INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team Reves

Rendering and virtual environments with sound

Sophia Antipolis - Méditerranée

Theme : Interaction and Visualization

Activity Report

2009

Table of contents

1. Team

Research Scientist

George Drettakis [Team Leader, Research Director (DR) Inria, HdR]

Sylvain Lefebvre [Research Associate (CR) Inria, joint appointment Reves and Alice]

Technical Staff

Manuel Asselot [Specialist Engineer, 2 year contract INRIA-funded, until Apr.27]

David Grelaud [Associated Engineer, 2 year contract INRIA-funded]

PhD Student

Nicolas Bonneel [CROSSMOD project, until Nov.15]

Marcio Cabral [INRIA CORDIS scholarship]

Pierre-Yves Laffont [INRIA CORDIS scholarship, since October]

Cécile Picard [INRIA/EdenGames collaboration, part-time in Grenoble (Evasion)]

Post-Doctoral Fellow

Ares Lagae [FWO Post-Doctoral Fellowship, Belgium, since October]

Peter Vangorp [INRIA-funded, since October]

Visiting Scientist

Karan Singh [Visiting Scientist, Associate Professor, Univ. of Toronto, Canada, June-Sept.]

Michiel van de Panne [Visiting Scientist, Associate Professor Univ. of British Columbia, Canada, until July 31]

Administrative Assistant

Sophie Honnorat [Administrative Assistant (AI) Inria, part-time]

Other

Matthäus G. Chajdas [Intern, Univ. of Erlangen, Germany, Apr-Oct.]

Emmanuelle Chapoulie [Intern, EPU Nice, June-Sept.]

Gaurav Chaurasia [Intern, IIT Delhi, India, May-Sept]

Georgios Koulieris [Intern, Univ. of Athens, Greece, until March 31]

2. Overall Objectives

2.1. General Presentation

Images, often accompanied by sound effects, have become increasingly present in our everyday lives; this has resulted in greater needs for content creation. Despite the fact that many traditional means exist, such as photography, artistic graphic design, audio mixing, they typically still remain the reserve of the expert, and require significant investment in time and expertise. Our main interest is computer image and sound synthesis, with an emphasis on automated methods. Our main goals include the simplification of the tasks required for the production of sound and images, as well as the development of new techniques for their generation.

The application domain is vast. It ranges from audiovisual production, which typically requires long, offline computation to obtain high quality results, all the way to real-time applications such as computer games or virtual reality, for which the main consideration is to guarantee 60 frames per second frame rates, or, in general the reduction of latency to user reaction. The process of generation of images and sound, generally called *rendering* is our primary interest; our second main interest are virtual environments (VE's) as well as augmented (AE's) or mixed environments (ME's), that is scenes containing both real objects (often digitized) as well as purely synthetic objects. We are interested in both the generation and the interaction with these environments. We use the term virtual environments for scenes with a certain degree of interactivity, potentially in a semi-immersive (stereo and tracking, workbench) or immersive (CAVE, RealityCenter) context.

2.2. Highlights

Nicolas Bonneel defended his Ph.D. thesis in September 2009 [14].

Cécile Picard defended her Ph.D. thesis in December 2009 [15].

3. Scientific Foundations

3.1. Rendering

We consider plausible rendering to be a first promising research direction, both for images and for sound. Recent developments, such as point rendering, image-based modeling and rendering, and work on the simulation of aging indicate high potential for the development of techniques which render *plausible* rather than extremely accurate images. In particular, such approaches can result in more efficient renderings of very complex scenes (such as outdoors environments). This is true both for visual (image) and sound rendering. In the case of images, such techniques are naturally related to image- or point-based methods. It is important to note that these models are becoming more and more important in the context of network or heterogeneous rendering, where the traditional polygon-based approach is rapidly reaching its limits. Another research direction of interest is realistic rendering using simulation methods, both for images and sound. In some cases, research in these domains has reached a certain level of maturity, for example in the case of lighting and global illumination. For some of these domains, we investigate the possibility of technology transfer with appropriate partners. Nonetheless, certain aspects of these research domains, such as visibility or high-quality sound still have numerous and interesting remaining research challenges.

3.1.1. Plausible Rendering

3.1.1.1. Alternative representations for complex geometry

The key elements required to obtain visually rich simulations, are sufficient geometric detail, textures and lighting effects. A variety of algorithms exist to achieve these goals, for example displacement mapping, that is the displacement of a surface by a function or a series of functions, which are often generated stochastically. With such methods, it is possible to generate convincing representations of terrains or mountains, or of non-smooth objects such as rocks. Traditional approaches used to represent such objects require a very large number of polygons, resulting in slow rendering rates. Much more efficient rendering can be achieved by using point or image based rendering, where the number of elements used for display is view- or image resolution-dependent, resulting in a significant decrease in geometric complexity. Such approaches have very high potential. For example, if all object can be rendered by points, it could be possible to achieve much higher quality local illumination or shading, using more sophisticated and expensive algorithms, since geometric complexity will be reduced. Such novel techniques could lead to a complete replacement of polygon-based rendering for complex scenes. A number of significant technical challenges remain to achieve such a goal, including sampling techniques which adapt well to shading and shadowing algorithms, the development of algorithms and data structures which are both fast and compact, and which can allow interactive or real-time rendering. The type of rendering platforms used, varying from the high-performance graphics workstation all the way to the PDA or mobile phone, is an additional consideration in the development of these structures and algorithms. Such approaches are clearly a suitable choice for network rendering, for games or the modelling of certain natural object or phenomena (such as vegetation, e.g. Figure 1, or clouds). Other representations merit further research, such as image or video based rendering algorithms, or structures/algorithms such as the "render cache" [40], which we have developed in the past, or even volumetric methods. We will take into account considerations related to heterogeneous rendering platforms, network rendering, and the appropriate choices depending on bandwith or application. Point- or image-based representations can also lead to novel solutions for capturing and representing real objects. By combining real images, sampling techniques and borrowing techniques from other domains (e.g., computer vision, volumetric imaging, tomography etc.) we hope to develop representations of complex natural objects which will allow rapid rendering. Such approaches are closely related to texture synthesis and image-based modeling. We believe that such methods will not replace 3D (laser or range-finger) scans, but could be complementary, and represent a simpler and lower cost alternative for certain applications (architecture, archeology etc.). We are also investigating methods for adding "natural appearance" to synthetic objects. Such approaches include *weathering* or *aging* techniques, based on physical simulations [27], but also simpler methods such as accessibility maps [36]. The approaches we intend to investigate will attempt to both combine and simplify existing techniques, or develop novel approaches founded on generative models based on observation of the real world.

3.1.1.2. Plausible audio rendering

Similar to image rendering, plausible approaches can be designed for audio rendering. For instance, the complexity of rendering high order reflections of sound waves makes current geometrical approaches inappropriate. However, such high order reflections drive our auditory perception of "reverberation" in a virtual environment and are thus a key aspect of a plausible audio rendering approach. In complex environments, such as cities, with a high geometrical complexity, hundreds or thousands of pedestrians and vehicles, the acoustic field is extremely rich. Here again, current geometrical approaches cannot be used due to the overwhelming number of sound sources to process. We study approaches for statistical modeling of sound scenes to efficiently deal with such complex environments. We also study perceptual approaches to audio rendering which can result in high efficiency rendering algorithms while preserving visual-auditory consistency if required.

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [26], representing a scene with trees, grass and flowers. We can achieve 7-8 frames per second compared to tens of seconds per image using standard polygonal rendering.

3.1.2. High Quality Rendering Using Simulation

3.1.2.1. Non-diffuse lighting

A large body of global illumination research has concentrated on finite element methods for the simulation of the diffuse component and stochastic methods for the non-diffuse component. Mesh-based finite element approaches have a number of limitations, in terms of finding appropriate meshing strategies and form-factor calculations. Error analysis methodologies for finite element and stochastic methods have been very different in the past, and a unified approach would clearly be interesting. Efficient rendering, which is a major advantage of finite element approaches, remains an overall goal for all general global illumination research. For certain cases, stochastic methods can be efficient for all types of light transfers, in particular if we require a view-dependent solution. We are also interested both in *pure* stochastic methods, which do not use finite element techniques. Interesting future directions include filtering for improvement of final image quality as well as beam tracing type approaches [38] which have been recently developed for sound research.

3.1.2.2. Visibility and Shadows

Visibility calculations are central to all global illumination simulations, as well as for all rendering algorithms of images and sound. We have investigated various global visibility structures, and developed robust solutions for scenes typically used in computer graphics. Such analytical data structures [31], [30], [29] typically have robustness or memory consumption problems which make them difficult to apply to scenes of realistic size. Our solutions to date are based on general and flexible formalisms which describe all visibility event in terms of generators (vertices and edges); this approach has been published in the past [28]. Lazy evaluation, as well as hierarchical solutions, are clearly interesting avenues of research, although are probably quite application dependent.

