INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team REVES

Rendering and Virtual Environments with Sound

Sophia Antipolis

THEME COG

Activity Report

2006

Table of contents

1. Team

Head of project-team
George Drettakis [Research Director (DR) Inria, HdR]

Vice-head of project team
Nicolas Tsingos [Research Associate (CR) Inria]

Member of project team
Sylvain Lefebvre [Research Associate (CR) Inria]

Postdoctoral Fellows
Guillaume Lemaitre [Post-doctoral fellowship RNTL/OPERA]
Carsten Dachsbacher [Post-doctoral fellowship]

Technical staff
David Geldreich [Part time IR, Workbench/DREAM]

Ph.D. students
Emmanuel Gallo [CIFRE Scholarship with the CSTB]
Nicolas Bonneel [CROSSMOD project]

Research scientist (external collaborators)
Kyle Hegeman [Ph.D. Stonybrook University]
Pau Estalella [Ph.D. student University of Girona]
Christian Linz [Ph.D. student University of Braunschweig]
Simon Premoze [Postdoctoral fellow, Columbia]
Ignacio Martin [Researcher University of Girona]
Dani Tost [Prof. Polytechnic University of Catalunya]
Marcus Magnor [Prof. University of Braunschweig]
Marc Stamminger [Prof. University of Erlangen]
Michael Ashikhmin [Assistant Professor, Stonybrook University]

Administrative assistant
Marie-Line Ramfos [Secretary (SAR) Inria]
Aicha Venturini [Secretary (SAR) Inria]
Monique Varlet [CROSSMOD assistant]

Interns
Julien Guyetat [CROSSMOD, Engineer, EIGM]
Raphaël Armati [Engineer, Polytech'UNSA]
Thomas Moeck [CROSSMOD, Master Erlangen University]
Christian Eisenacher [Master Erlangen University]

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.

3. Scientific Foundations

3.1. Rendering

Keywords: *high-quality rendering*, *image rendering*, *plausible rendering*, *rendering*, *sound 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 outoors 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 is 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 fucntion or a series of functions, which are often generating 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, or clouds). Other representations merit further research, such as image or video based rendering algorithms, or structures/algorithms such as the "render cache" [50], 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, acheology etc.).

We are also investigating methods for adding "natural appearance" to synthetic objects. Such approaches include *weathering* or *aging* techniques, based on physical simulations [38], but also simpler methods such as accessibility maps [47]. The approaches we intend to investigate will attempt to both combine and simplify existing techniques, or develop novel approaches which are based 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.

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 [48] 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 [42] [41], [40] 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 [39]. Lazy evaluation, as well as hierarchical solutions are clearly interesting avenues of research, although are probably quite application dependent.

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [37], 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.

3.1.2.3. 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 [43] [48]), 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 [44] 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).

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 [49].

3.2. Virtual and Augmented Environments with Sound

Keywords: *augmented environments*, *augmented reality*, *auralisation*, *inverse rendering*, *re-lighting*, *sound "ambiance"*, *virtual environments*, *virtual reality*.

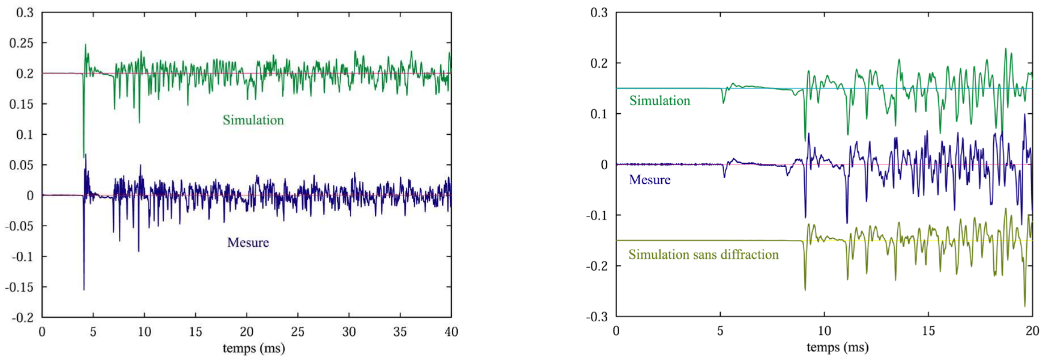

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

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 maniuplate 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.

