resources.

Though both problems cannot be solved without considering hardware architecture specificities, they raise the same fundamental issue of fast and efficient computation of closed viewpoints in the spatio-temporal domain. This research topic is not new in computer graphics and the problem of exploiting spatio-temporal coherence to speed up computation time has been extensively addressed in the literature e.g.[GBP07b]. However, with the constraints of

[BWS+07] P. Benzie, J. Watson, P. Surman, I. Rakkolainen, K. Hopf, H. Urey, V. Sainov, C. von Kopylow, "A Survey of 3DTV Displays: Techniques and Technologies", *IEEE Transactionns on Circuits and Systems for Video Technology 17*, 11, 2007, p. 1647–1658.

[DML+01] N. A. Dodgson, J. R. Moore, S. R. Lang, G. Martin, P. Canepa, "A 50" Time-Multiplexed Autostereoscopic Display", *in : Proc. SPIE Stereoscopic Displays & Applications XI*, *3957*, SPIE, p. 177–183, January 2001.

[TN10] Y. Takaki, N. Nago, "Multi-projection of lenticular displays to construct a 256-view super multi-view display.", *Opt Express 18*, 9, Apr 2010, p. 8824–8835.

[GBP07b] P. Gautron, K. Bouatouch, S. Pattanaik, "Temporal Radiance Caching", *IEEE Transactions on Visualization and Computer Graphics 13*, 5, 2007, p. 891–901.

the applications mentioned above, the complexity and the dimensionality of the problem has considerably changed, which leads us to re-examine the problem in depth.

3.4 Automatic computation of cinematographic sequences

Keywords: 3D Graphics, Realism, Animation, Cinematography.

Automatically computing a cinematographically consistent sequence of shots over a set of actions occurring in a 3D world is a complex task which requires not only the computation of appropriate shots (viewpoints) and appropriate transitions between shots (cuts), but the ability to encode and reproduce elements of cinematographic style. Models proposed in the literature, generally base on finite state machine representations, provide limited functionalities to build sequences of shots and are no designed in mind to generate sequences reproducing a given cinematographic style, nor do they allow to perform significant variations in style over the same sequence of actions. In collaboration with Marc Christie (member of the Mimetic team at INRIA Rennes Bretagne Atlantique) we first proposed an expressive automated cinematography model that can compute significant variations in terms of cinematographic style, with the ability to control the duration of shots and the possibility to add specific constraints on the desired sequence. Second, we parametrized the model in a way that facilitates the application of learning techniques so as to reproduce elements of style extracted from real movies by using a Hidden Markov Model representation of the editing process. The proposed model is both more general than existing representations, and proves to be more expressive in its ability to precisely encode elements of cinematographic style. Results comparing our model with state-of-the-art finite state machine representations illustrate these features. The next step is to add some lighting constraints which would influence the determination of shot sequences.

4 Application Domains

4.1 Applications

Participants: all participants.

The applications of our works are:

- Lighting simulation: fast and physics-based;
- High quality video games: real-time, high realism, perception-based, help game designers to choose lighting and material properties interactively, stereoscopic and autostereoscopic;
- 3D movies: interactivity, high photorealism, perception-based, help for artists to choose lighting and material properties interactively, stereoscopic and autostereoscopic;
- Virtual and augmented reality: high realism, fast, stereoscopic and autostereoscopic.

5 New Results

5.1 Camera placement and Cinematography

Participants: Billal Merabti, Rémi Cozot, Kadi Bouatouch.

The important advances in the quality and performances of 3D rendering allow to create 3D animated sequences in real-time with the aim to generate in real-time movies of cinematographic quality. However, the integration of 3D cinematographic elements (camera placement, light source placement, actor placement) requires the formalization of a knowledge that exists empirically in industry and in the literature. The approach that we adopt is to study, from real data (movies), the combination of framing, lighting and placement of actors, given elements of cinematographic style (duration of framings, their frequency, transition between framings, etc.) Our approach will provide an efficient computational model allowing the placement of actors, camera and light sources, with the aim of automatic staging satisfying elements of specified styles.