3.1.2.3. Radiosity

For purely diffuse scenes, the radiosity algorithm remains one of the most well-adapted solutions. This area has reached a certain level of maturity, and many of the remaining problems are more technology-transfer oriented We are interested in interactive or real-time renderings of global illumination simulations for very complex scenes, the "cleanup" of input data, the use of application-dependent semantic information and mixed representations and their management. Hierarchical radiosity can also be applied to sound, and the ideas used in clustering methods for lighting can be applied to sound.

3.1.2.4. High-quality audio rendering

Our research on high quality audio rendering is focused on developing efficient algorithms for simulations of geometrical acoustics. It is necessary to develop techniques that can deal with complex scenes, introducing efficient algorithms and data structures (for instance, beam-trees [32] [38]), especially to model early reflections or diffractions from the objects in the environment. Validation of the algorithms is also a key aspect that is necessary in order to determine important acoustical phenomena, mandatory in order to obtain a high-quality result. Recent work by Nicolas Tsingos at Bell Labs [33] has shown that geometrical approaches can lead to high quality modeling of sound reflection and diffraction in a virtual environment (Figure 2). We will pursue this research further, for instance by dealing with more complex geometry (e.g., concert hall, entire building floors).

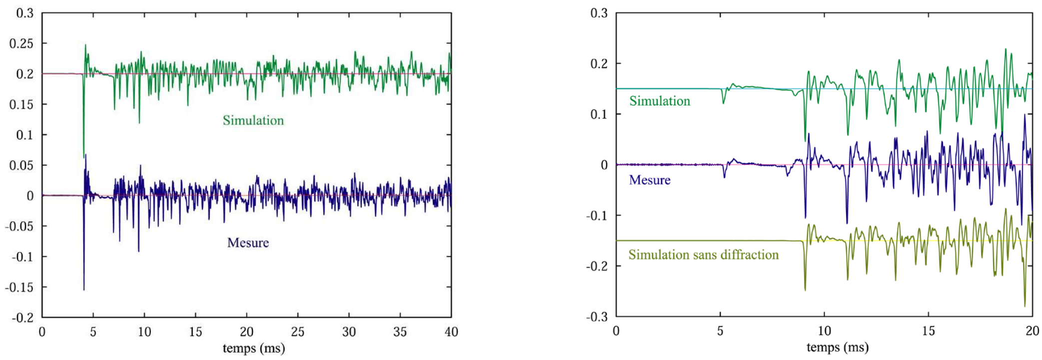

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and diffraction off a panel introduced in the room.

Finally, several signal processing issues remain in order to properly and efficiently restitute a 3D soundfield to the ears of the listener over a variety of systems (headphones, speakers). We would like to develop an open and general-purpose API for audio rendering applications. We already completed a preliminary version of a software library: AURELI [39].

3.2. Virtual and Augmented Environments with Sound

The second major research direction of our group is on virtual, augmented or mixed environments, which include both visual and sound representations. We are mainly interested in interactive environments, permitting the user to create and manipulate scenes consisting of both real and synthetic objects. As a first step, we consider *real* objects to be digitised representations of reality, rather than the real world. Our first goal is to apply and adapt our rendering expertise, presented in the previous paragraphs to virtual and augmented reality. There are three areas in which we concentrate our efforts: consistent lighting between real and synthetic illumination, for shadows and reflections, enriching virtual and augmented environments with sound, in a consistent manner and finally appropriate interaction and visual paradigms for virtual and augmented environments.

3.2.1. Efficient and Simple Relighting

We wish to develop relighting and consistent real/virtual lighting methods which have simple input requirements: i.e., a small number of input images, and the smallest number of restrictions on the lighting conditions. The goal is to get high quality results for both interior and outdoors environments. To achieve these goals, we investigate ways to extract approximate reflectances in real scenes, potentially using scene or image statistics, and by including some level of user interaction in the process. For efficient display, texture capacities of modern graphics hardware will definitely be advantageous.

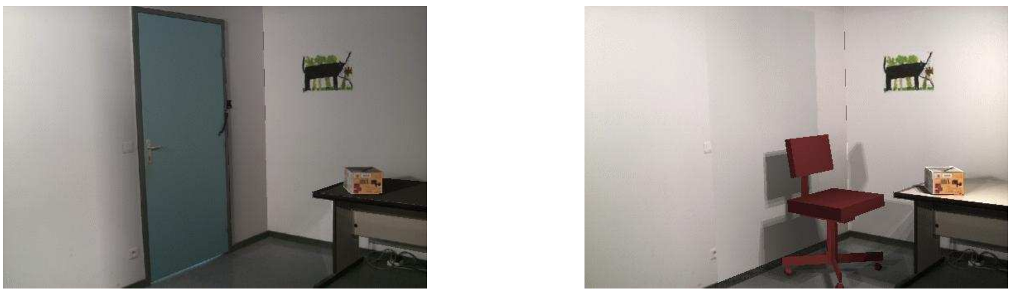

Figure 3. (a) Original conditions (b) The door has been removed virtually, and a virtual object and light have been added (method of [35])

Our previous work on interior relighting (Figure 3) has given satisfactory solutions, allowing us to add virtual object with consistent lighting, but implies severe restrictions on the lighting conditions of input images [34], [35]. Such approaches are based on the creation of "shadow free" base textures using heuristics, and a relatively precise reconstruction of the geometry. For outdoors scenes, geometric complexity and the fact that lighting conditions cannot be easily manipulated render such approaches less appropriate. However, some of the techniques developed can be applied, and we believe that the key is to combine automated techniques with user interaction at the various stages of the process. The long-term goal is to turn on a video camera in a scene (potentially with partially pre-reconstructed geometry), and be able to add virtual objects or light sources interactively in a consistent manner into the video stream. Relighting could also be achieved in this manner, or using semi-transparent glasses or headsets. Applications of such an approach are numerous, for archeology, architecture and urban planning, special effects, manufacturing, design, training, computer games etc. This long term vision will require a way to smoothly vary from low-quality methods [34], [35] to high quality approaches [41], in a manner which is much less complex in terms of capture, processing for relighting and (re)rendering.

3.2.2. Enriching virtual environments with sound

Consistent rendering of real and synthetic sounds is a key aspect for virtual reality applications. Solving the problem would make it possible to mix natural sounds with synthesized spatial audio for augmented reality applications. This can be used to enrich the natural soundscape with additional auditory information through wearable devices (e.g., virtual museums, etc.). Another application would be to provide auditory feedback to visually-impaired people while preserving their natural auditory perception. Another future direction of research is active control of rooms and listening spaces. Such control can be achieved by coupling microphones and speaker arrays and allow for modifying the natural acoustical properties of the space (e.g., reverberation time) in real-time. Such technologies have already been used to improve acoustics in concert halls that, for a variety of reasons, do not sound as good as designed for. They appear to be promising for VR/AR applications. However, existing techniques yet have to be improved to be applied in this context.

3.2.3. Interaction and Visual Paradigms for Virtual and Augmented Environments

The use of immersive or semi-immersive systems opens a large number of new types of interaction with virtual or augmented environments. There is a vast body of research on interfaces for 3D environments, and in particular for immersive systems. Our focus will be on specific interfaces, interaction or visual paradigm problems which inevitably appear in the course of our research. When necessary, we will work with complementary partners in Computer-Human Interaction to find solutions to these problems. One question we consider important is finding appropriate interface paradigms which replace 2D (menu or button-based) interfaces both in the context of the actual rendering research process and for the applications we investigate. Despite significant previous work in the domain, a standard that can be widely adopted has yet to be defined. It may be that the lack of standard interfaces is part of the reason why immersive systems are not being adopted as widely nor as rapidly as their inventors would have hoped. In terms of visual representation, non-photorealistic (NPR) or expressive, renderings are an interesting avenue of investigation. In particular, NPR can allow abstraction of unimportant details and more efficient communication of certain concepts. Since a number of the algorithms developed are based on inherently 2D drawing, their transposition to immersive, stereo-display environments poses a number of very interesting and challenging questions. There are also some applications domains, for example archeology or architecture, where drawing-style renderings are part of the current workflow, and which will naturally fit into a EVs adapted to these domains. Virtual storytelling is another domain in which NPR has a natural application. Immersive, stereo-based systems seem a well-adapted platform for more intuitive interactive modelling in 3D. The development of efficient and flexible structures such as procedural point-based representations, or rapid aging techniques in a true 3D context could result in systems which are much more efficient that 2D displays, in which the sensation of 3D depth and immersion is missing. Finally, the inclusion of spatialised sound for 3D interfaces is clearly a promising research direction. The benefit of consistent 3D sound is evident, since it results in better spatial perception for the user, can help for example in determining spatial or visibility relationships, resulting in improved usability. The actual inclusion of sound effects or sound metaphors in interface design is clearly an interesting challenge.

4. Software

4.1. APF: state-of-the-art 3D audio library

Participants: David Grelaud, George Drettakis.

This work was performed in collaboration with Jean-Christophe Lombardo of the DREAM research engineer service at INRIA Sophia-Antipolis Méditerranée. REVES has several audio research publications: the first component is the so called masking or culling algorithm, which aims at removing all the inaudible audio sources from a virtual scene based on perceptual metrics. The second component, called clustering, aims at grouping audio sources that are spatially close to each other and premix them to a representative cluster source, so that all spatialization related processing can be applied only on the representative premixed source. Other audio topics were also considered and developed, like progressive and scalable frequency domain mixing, sound propagation, scalable reverberation, modal sound synthesis and contact sounds generation.

In order to maintain all the knowledge in the group and re-use these technologies in the Immersive Space, David Grelaud has written a fully documented audio library (APF) which gathers about 10 audio publications and 1 US patent. APF is a cross-platform, object oriented C++ API available on GForge. All the code has been re-implemented and a completely new software architecture resulted in a twofold increase in the speed of our algorithms. APF runs in the Immersive Space and uses the tracking system to spatialize virtual audio sources around the listener. It can also exploit personal Head Related Transfer Functions (HRTF). We are currently implementing a network communications layer to create an audio rendering server on a separate machine.

This audio library has already been used to establish a partnership with the University of Aachen in Germany. They will share their audio technologies which allows people to hear near-to-head sources realistically. This will drastically enhance the sense of immersion in the virtual reality platform and it is crucial for the success of the immersive space. APF has also been critical in establishing collaborations in the context of various grant proposals.

4.2. Sound Grains Software

Participants: David Grelaud, Cécile Picard, George Drettakis.

In the context of the Ph.D. thesis of C. Picard, we developed a novel solution for sound synthesis based on grains (see section below on new results). The prototype software developed in the context of the thesis has been rewritten cleanly by D. Grelaud as will be an independent and documented library in C++. This library will be transferred to EdenGames as part of the overall agreement governing the Ph.D. thesis of C. Picard.

4.3. OgreVR - A Virtual Reality Framework

Participant: David Grelaud.

We have upgraded OgreVR which is a framework to use Ogre3D, an open-source graphics rendering engine, on virtual reality platform. OgreVR is now compatible with the I-Space and the CADWall. We have also implemented a special version of LibSL, a C++ graphics-programming toolbox originally developed by Sylvain Lefebvre, which allows us to deploy rapidly a 3d application in the Immersive Space. These two libraries are available on GForge.

4.4. Example code for texture synthesis by example

Participant: Sylvain Lefebvre.

Sylvain Lefebvre created a simple implementation of his prior work on texture synthesis for research and educational purposes. This program is available online at http://www-sop.inria.fr/members/Sylvain.Lefebvre/wiki/Main/TSynEx It has been downloaded 975 times at this date (27/11/2009).

4.5. LibSL - Simple library for graphics

Participant: Sylvain Lefebvre.

Sylvain Lefebvre continued development of the LibSL graphics library (under CeCill-C licence, filed at the APP). The library is actively used in both the REVES / INRIA Sophia-Antipolis and the Alice / INRIA Nancy Grand-Est teams. It has been used in several publications including two this year [18], [17].

4.6. GaborNoise Software

Participants: Ares Lagae, Sylvain Lefebvre, George Drettakis.

We have recently proposed a new procedural noise function, Gabor noise (see new results section below). In the context of this project, we have developed an extensive software package, which includes a CPU reference implementation of the 2D noise, and a complete GPU implementation of the 2D noise, surface noise, and 3D noise. This software package is in the process of being filed for APP protection.

5. New Results

5.1. Plausible Image Rendering

5.1.1. Procedural Noise using Sparse Gabor Convolution

Participants: Ares Lagae, Sylvain Lefebvre, George Drettakis.

Faithfully modeling, rendering and animating the world we live is the driving problem in computer graphics. Our world is very complex, and it was recognized early on that randomness is a very useful tool for modeling this complexity. This randomness is typically expressed using noise functions, which generate random yet controllable noise, that is used to model the appearance of objects. Several desirable properties of noise were identified in previous work. Over the last decades, many different noise functions have been proposed, each with a different subset of these properties. However, none of the existing noise functions has all of these properties.

In collaboration with K.U.Leuven university, we have developed a new noise function based on sparse convolution and the Gabor kernel. Our new noise function offers accurate spectral control with intuitive parameters such as orientation, principal frequency and bandwidth. Our noise supports two-dimensional and solid noise, but we also introduce setup-free surface noise. This is a method for mapping noise onto a surface, complementary to solid noise, that maintains the appearance of the noise pattern along the object and does not require a texture parametrization. Our approach requires only a few bytes of storage, does not use discretely sampled data, and is nonperiodic. It supports anisotropy and anisotropic filtering.

Our new noise function combines all of the desirable properties of noise functions. This allows users of noise functions to benefit from the advantages implied by all properties, rather than having to chose a single noise function which only has a subset of the properties. We have demonstrated our noise using an interactive tool for noise design, which is illustrated in figure 4.

Figure 4. (Left) We have developed a new procedural noise function that offers accurate spectral control. The user can interactively manipulate the power spectrum. (Middle) We apply the noise to a surface without the need for texture coordinates, and provide high-quality anisotropic filtering. Thanks to increased expressiveness of the noise, a simple colormap is enough to produce visually interesting textures. (Right) Since our surface noise does not require a texture parametrization, the surface can evolve dynamically and even sustain topology changes.

This work is a collaboration with Philippe Dutré (Computer Graphic Group, KU Leuven), and has been published in ACM Transactions on Graphics, presented also at ACM SIGGRAPH 2009 [18].

5.1.2. Improving Gabor Noise

Participants: Ares Lagae, Sylvain Lefebvre.

We have recently proposed a new procedural noise function, Gabor noise (see above). In a follow up project, we present three significant improvements to Gabor noise: (1) an isotropic kernel for Gabor noise, which speeds up isotropic Gabor noise with a factor of roughly two, (2) an error analysis of Gabor noise, which relates the kernel truncation radius to the relative error of the noise, and (3) spatially varying Gabor noise, which enables spatial variation of all noise parameters. These improvements make Gabor noise an even more attractive alternative for existing noise functions.

This work is a collaboration with Philippe Dutré (Computer Graphic Group, KU Leuven) and is currently submitted for publication.

5.1.3. Gabor Noise by Example

Participants: Ares Lagae, Peter Vangorp, George Drettakis, Sylvain Lefebvre.

In another follow up project to the Gabor Noise work presented above, we are trying to derive the parameters for a procedural noise function from exemplars. This has important applications in, for example, texture synthesis.

We are currently investigating the extension of Gabor noise and other stochastic texture synthesis methods to include certain aspects of structure, such as regularity, edges, lines, corners, and gradients. Research in human visual perception [37] has shown that a texture's first and second order statistics, used in most synthesis-by-analysis methods, are insufficient to reproduce structure beyond simple anisotropy, and that subband correlations and phase statistics do capture certain aspects of structure.

5.1.4. Single Photo Estimation of Hair Appearance

Participants: Nicolas Bonneel, Michiel van de Panne, George Drettakis.

Significant progress has been made in high-quality hair rendering, but it remains difficult to choose parameter values that reproduce a given real hair appearance. In particular, for applications such as games where naive users want to create their own avatars, tuning complex parameters is not practical. We developed an approach using a single flash photograph which estimates model parameters that reproduce the visual likeness of the observed hair. We developed an example usage scenario where a photograph is used to infer the hair appearance to an avatar (Fig.5).

We estimate the color absorption coefficients of the hair fibers, three reflectance lobe parameters of a multiple-scattering rendering model, and a geometric noise parameter that we introduce.

We use a novel melanin-based model to capture the natural subspace of hair absorption parameters. At its core, the method assumes that images of hair with similar color distributions are also similar in appearance. This allows us to recast the issue as an image retrieval problem where the photo is matched with a dataset of rendered images; we thus also match the model parameters used to generate these images (Fig.6, far left). An Earth-mover's distance is used between luminance-weighted color distributions to gauge similarity. We conducted a perceptual experiment to evaluate this metric in the context of hair appearance and demonstrate the method on 64 photographs, showing that it can achieve a visual likeness for a large variety of input photos. The estimated parameters can further be used in new natural lighting conditions (Fig.6).

This work is a collaboration with Sylvain Paris (Adobe) and Fredo Durand (MIT), and has been published in the special issue of the journal Computer Graphics Forum, and presented at the Eurographics Symposium on Rendering 2009 [16].