In collaboration with Marc Christie (member of the Mimetic team at INRIA Rennes Bretagne Atlantique) we first proposed an expressive automated cinematography model that can compute significant variations in terms of cinematographic style, with the ability to control the duration of shots and the possibility to add specific constraints on the desired sequence. Second, we parametrized the model in a way that facilitates the application of learning techniques so as to reproduce elements of style extracted from real movies by using a Hidden Markov Model representation of the editing process. The proposed model is both more general than existing representations, and proves to be more expressive in its ability to precisely encode elements of cinematographic style. Results comparing our model with state-of-the-art finite state machine representations illustrate these features. The next step is to add some lighting constraints (expertise of the FRVSense team) which would influence the determination of shot sequences.

This work is described in a paper that will be submitted to a journal very shortly.

5.2 Global Illumination

Participants: Adrien Gruson, Remi Cozot, Mickael Ribardiere, Kadi Bouatouch.

We designed and implemented a global illumination method that is capable to handle different kinds of surface materials (specular, glossy, transparent) as well as participating media.

When rendering surfaces, the Progressive Photon Mapping method assigns a disk to every visible point lying on a surface. For volumes, we use beams rather than disks. To each ray of a ray path, cast through a pixel and traversing a medium, we assign a cylinder whose radius gets smaller after each pass. This cylinder is called beam from now on. These beams get reflected/refracted if the associated ray hits a specular or refracting surface. We build a beam Kd-tree whose leaves are the resulting beams. When a ray hits a diffuse or a glossy surface, directly or indirectly through specular reflection or refraction, the resulting hit point is registered and assigned a disk-shaped influence zone. The radius of this disk gets smaller

after each pass.We build another Kd-tree, called view Kd-Tree, whose leaves are those hit points.We also precompute a 2D array storing ray paths traced from the viewpoint through each pixel, as explained later. To make our method computationally efficient, we decided to perform this preprocessing every N passes. For every new N passes, we jitter for each pixel its associated view ray then perform again the preprocessing. This allows to solve the aliasing artifacts with the help of oversampling. At each pass, a set of photons are shot. Whenever one of those photons interacts with a volume (or a surface), the top-down traversal of the associated Kd-tree leads to beams (or visible points), which allows to compute the current radiance, due to the photon, of the hit points and the beams. Then, the photon is discarded instead of stored in a photon map. Next, these current radiances are updated to account for all the successive shot photons. Once all the photons have been shot, the image is computed. For more efficiency, every contributive light path is mutated using Metropolis. A light path is considered as contributive if one of its vertex contributes to visible points or beams. This increases the efficient of the shooting process. To speed up the traversal of the view Kd-tree, we use a same radius for all disks. We also use a same radius for all beams for the same reason. This slightly affects the efficient of the method but the method keeps converging to an unbiased result. Finally, at the end of each pass, the resulting image is updated, and so are the radii of the disks and beams.

This work has been published in The Visual Computer Journal [2] and presented at the CGI conference [6].

5.3 Bayesian and Quasi Monte Carlo

Participants: Ricardo Marques, Christian Bouville, Kadi Bouatouch.

Bayesian Monte Carlo (BMC) techniques are widely used in the domain of Machine Learning, and relies on priors over the function of interest to improve Monte Carlo computations. We have used BMC integration to speed-up the final gathering operation needed in global illumination computation but only for diffuse objects. We have extended BMC to glossy surfaces. This work has been published in the IEEE TVCG journal [4] .

Quasi-Monte Carlo (QMC) methods exhibit a faster convergence rate than that of classic Monte Carlo methods. This feature has made QMC prevalent in image synthesis, where it is frequently used for approximating the value of spherical integrals (e.g. illumination integral). The common approach for generating QMC sampling patterns for spherical integration is to resort to unit square low-discrepancy sequences and map them to the hemisphere. However such an approach is suboptimal as these sequences do not account for the spherical topology and their discrepancy properties on the unit square are impaired by the spherical projection. In this paper we present a strategy for producing high-quality QMC sampling patterns for spherical integration by resorting to spherical Fibonacci point sets. We show that these patterns, when applied to illumination integrals, are very simple to generate and consistently outperform existing approaches, both in terms of root mean square error (RMSE) and image quality. Furthermore, only a single pattern is required to produce an image, thanks to a scrambling scheme performed directly in the spherical domain. This work has been published in the Computer Graphics Forum journal [3].