5.1.5. Gigavoxels

Participant: Sylvain Lefebvre.

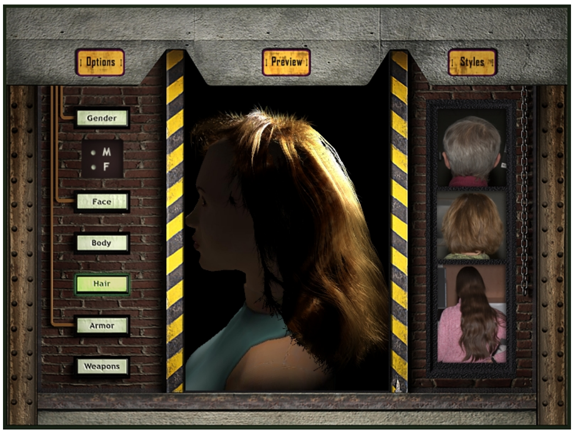

Figure 5. A typical usage of the proposed method: a user selects a photographs to infer the hair color to the avatar.

Figure 6. Far left: The histogram of the photograph is matched with the closest histogram of a database of pre-rendered hair images. A sample result. From left to right: the original photograph; the nearest match in the rendering databased providing the hair parameters; a photograph of the same person in a new lighting condition; a rendering of the CG hair using a similar lighting environment and the estimated hair parameters.

The goal of the GigaVoxel project [20] is to render large environments made of voxel volume data - a massive grid of opacity values - in real time, using the GPU. Our approach performs both rendering and asynchronous streaming of voxel data at interactive rates. This lets the user explore the data freely, at very different scales. Applications range from medical data display to large environments for computer graphics and entertainment.

This is a collaboration between Cyril Crassin (AERIS / INRIA Grenoble), Fabrice Neyret (AERIS / INRIA Grenoble), Elmar Eisemann (MPI) and Sylvain Lefebvre. The work has been published at the I3D 2009 conference. It was also presented, together with improvements of the algorithm, during a technical talk at SIGGRAPH 2009. [20]

5.1.6. EG STAR report on example-based texture synthesis

Participant: Sylvain Lefebvre.

Texture synthesis from example has been a very active area of computer graphics, with several major advances over the past 10 years. The domain has now reach maturity and the community was lacking a survey discussing the various approaches. We thus proposed a state of the art report on example-based texture synthesis methods.

This work is a collaboration with Li-Yi Wei (Microsoft Research Asia), Vivek Kwatra (Google), and Greg Turk (Georgia Tech), and a report has been published and presented at the Eurographics 2009 conference. [24]

5.1.7. Assisted Texture Assignement

Participants: Matthäus G. Chajdas, Sylvain Lefebvre.

Matthäus G. Chajdas was an intern within the REVES team from April to September 2009. He was supervised by Sylvain Lefebvre and worked on an algorithm to help modelers texture large virtual environments. Modelers typically manually select a texture from a database of materials for each and every surface of a virtual environment. Our algorithm automatically propagates user input throughout the entire environment as the user is applying textures to it. After choosing textures for only a small subset if the surface, the entire scene is textured. This work has been submitted for publication.

5.1.8. Relighting Photographs of Tree Canopies

Participants: Marcio Cabral, Nicolas Bonneel, Sylvain Lefebvre, George Drettakis.

In this ongoing work, we present an image-based approach to relighting photographs of tree canopies. Given a set of input photographs taken at a single time of day, we are able to relight the tree canopy at any other time of day. We first present a careful analysis of lighting in a tree canopy. We propose several simplifying assumptions, including diffuse reflectance for leaves, and a single-scattering volumetric lighting model for the tree canopy. We use globally illuminated synthetic renderings to verify our assumptions. Based on this analysis, we use a volumetric representation of the tree, created from 10-12 images taken at a single time of day and represent lighting using a single-scattering participating media lighting model. We relight the original input image by applying a ratio of our lighting representation at the target and input lighting times. We also present an efficient way to compute this representation by simultaneously coding transmittance from the sky and to the eye in spherical harmonics. We validate our method by relighting images of synthetic trees and their corresponding ground truth. We also present results for real tree photographs where sparse time-lapse sequences are used for validation.

5.1.9. A Texture-Synthesis Approach for Casual Modeling

Participants: Nicolas Bonneel, Sylvain Lefebvre, Michiel Van de Panne, George Drettakis.

Developing visually-rich depictions of natural scenes is a challenging problem. We developed a method which uses a simple 3D proxy geometry and a source photo to create such images for casual modeling applications. The method is inspired by guidance-based texture synthesis techniques.

A simple proxy model is first rapidly created by a non expert user, then rendered to obtain an image-space guidance map of the required texture regions. To obtain silhouettes and borders with the same rich detail as the source photo, we use a first texture synthesis pass to produce an enriched guidance map. This is then used in a subsequent guided synthesis step to produce the final image. Our use of Chamfer distance metric provides a way for dealing with discrete texture labels in a principled fashion and provides high quality results. Because texture synthesis is invoked as a shading operation, including for consistent depth, our method allows for the consistent integration of standard CG elements and texture-synthesized elements in a scene. This work is currently in progress.

5.2. Plausible Audio Rendering

5.2.1. Retargetting Example Sounds to Interactive Physics-Driven Animations

Participants: Cécile Picard, Nicolas Tsingos.

We propose a new method to generate audio in the context of interactive animations driven by a physics engine. Our approach aims at bridging the gap between direct playback of audio recordings and physically-based synthesis by retargetting audio grains extracted from the recordings according to the output of a physics engine. It combines an off-line analysis process with an interactive on-line resynthesis, as illustrated in the Figure 7.

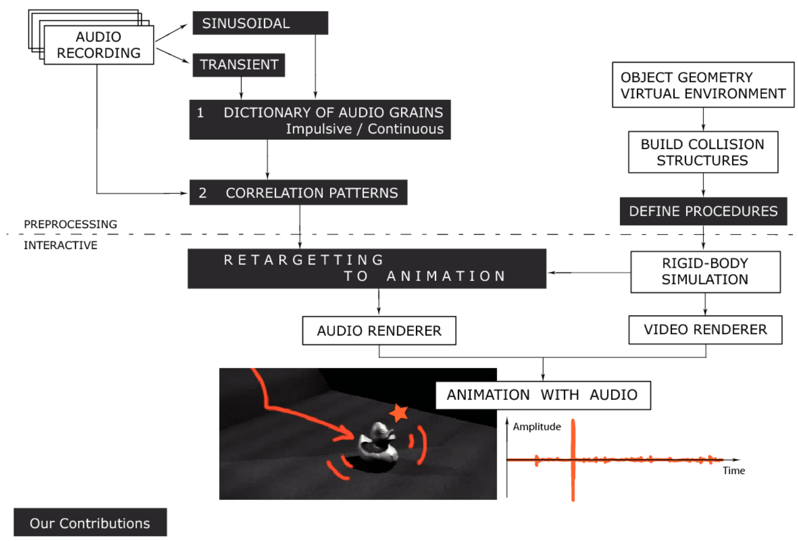

Figure 7. Retargetting Example Sounds to Interactive Physics-Driven Animations; overview of the method that includes an off-line analysis of recodings and an interactive on-line audio synthesis.

In an preprocessing analysis step, we automatically segment audio recordings into atomic grains. The segmentation depends on the type of contact event and we distinguished between impulsive events, e.g. impacts or breaking sounds, and continuous events, e.g. rolling or sliding sounds. We segment recordings of continuous events into sinusoidal and transient components, which we encode separately. In the second step of the analysis process, original audio recordings are decomposed onto the basis formed by the previously extracted audio grains to obtain correlation patterns. The correlation patterns will be used to encode each original recording as a series of audio grains.

A direct benefit of our analysis step is a compact representation of the original sound database since the dictionary of grains and the correlation patterns typically have a smaller memory footprint than the original assets. During interactive animations, the correlation patterns obtained in the off-line analysis process can be directly used to reconstruct the original recordings. Our approach also supports additional, user-defined, resynthesis schemes. For instance, time-scaling of the original audio recordings can be easily implemented using the previously calculated correlation patterns. The method for concatenating grains is flexible and allows the synthesis of an infinite number of similar audio events can be matched to the characteristic rhythmic pattern of a source recording using its correlation pattern. Compelling re-synthesis for physics-driven animations implies adequate matching between audio content and contact parameters to parameters reported from the physics engine and/or user-defined procedures.

Retargetting of extracted audio grains was used to generate interactive audio in the context of simulations driven by the *Sofa*[1] physics engine, see Figure 8. We used the precomputed correlation patterns generated in the analysis step. The penetration force and the relative velocity of the objects in interaction were studied across time in order to detect impacts and continuous contacts. The grains were chosen from an appropriate correlation pattern and triggered, in a random manner or in series.

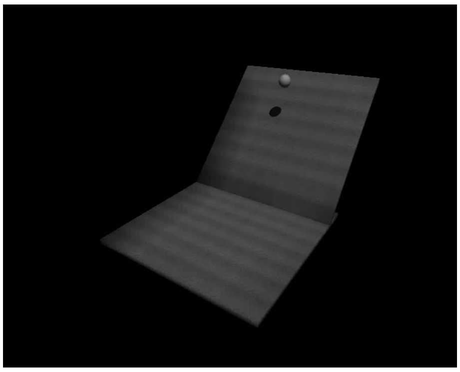

Figure 8. Generating interactive audio in the context of simulations driven by the Sofa physics engine. Audio grains were chosen from an appropriate correlation pattern and triggered, in a random manner or in series, according to the penetration force and the relative velocity reported from the physics engine.

The proposed framework allows the generation of compelling and controllable soundtracks in the context of physics-driven animations. Key aspects of our approach are its low memory usage and the different levels of authoring available to both audio programmers and designers.

This work is a collaboration with François Faure from the Evasion INRIA project-team in Grenoble, and was presented at the *35th AES International Conference - Audio for Games* London, UK, Feb. 2009 [23] also referred in *Sound and Music for Games*, JAES 57(6), June 2009.

5.2.2. A Robust And Multi-Scale Modal Analysis For Sound Synthesis

Participants: Cécile Picard, George Drettakis.

We present a new approach to modal synthesis for rendering sounds of virtual objects. We propose a generic method to efficiently extract modal parameters for any given geometry. Our method preserves sound variety across the surface of an object, at different scales of resolution and for a variety of complex geometries.

[1] Simulation Open Framework Architecture: http://www.sofa-framework.org/

For performing modal decomposition, we must select a deformable modeling method to generate the stiffness and the mass matrices of the mechanical system. Because the object being analyzed can be of arbitrary shape, the Finite Element Method (FEM) and, in particular, the tetrahedral finite element method, is commonly used. However, tetrahedral meshes are computationally expensive for complex geometries, and can be difficult to tune. Our technique performs automatic voxelization of a surface model and automatic tuning of the parameters of hexahedral finite elements, based on the distribution of material in each cell. Thus, we can easily deal with non-manifold geometry which includes both volumetric and surface parts, see Figure 9 (left).

The voxelization is performed using a sparse regular grid embedding of the object, which easily permits the construction of plausible lower resolution approximations of the modal model, see Figure 9 (right).

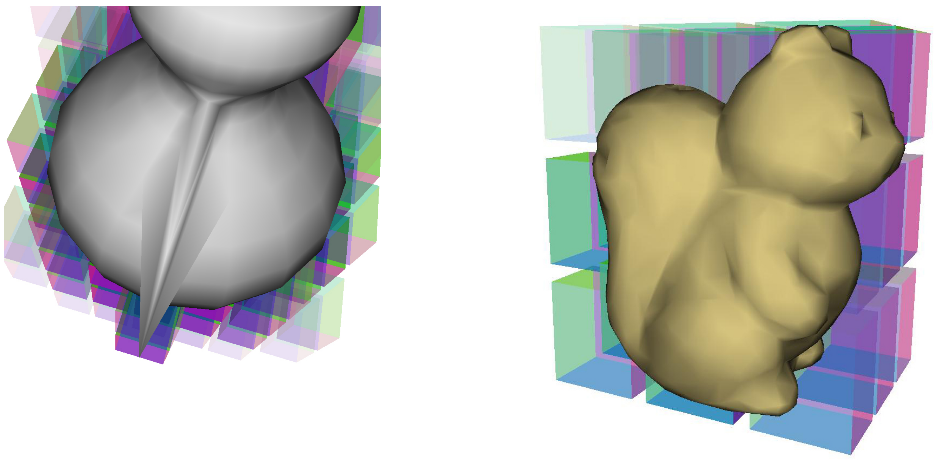

Figure 9. Left: Example of problematic geometry when using tetrahedralization for modal analysis. In contrast, our method can easily deal with non-manifold geometry which includes both volumetric and surface parts. Right: To study the influence of the number of hexahedral finite elements on the sound rendering, we model a sounding object with different resolutions of hexahedral finite elements. We have created a squirrel model with 999 vertices which we use as our test sounding object. Its material is pine wood.

With our approach, we can compute the audible impulse response of a variety of objects, such as those used in games, training simulations, and other interactive virtual environment. The interface between physics engines and audio has often been one of the obstacles for the adoption of physically based sound synthesis. We again use the Sofa Framework. Thus, problems of coherence between physics simulation and sound synthesis are avoided by using exactly the same model for simulation and sound modeling.

This work is a collaboration with François Faure and Paul G. Kry from the Evasion INRIA project-team in Grenoble, and was presented at the *12th International Conference on Digital Audio Effects DAFx-09* Como, Italy in September 2009 [22].

5.3. Virtual Environments with Sound

5.3.1. Structure-Preserving Reshape for Textured Architectural Scenes

Participants: Marcio Cabral, Sylvain Lefebvre, Carsten Dachsbacher, George Drettakis.

Modeling large architectural environments is a difficult task due to the intricate nature of these models and the complex dependencies between the structures represented. Moreover, textures are an essential part of architectural models. While the number of geometric primitives is usually relatively low (i.e., many walls are flat surfaces), textures actually contain many detailed architectural elements. We present an approach for modeling architectural scenes by reshaping and combining existing textured models, where the manipulation of the geometry and texture are tightly coupled. For geometry, preserving angles such as floor orientation or vertical walls is of key importance. We thus allow the user to interactively modify lengths of edges, while constraining angles. Our texture reshaping solution introduces a measure of directional autosimilarity, to focus stretching in areas of stochastic content and to preserve details in such areas. We show results on several challenging models, and show two applications: Building complex road structures from simple initial pieces and creating complex game-levels from an existing game based on pre-existing model pieces.

As a first offline step, our approach extracts features that characterize the model as an architectural piece, i.e., angles and edge lengths. These features are then expressed as a set of linear equations. At runtime, the user can apply modifications to the original model by moving one or more vertices. The systems of equations are then solved in a least square sense to comply with user's motions while preserving its original characteristics. Since we pre-factor the linear system's matrices, user editing is interactive and allows fast manipulation of the geometry giving the user instant visual feedback. Textures are also handled by means of solving a linear system that preserves their characteristics when stretching is performed: in a off-line step, we analyze each texture and extract from them an autosimilarity map which indicates which areas are easier to resynthesize. A grid is then overlaid on top of this map to produce the systems of equations that will allow more or less deformation on the underlying texture depending on the values in the autosimilarity map. We later reintroduce texture details in regions that receive more deformation by identifying tiles of stochastic texture in the autosimilarity map. Finally, we assume that each individual model has a "portal" associated at each extremity, and that portals are compatible, i.e., have the same number of vertices. In a editing session, the user deforms and connects pieces to achieve the desired result. When modifications are made on a given piece, deformations and reshape are appropriately propagated along the chain of connected pieces, by solving the systems to their new portal positions. Results can be seen in Figure 10.

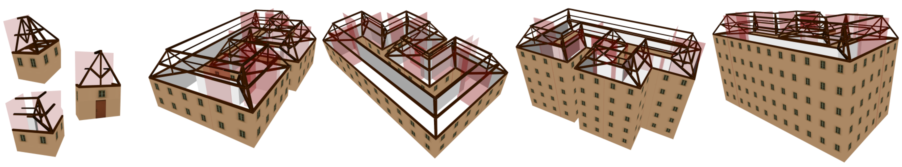

Figure 10. An example of houses with a complex roof frame using three pieces shown left. We show two views of two different variants of the resulting house.

This work is a collaboration with Carsten Dachsbacher (Univ. of Stuttgart) and was presented at the Eurographics 2009 international conference [17], and is published in a special issue of the journal Computer Graphics Forum.

5.3.2. Bimodal perception of audio-visual material properties for virtual environments

Participants: Nicolas Bonneel, George Drettakis.

In this project we study the audio-visual interaction for the estimation of material quality when varying audio LOD (number of modes for contact sounds) and visual LOD (number of spherical harmonic coefficients for BRDF rendering). We performed an experiment with two models (a bunny and a dragon) falling on a table, made of plastic or gold, and showing such a crossmodal interaction, and showing that for all audio qualities, 7 or 9 spherical harmonic bands was not perceived differently as the high quality approximation for material quality estimation.

This work is a collaboration with Clara Suied and Isabelle Viaud-Delmon from CNRS/IRCAM, and has been accepted for publication in 2010 for ACM Transactions on Applied Perception [25].