5.4 Multiview Rendering

Participants: Mohamed Djebbar, Christian Bouville, Kadi Bouatouch.

Multi-view auto-stereoscopic displays are now available at affordable cost and are getting widely used in virtual reality applications and 3D games. With their wide viewing zone, this type of display easily accommodates multiple viewers and does not require any head tracking. Real-time rendering on these displays poses a number of difficult problems: (1) real-time generation of several views of the same 3D scene, (2) choice of a particular sampling pattern of the displayed image requiring specific anti-aliasing procedures that results in a limitation of the usable depth range. We have proposed some solutions to the rendering of multiple views in real time relying on image based rendering while handling the occlusion and dis-occlusion problems using a frequency approach and light field sampling techniques.

5.5 HDR Imaging

Participants: Ronan Boitard, Mickael Ribardiere, Rémi Cozot, Kadi Bouatouch.

Several TMOs have been proposed over the last decade, from the simple global mapping to the more complex one simulating the human vision system. While these solutions work generally well for still pictures, they are usually less efficient for video sequences as they are source of visual artifacts. Only few of them can be adapted to cope with a sequence of images. In this paper we present a major problem that a static TMO usually encounters while dealing with video sequences, namely the temporal coherency. Indeed, as each tone mapper deals with each frame separately, no temporal coherency is taken into account and hence the results can be quite disturbing for high varying dynamics in a video. We have proposed a temporal coherency algorithm that is designed to analyze a video as a whole, and from its characteristics adapts each tone mapped frame of a sequence in order to preserve the temporal coherency. This temporal coherency algorithm has been tested on a set of real as well as Computer Graphics Image (CGI) content and put in competition with several algorithms that are designed to be time-dependent. Results have shown that temporal coherency preserves the overall contrast in a sequence of images. Furthermore, this technique is applicable to any TMO as it is a post-processing that only depends on the used TMO. This work gave rise to a paper presented at the EUSIPCO conference [5] and in "Signal Processing: Image Communication, Elsevier journal" [1].

6 Contracts and Grants with Industry

6.1 Contracts

We are part of a project, named Nevex, with industrial partners such as: Technicolor, TF1, and other companies involved in HDR acquisition and display. This last project is concerned with the design of a complete High Dynamic Range chain: acquisition, compression, transmission and display. We have also two CIFRE grants with Technicolor.

7 Other Grants and Activities

7.1 International Collaborations

We have several collaborations with other laboratories and universities: University of Central Florida (UCF) in Orlando (US), University of Utah in Salt Lake City (US), Technical University of Prague (Czech Republic), University of Minho in Braga (Portugal), University of Pretoria (South Africa). Within the framework of these collaborations, several PhD students have been or are jointly supervised.

We are also in an european COST action (HDRi IC-1005) whose leader is the university of Warwick, England.

7.2 National Collaborations

8 Dissemination

8.1 Involvement in the Scientific Community

Most of the team members have published in:

- high level journals: IEEE Trans. Visualization and Computer Graphics, IEEE Comp. Graphics and Applications, Computer Graphics Forum, Journal of Graphics Tools, The Visual Computer, etc.

- and high level conferences: EGSR, Pacific Graphics, IEEE Virtual Reality, Computer Graphics International, Web3D, etc.

They have been reviewers for several conferences and symposia such as: Siggraph, Eurographics Pacific graphics, etc., and for journals such as: ACM TOG, IEEE TVCG, The Visual Computer, Computer Graphics Forum, etc.

Kadi Bouatouch is an associate editor of the journal The Visual Computer. He has also been external examiners for several PhD in Europe (England, Germany, Belgium, The Netherlands, Spain, Cyprus) and the United States (UCF, Utah).

Christian Bouville and Kadi Bouatouch have been the co-chairs of the Web3D'2003 conference which has been held in Saint Malo, France. We have also founded the Eurographics Workshop on Rendering in 1990. This workshop became EGSR (EuroGraphics Symposium on Rendering) in 2003.