5.3.3. *Efficient and Practical Audio-Visual Rendering for Games using Crossmodal Perception*

Participants: David Grelaud, Nicolas Bonneel, Manuel Asselot, George Drettakis.

Interactive applications such as computer games, are inherently audio visual, requiring high-quality rendering of complex 3D audio soundscapes and graphics environments. A frequent source of audio events is impact sounds, typically generated with physics engines. We first developed an optimization allowing efficient usage of impact sounds in a unified audio rendering pipeline, also including prerecorded sounds. This optimization allows impact sound per-frame energy computation more than 200 times faster than previous methods (Figure 11), by considering an exponentially decaying energy function. It also allows integrating good quality attacks in the previous pipeline, thus handling recorded sounds with high quality impact sounds.

We also exploited the abovementioned result on audio-visual crossmodal perception ([25]) to introduce a new level-of-detail selection algorithm, which jointly chooses the quality level of audio and graphics rendering. This level-of-detail (LOD) selection is based on our previous crossmodal perceptual study and determines the visual LOD (spherical harmonics of the BRDF) and audio LOD (number of modes) which give best material quality for a given cost.

We have integrated these two techniques as a comprehensive crossmodal audio-visual rendering pipeline in an in-house game engine "PentaG", thus demonstrating the potential utility of our approach.

Figure 11. Using our approach, we can render sounds for larger number of impact sounds in real time.

This work was presented at the I3D 2009 international conference [21] and is a collaboration with Michael Wimmer from the Vienna Institute of Technology.

5.3.4. *Integration of auditory and visual information in the recognition of realistic objects*

Participant: Nicolas Bonneel.

This work presents an experiment aimed a determining crossmodal interactions in reaction times, in a recognition task. The experiment consists in presenting a telephone or a frog during 500ms, with or without sound. The sound can be congruent with the visual or incongruent (a frog sound with a telephone visual), well spatialized or mislocated by 40° to the right of the participant. The goal of the participant is to press a button as fast as possible when a frog is perceived (seen, heard or both). Reaction time is measured. Results show such a crossmodal effect on the semantics, i.e., better reaction times were obtained when the frog visual is displayed with the frog sound. Results also shown that wrong spatialization did not influence reaction time more than the Simon effect (shorter times for objects displayed on the same side as the hand handling the response device).

This work is a collaboration with Clara Suied and Isabelle Viaud-Delmon from CNRS/IRCAM, and has been published in the journal Experimental Brain Research [19].

6. Contracts and Grants with Industry

6.1. ANR Similar Cities

Participant: Sylvain Lefebvre.

The ANR SIMILAR-CITIES started late January 2009. The project advanced well and the first work packages have been completed, including the creation by our partner Allegorithmic of a database of textures. Two meetings have been held, the kickoff meeting in Sophia-Antipolis and a second meeting in Clermont-Ferrand. A server was installed at INRIA Sophia-Antipolis by SEMIR for the project, and the intellectual property agreement has been prepared by the legal department of INRIA. It is currently being signed by all partners.

The project hired Anass Lasram as a PhD student. Anass started in October in the Alice team of INRIA Nancy (since Sylvain Lefebvre joined the Alice team). We are currently working on an algorithm for fast synthesis of structural textures, with a particular focus on urban imagery. The main application is to give a more realistic appearance to the many anonymous buildings of large virtual cities, which are typically only crudely modeled. A key advantage of our approach is the compactness of the results. We are currently working on submitting this work for publication. The project also involves Samuel Hornus (Alice/INRIA Nancy) and Anass Lasram (Alice/INRIA Nancy).

6.2. Alias|Wavefront

We are part of the Alias|Wavefront software donation program. We use Maya extensively to model various objects used in our animations and results, for many of our projects.

6.3. Eden Games

We have an ongoing collaboration with Eden Games (an ATARI game studio in Lyon, France). Eden had previously licensed our audio masking/clustering technology for TestDrive Unlimited and Alone in the Dark Near Death Experience on Xbox360, PCs and PS3. We extended our licensing agreement to all platforms and two new titles. The agreement funds the PhD thesis of C. Picard to co-develop novel real-time synthesis algorithms for interactive audio rendering.

In the context of this agreement, we are in the process of transferring the Sound Grains software module which implements the work presented in [23] on audio grain retargetting.

6.4. Audio Technology Transfer Games

In 2008, Manuel Asselot worked on a project to transfer our audio research of the last few years as an INRIA spin-off. As part of this effort, we contracted a market research study with the company Actemis. The result of the study revealed that while the technology has some promise, for example as a component in a third-party existing audio software library, it is not viable as a spin-off company.

As a consequence, the spin-off project was abandoned in March 2009.

7. Other Grants and Activities

7.1. Regional/Local Projects

7.1.1. Collaboration with CNRS and IRCAM

In the context of the CROSSMOD European project we have developed continuous and active collaboration with Isabelle Viaud-Delmon of CNRS (Laboratoire Vulnérabilité, adaptation et psychopathologie, UMR 7593 and IRCAM) and Olivier Warusfel of IRCAM). This notably resulted in the publication to appear in 2010 in ACM Transactions on Applied Perception [25].

7.1.2. Workbench and Immersive space

Participants: David Grelaud, Nicolas Bonneel, George Drettakis.

New Immersive Space. Under the programme CPER Telius, the INRIA Sophia Antipolis - Mediterranée Research Center built an immersive space (an I-Space and a CADWall); this project was supervised by Jean-Christophe Lombardo of the DREAM research engineer service. David Grelaud was strongly involved in this project from the needs analysis to the preparation of audio-visual demonstrators which will emphasize the teams work. We are currently (end of 2009) in the test phase of all the systems. The equipments are complex and it took a lot of time to set up the first audio-visual demonstrator. David Grelaud participated actively in resolving hardware and software issues on the audio part and the Windows machines. The OgreVR software framework upgrade (see Software section), is part of this effort.

We are also preparing new virtual reality demonstrators. For instance, we have modified completely the Japan Scene application (Figure 12) in order to create a high quality immersive application which shows up our recent scientific results on modal sound synthesis.

Figure 12. Japan Scene demonstrator for the Immersive Space [14]

7.2. Visiting Researchers

Michiel van de Panne, Professor at the University of British Columbia, Canada, visited the Reves research group on a sabbatical for a year (August 2008-July 2009). Karan Singh, Professor at the University of Toronto, Canada, visited the Reves research group on a sabbatical for 4 months (June-Sept.2009)

Ares Lagae and Peter Vangorp, from KU Leuven (FWO Belgium), are visiting the Reves research group for a year since October 2009, as post-doctoral researchers. Matthäus G. Chajdas, from the Univ. of Erlangen, Germany, completed his masters internship (April-October 2009) on Texturing Templates. Gaurav Chaurasia, from IIT Delhi, India, completed his M1 (4th year) internship (May-Sept.2009) on High-Quality View-Dependent Texture Mapping.

We hosted several visiting researchers this year: Ricardo Maroquim (CNR, Pisa, ERCIM fellow), one week in March, Mikio Shinya (University of Tokyo) in March, Pierre Poulin (University of Montreal) in May, Ryan Schmidt (University of Toronto) in July, Fredo Durand (MIT) in September, Olga Sorkine (University of New-York) in November, Alexei Efros (Carnegie Mellon University), Davide Rocchesso (University IUAV Venice), Xavier Serra (Universitat Pompeu Fabra, Barcelona) and Jonathan Ragan-Kelley (MIT), in December.

7.3. Bilateral Collaborations

7.3.1. France-USA

We have an ongoing collaboration with MIT and Adobe Research, in the context of the hair BRDF estimation project (Fredo Durand and Sylvain Paris). This resulted in a EGSR 09 paper [16].

We have recently initiated a new collaboraton with Yale University (Holly Rushmeier and Julie Dorsey), on the topic of effective and efficient weathering techniques, in collaboration with Carles Bosch who is at Yale Oct 09-Mar 10 and will be an ERCIM postdoc at REVES in 2010.

7.3.2. France-Canada

Cecile Picard's visit to the McGill in Montreal, Quebec, Canada resulted in a succesful collaboration with Paul Kry, resulting in the publication [22].

7.3.3. France-Korea

With the arrival of Pierre-Yves Laffont, we have initiated a collaboration with the Scalable Graphics Lab in KAIST, South Korea (Sung-Eui Yoon), on the topic of interactive content-aware image resizing and zooming.

7.3.4. France-Belgium

In the context of the Gabor Noise project we collaborated with Phil Dutre and the Catholic University of Leuven. This resulted in a SIGGRAPH paper [18] and the arrival of two postdocs (P. Vangorp and A. Lagae) from this lab.

7.3.5. France-Germany

In the context of the Master's thesis of Matthäus G. Chajdas, we continuted our long-standing collaboration with the University of Erlangen (Marc Stamminger), resulting in a submitted publication.

7.3.6. France-Spain

Sylvain Lefebvre started a collaboration with Gustavo Patow (researcher) and Ismael Garcia (PhD student) of Girona University, Spain, on the topic of dynamic tiletree; a method to progressively texture objects as they appear or are modified on screen. Ismael Garcia is currently visiting the Alice team in INRIA Nancy which S. Lefebvre has joined.

8. Dissemination

8.1. Participation in the Community

8.1.1. Program Committees

George Drettakis is chair of the program committee of SIGGRAPH Asia 2010, which will take place in Seoul (Korea) in December 2010; as part of his responsibilities in this context he also served on the SIGGRAPH and the SIGGRAPH Asia 2009 program committees.

In 2009, Sylvain Lefebvre served on the Eurographics Rendering Symposium papers committee and on the Symposium on Interactive 3D Graphics and Games committee. He was reviewer for SIGGRAPH, SIGGRAPH Asia, Eurographics paper and short paper tracks), VIS, HPG, CGA, Pacific Graphics and CAGD. He also reviewed an 'Equipe Associée' proposal for INRIA.

8.1.2. Invited Talks

Sylvain Lefebvre gave invited talks at the 2009 Italian EUROGRAPHICS chapter (EG-IT), the Mathematics and Image Analysis 2009 workshop (MIA) and the Paris chapter of ACM SIGGRAPH (15/12/2009). Sylvain Lefebvre presented the Gabor noise work at a seminar at INRIA Grenoble (June 8, 2009) and during a talk at the GT-Rendu on October 2nd, 2009. He also presented his research program at the CP of INRIA Nancy Grand-Est on May 4, 2009.

8.1.3. Thesis Committees

G. Drettakis was an evaluator for the habilitation thesis of Stéphane Mérillou (Limoges), and Ph.D. thesis of Adrien Bousseau (Grenoble).

8.1.4. Community service

G. Drettakis is a member of the “Bureau de Comité de Projets” and a member of the “comités d’animation scientifique” for the Interaction, Cognition and Perception theme of INRIA since October 2009. He also participated in the AERES evaluation of the IRIT laboratory in Toulouse (Nov. 2009).

8.1.5. Web server

Participants: George Drettakis, Sylvain Lefebvre.

http://www-sop.inria.fr/reves/

The project web-server is constantly updated with our research results. Most of the publications of REVES can be found online, and often include short movies demonstrating our research results. See

http://www-sop.inria.fr/reves/publis/. Sylvain Lefebvre installed the bibliography server Basilic (from Grenoble), which allows a much improved presentation of our publications.

8.2. Teaching

8.2.1. University teaching

George Drettakis is the organizer of the Computer Graphics class at Ecole Centrale Paris, and taught for 6 hours at this course. Together with Sylvain Paris they gave an introductory course (9 hours) image synthesis for the ENS master in Sophia-Antipolis (M1) and the master program of EPU UNSA (Jeux Video).

Sylvain Lefebvre also taught in the advanced Computer Graphics class at Ecole Centrale Paris, for 12 hours.

Nicolas Bonneel conducted the laboratory excercices for the M1 graphics class at the EPU UNSA in the Spring of 2009.

8.2.2. PhD Thesis Completed and Continuing

Cécile Picard was in her 3rd year working on physically-based audio synthesis. Her defense took place in December 2009 [15]. Nicolas Bonneel was also in his 3rd year as part of the CROSSMOD project working on audiovisual crossmodal effects. His defense took place in September 2009 [14]. Marcio Cabral was in his 2nd year working on procedural techniques for modelling and on image-based relighting.

8.3. Participation at conferences

8.3.1. Presentations at Conferences

David Grelaud presented his paper "Efficient and Practical Audio-Visual Rendering for Games using Crossmodal Perception" at the I3D conference in Boston in February 2009.

Sylvain Lefebvre presented the state of the art on example-based texture synthesis at the EUROGRAPHICS conference, Munich Germany, March 2009.

Marcio Cabral presented the paper "Structure-Preserving Reshape for Textured Architectural Scenes" at the Eurographics 2009 Conference in Munich, Germany - 30/03/2009 - 03/04/2009.

Nicolas Bonneel presented his paper "Single Photo Estimation of Hair Appearance" at the Eurographics Symposium on Rendering in Girona, June 2009.

Cécile Picard presented the paper [23] "Retargetting Example Sounds to Interactive Physics-Driven Animations" at the 35th AES International Conference - Audio for Games in London, UK, in Feb. 2009.

Cécile Picard presented the paper [22] "A Robust And Multi-Scale Modal Analysis For Sound Synthesis", at the 12th International Conference on Digital Audio Effects DAFx-09, Como, Italy in September 2009.

8.3.2. Participation at Conferences and Workshops

George Drettakis attended EGSR 2009 (Spain, June), SIGGRAPH (USA, August), and SIGGRAPH Asia 2009 (Yokohama, Japan, December). Sylvain Lefebvre attended Eurographics 2009 (Germany, April), EGSR 2009 (Spain, June), SIGGRAPH 2009 (USA, August) and Eurographics Italian Chapter (Italy, October). Nicolas Bonneel attended EGSR 2009 (Spain, June). Cécile Picard attended DAFx 2009 (Italie, September). Marcio Cabral attended Eurographics 2009 (Germany, April) and EGSR 2009 (Spain, June)

8.4. Demonstrations and Press

8.4.1. Demonstrations

Participants: David Grelaud, George Drettakis.

We demonstrated the workbench and the immersive space to ten members of the European project TEMPUS with the LOGNET team (May 11), as well as on numerous occasions to INRIA's people and/or visitors.

8.4.2. Press

Participant: George Drettakis.

An article on the research of REVES in 3D sound appeared in the magazine "Les Cahiers de l'Inria: La Rercherche" in February.

9. Bibliography

Major publications by the team in recent years

[1] N. BONNEEL, G. DRETTAKIS, N. TSINGOS, I. VIAUD-DELMON, D. JAMES. *Fast Modal Sounds with Scalable Frequency-Domain Synthesis*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", vol. 27, n^{o} 3, August 2008, http://www-sop.inria.fr/reves/Basilic/2008/BDTVJ08.

[2] C. DACHSBACHER, M. STAMMINGER, G. DRETTAKIS, F. DURAND. *Implicit Visibility and Antiradiance for Interactive Global Illumination*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", vol. 26, n^{o} 3, August 2007, http://www-sop.inria.fr/reves/Basilic/2007/DSDD07.

[3] O. DEUSSEN, C. COLDITZ, M. STAMMINGER, G. DRETTAKIS. *Interactive visualization of complex plant ecosystems*, in "Proceedings of the IEEE Visualization Conference", IEEE, October 2002, http://www-sop.inria.fr/reves/publications/data/2002/DCSD02/Vis2002Deussen.pdf.

[4] F. DUGUET, G. DRETTAKIS. *Robust Epsilon Visibility*, in "Proceedings of ACM SIGGRAPH 2002", J. HUGHES (editor), ACM Press / ACM SIGGRAPH, July 2002, http://www-sop.inria.fr/reves/publications/data/2002/DD02.

[5] F. DUGUET, G. DRETTAKIS. *Flexible Point-Based Rendering on Mobile Devices*, in "IEEE Computer Graphics and Applications", vol. 24, n^{o} 4, July-August 2004.

[6] P. ESTALELLA, I. MARTIN, G. DRETTAKIS, D. TOST. *A GPU-driven Algorithm for Accurate Interactive Reflections on Curved Objects*, in "Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)", T. AKENINE-MÖLLER, W. HEIDRICH (editors), Eurographics/ACM SIGGRAPH, June 2006, http://www-sop.inria.fr/reves/Basilic/2006/EMDT06.

[7] T. FUNKHOUSER, N. TSINGOS, I. CARLBOM, G. ELKO, M. SONDHI, J. WEST, G. PINGALI, P. MIN, A. NGAN. *A Beam Tracing Method for Interactive Architectural Acoustics*, in "The Journal of the Acoustical Society of America (JASA)", vol. 115, n^{o} 2, February 2004, http://www-sop.inria.fr/reves/Basilic/2003/FTCESWPMN03.

[8] X. GRANIER, G. DRETTAKIS. *A Final Reconstruction Framework for an Unified Global Illumination Algorithm*, in "ACM Transactions on Graphics", vol. 23, n^{o} 2, April 2004.

[9] A. RECHE, I. MARTIN, G. DRETTAKIS. *Volumetric Reconstruction and Interactive Rendering of Trees from Photographs*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", vol. 23, n^{o} 3, July 2004, http://www-sop.inria.fr/reves/publications/data/2004/RMD04.