8.2 Teaching

Kadi Bouatouch is Professor at ISTIC, university of Rennes 1. He gives courses on programming (Mathematica, JAVA) and on computer graphics. He is responsible for Digital Image track of the Master of computer science delivered by the university of Rennes 1.

Remi Cozot is an assistant professor at ESIR, university of Rennes 1. He gives courses on programming (JAVA, C++) and on computer graphics.

9 Bibliography

Major publications by the team in recent years

[1] D. K. B. Duvenhage, K. Bouatouch, "Verification of Bidirectional Reflectance Distribution-Function Implementations for Physical Plausibility", South African Institute for Computer Scientists and Information Technologists, October 2013.

[2] R. Boitard, D. Thoreau, R. Cozot, K. Bouatouch, "Impact of Temporal Coherence-Based Tone Mapping on Video Compression", *in : Eusipco*, Marrakech, Marocco, September 2013.

[3] R. Boitard1, R. Cozot, D. Thoreau, K. Bouatouch, "Zonal Brightness Coherency for Video Tone Mapping", *Signal Processing: Image Communication, Elsevier journal, 29*, 2, February 2014.

[4] C. Collin, K. Chen, A. Hakke-Patil, S. Pattanaik, K. Bouatouch, "Green's function solution to subsurface light transport for BRDF computation", Spring Conference on Computer Graphics, May 2013.

[5] C. Collin, M. Ribardiere, A. Gruson, R. Cozot, K. Bouatouch, S. Pattanaik, "Visibility-driven progressive volume photon tracing", *in : Computer Graphics International*, Hannover, Germany, June.

[6] C. Collin, M. Ribardiere, A. Gruson, R. Cozot, K. Bouatouch, S. Pattanaik, "Visibility-driven progressive volume photon tracing,", *The Visual Computer 29*, 9, 2013.

[7] A. Gruson, M. Ribardiere, R. Cozot, "Eye-Centered Color Adaptation in Global Illumination", *Computer Graphics Forum 32*, 7, 2013, p. 111–120.

[8] R. Marques, C. Bouville, M. Ribardiere, L. Santos, K. Bouatouch, "Spherical Fibonacci Point Sets for Illumination Integrals", *Computer Graphics Forum 32*, 8, 2013.

[9] R. Marques, C. Bouville, M. Ribardiere, L. Santos, K. Bouatouch, "A Spherical Gaussian Framework for Bayesian Monte Carlo Rendering of Glossy Surfaces", *IEEE Transaction On Visualization and Computer Graphics 29*, 10, 2013.

Articles in referred journals and book chapters

[1] R. Boitard1, R. Cozot, D. Thoreau, K. Bouatouch, "Zonal Brightness Coherency for Video Tone Mapping", *Signal Processing: Image Communication, Elsevier journal, 29*, 2, February 2014.

[2] C. Collin, M. Ribardiere, A. Gruson, R. Cozot, K. Bouatouch, S. Pattanaik, "Visibility-driven progressive volume photon tracing,", *The Visual Computer 29*, 9, 2013.

[3] R. Marques, C. Bouville, M. Ribardiere, L. Santos, K. Bouatouch, "Spherical Fibonacci Point Sets for Illumination Integrals", *Computer Graphics Forum 32*, 8, 2013.

[4] R. Marques, C. Bouville, M. Ribardiere, L. Santos, K. Bouatouch, "A Spherical Gaussian Framework for Bayesian Monte Carlo Rendering of Glossy Surfaces", *IEEE Transaction On Visualization and Computer Graphics 29*, 10, 2013.

Publications in Conferences and Workshops

[5] R. Boitard, D. Thoreau, R. Cozot, K. Bouatouch, "Impact of Temporal Coherence-Based Tone Mapping on Video Compression", *in : Eusipco*, Marrakech, Marocco, September 2013.

[6] C. Collin, M. Ribardiere, A. Gruson, R. Cozot, K. Bouatouch, S. Pattanaik, "Visibility-driven progressive volume photon tracing", *in : Computer Graphics International*, Hannover, Germany, June.