[10] M. STAMMINGER, G. DRETTAKIS. *Interactive Sampling and Rendering for Complex and Procedural Geometry*, in "Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)", K. MYSZKOWSKI, S. GORTLER (editors), 12th Eurographics workshop on Rendering, Springer Verlag, Eurographics, 2001, http://www-sop.inria.fr/reves/Basilic/2001/SD01.

[11] M. STAMMINGER, G. DRETTAKIS. *Perspective Shadow Maps*, in "Proceedings of ACM SIGGRAPH 2002", J. HUGHES (editor), ACM Press/ ACM SIGGRAPH, July 2002, http://www-sop.inria.fr/reves/publications/data/2002/SD02.

[12] N. TSINGOS, E. GALLO, G. DRETTAKIS. *Perceptual Audio Rendering of Complex Virtual Environments*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", vol. 23, n^{o} 3, July 2004, http://www-sop.inria.fr/reves/publications/data/2004/TGD04.

[13] B. WALTER, G. DRETTAKIS, D. GREENBERG. *Enhancing and Optimizing the Render Cache*, in "Proceedings of the Eurographics Workshop on Rendering", P. DEBEVEC, S. GIBSON (editors), ACM Press, Eurographics, June 2002, http://www-sop.inria.fr/reves/Basilic/2002/WDG02.

Year Publications

Doctoral Dissertations and Habilitation Theses

[14] N. BONNEEL. *Audio and Visual Rendering with Perceptual Foundations*, Université de Nice - Sophia Antipolis, Sept. 2009, http://www-sop.inria.fr/reves/Basilic/2009/Bon09, Ph. D. Thesis.

[15] C. PICARD. *Expressive Sound Synthesis for Animation*, Université de Nice - Sophia Antipolis, Dec. 2009, http://www-sop.inria.fr/reves/Basilic/2009/Pic09, Ph.D. Thesis.

Articles in International Peer-Reviewed Journal

[16] N. BONNEEL, S. PARIS, M. VAN DE PANNE, F. DURAND, G. DRETTAKIS. *Single Photo Estimation of Hair Appearance*, in "Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)", 2009, http://www-sop.inria.fr/reves/Basilic/2009/BPVDD09.

[17] M. CABRAL, S. LEFEBVRE, C. DACHSBACHER, G. DRETTAKIS. *Structure Preserving Reshape for Textured Architectural Scenes*, in "Computer Graphics Forum (Proceedings of the Eurographics conference)", 2009, http://www-sop.inria.fr/reves/Basilic/2009/CLDD09 DE .

[18] A. LAGAE, S. LEFEBVRE, G. DRETTAKIS, P. DUTRÉ. *Procedural Noise using Sparse Gabor Convolution*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", vol. 28, n^{o} 3, August 2009, http://www-sop.inria.fr/reves/Basilic/2009/LLDD09 BE .

[19] C. SUIED, N. BONNEEL, I. VIAUD-DELMON. *Integration of auditory and visual information in the recognition of realistic objects*, in "Experimental Brain Research", 2009, http://www-sop.inria.fr/reves/Basilic/2009/SBV09.

International Peer-Reviewed Conference/Proceedings

[20] C. CRASSIN, F. NEYRET, S. LEFEBVRE, E. EISEMANN. *GigaVoxels : Ray-Guided Streaming for Efficient and Detailed Voxel Rendering*, in "Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games", 2009, http://hal.inria.fr/inria-00345899/en/DE.

[21] D. GRELAUD, N. BONNEEL, M. WIMMER, M. ASSELOT, G. DRETTAKIS. *Efficient and Practical Audio-Visual Rendering for Games using Crossmodal Perception*, in "Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games", 2009, http://www-sop.inria.fr/reves/Basilic/2009/GBWAD09 AT .

[22] C. PICARD, F. FAURE, G. DRETTAKIS, P. KRY. *A Robust and Multi-Scale Modal Analysis For Sound Synthesis*, in "Proc. of the 12th Int. Conference on Digital Audio Effects (DAFx-09), Italie Como", 2009, http://hal.inria.fr/inria-00410314/en/CA.

[23] C. PICARD, N. TSINGOS, F. FAURE. *Retargetting Example Sounds to Interactive Physics-Driven Animations*, in "n AES 35th International Conference, Audio in Games., Royaume-Uni London", 2009, http://hal.inria.fr/inria-00394469/en/.

Workshops without Proceedings

[24] L.-Y. WEI, S. LEFEBVRE, V. KWATRA, G. TURK. *State of the Art in Example-based Texture Synthesis*, in "Eurographics 2009, State of the Art Report, EG-STAR", Eurographics Association, 2009, http://www-sop.inria.fr/reves/Basilic/2009/WLKT09 US .

References in notes

[25] N. BONNEEL, C. SUIED, I. VIAUD-DELMON, G. DRETTAKIS. *Bimodal perception of audio-visual material properties for virtual environments*, in "ACM Transactions on Applied Perception", 2010, http://www-sop.inria.fr/reves/Basilic/2009/BSVD09, to appear.

[26] O. DEUSSEN, C. COLDITZ, M. STAMMINGER, G. DRETTAKIS. *Interactive visualization of complex plant ecosystems*, in "Proceedings of the IEEE Visualization Conference", IEEE, October 2002, http://www-sop.inria.fr/reves/publications/data/2002/DCSD02/Vis2002Deussen.pdf.

[27] J. DORSEY, H. K. PEDERSEN, P. HANRAHAN. *Flow and Changes in Appearance*, in "ACM Computer Graphics (SIGGRAPH'96 Proceedings)", Aout 1996, p. 411–420.

[28] F. DUGUET, G. DRETTAKIS. *Robust Epsilon Visibility*, in "ACM Transactions on Computer Graphics (Proceedings of ACM SIGGRAPH 2002)", July 2002, http://www-sop.inria.fr/reves/publications/data/2002/DD02.

[29] F. DURAND, G. DRETTAKIS, C. PUECH. *The 3D Visibility Complex, a new approach to the problems of accurate visibility*, in "Rendering Techniques'96 (7th Eurographics Workshop on Rendering)", Springer Verlag, June 1996, p. 245–257.

[30] F. DURAND, G. DRETTAKIS, C. PUECH. *The Visibility Skeleton: A Powerful and Efficient Multi-Purpose Global Visibility Tool*, in "ACM Computer Graphics (SIGGRAPH'97 Conference Proceedings)", Aout 1997.

[31] F. DURAND, G. DRETTAKIS, C. PUECH. *Fast and Accurate Hierarchical Radiosity Using Global Visibility*, in "ACM Transactions on Graphics", vol. 18, April 1999, p. 128–170.

[32] T. FUNKHOUSER, I. CARLBOM, G. ELKO, G. PINGALI, M. SONDHI, J. WEST. *A Beam Tracing Approach to Acoustic Modeling for Interactive Virtual Environments*, in "ACM Computer Graphics (SIGGRAPH'98 Proceedings)", July 1998.

[33] T. FUNKHOUSER, N. TSINGOS, I. CARLBOM, G. ELKO, M. SONDHI, J. WEST. *Modeling Sound Reflection and Diffraction in Architectural Environments with Beam Tracing*, in "Forum Acusticum", September 2002.

[34] C. LOSCOS, G. DRETTAKIS, L. ROBERT. *Interactive Virtual Relighting of Real Scenes*, in "IEEE Transaction of Visualization and Computer Graphics", vol. 6, October 2000, p. 289–305.

[35] C. LOSCOS, M. FRASSON, G. DRETTAKIS, X. GRANIER, B. WALTER. *Interactive Virtual Relighting and Remodeling of Real Scenes*, in "Rendering Techniques'99 (10th Eurographics Workshop on Rendering)", Springer Verlag, June 1999, p. 329–340.

[36] G. MILLER. *Efficient Algorithms for Local and Global Accessibility Shading*, in "ACM Computer Graphics (SIGGRAPH'94 Proceedings)", July 1994, p. 319–326.

[37] J. PORTILLA, E. P. SIMONCELLI. *A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients*, in "International Journal of Computer Vision", vol. 40, n^o 1, 2000, p. 49-70.

[38] N. TSINGOS, T. FUNKHOUSER, I. CARLBOM. *Modeling Acoustics in Virtual Environments Using the Uniform Theory of Diffraction*, in "ACM Computer Graphics (SIGGRAPH 2001 Proceedings)", July 2001.

[39] N. TSINGOS. *Artifact-free asynchronous geometry-based audio rendering*, in "ICASSP'2001", May 2001.

[40] B. WALTER, G. DRETTAKIS, S. PARKER. *Interactive Rendering using the Render Cache*, in "Rendering Techniques'99 (10th Eurographics Workshop on Rendering)", Springer Verlag, June 1999.

[41] Y. YU, P. DEBEVEC, J. MALIK, T. HAWKINS. *Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs*, in "ACM Computer Graphics (SIGGRAPH'99 Proceedings)", 1999.