Our previous work on interior relighting has given satisfactory solutions, allowing us to add virtual object with consistent lighting, but implies severe restrictions on the lighting conditions of input images [45], [46]. 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.

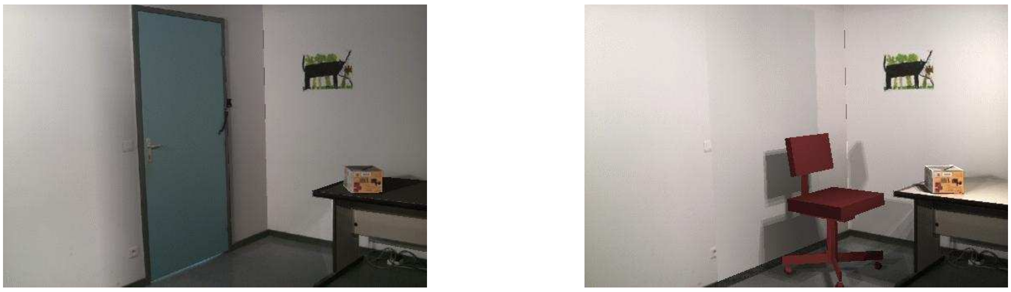

Figure 3. (a) Original conditions (b) The door has been removed virtually, and a virtual object and light have been added (method of [46])

This long term vision will require a way to smoothly vary from low-quality methods [45], [46] to high quality approaches [51], 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, there is yet to be a standard which has been widely adopted. 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. AURELI: AUdio REndering LIbrary/AUDILE

Participants: Nicolas Tsingos, Emmanuel Gallo, Thomas Moeck.

REVES is developing an API, AURELI (AUdio REndering LIbrary), as a tool supporting our research in acoustic modeling and audio rendering. Several prototype algorithms for sound spatialization, geometrical and statistical reverberation modeling, sound source clustering and audio rendering server have been implemented using AURELI's core functionnalities or as an extension to the API itself. Core fonctionalities include audio i/o plug-ins, audio buffer handling and basic signal processing. Higher level functions perform geometrical processing and audio rendering on a variety of restitution systems. AURELI is a cross-platform, object oriented, C++ API. It runs on LINUX/Windows/IRIX and also features primitives for parallel signal processing on multi-processor systems and network communication (used for instance to design audio rendering servers).

We are investigating possibilities for public release of the API as a tool for researchers in acoustics and virtual/augmented reality.

AUDILE is the codename for the clustering, perceptual masking and scalable processing technology, part of which has been transferred to EdenGames. As part of the work on CROSSMOD, an independent library has been created which interfaces to both OpenAL and AURELI in a transparent manner. This software is available on the CROSSMOD repository server.

4.2. OgreVR, OgreAL and OgreVRScript

Participants: David Geldreich, Nicolas Bonneel, Raphaël Armati, Julein Gueytat, George Drettakis.

In the context of his DREAM "mission" David Geldreich has written a framework (OgreVR) to use Ogre3D on virtual reality platform. OgreVR abstracts displays/devices setups through a configuration file. OgreVR/Ogre3D is the new software basis for our VR platform and is currently being used by several projects within the group. It is avalaible on the REVES project CVS and is available in SDK form in the CROSSMOD SVN.

As part of his internship, J. Gueytat, developed *OgreAL*, the integration of OpenAL into Ogre3D. This involves two main aspects: A binding for sound sources and their control, which is accessible to the graphics application, and an update phase when graphics objects are moved so that their sound equivalent is also moved. This integration lets us test and develop sound and graphics algorithms within the OgreVR framework.

Concrete experiments require not only 3D models, but also a specific interaction behavior. The goal is to allow also non-CG-experts (e.g., Neuroscience partners of the CROSSMOD project) to specify the interaction in a 3D environment without the need to actually program and compile the code.

During his summer internship, R. Armati, in collaboration with D. Geldreich, developed a set of specifications and requirements for such an authoring tool, in the context of the CROSSMOD project. Different alternatives have been then evaluated. The final decision was to take an existing scripting language (XP) and to re-implement it in Ogre. We also extended the scripting language, leading to the new name OgreVRScript. OgreVRScript allows specifying both the audio and the visual part of an experiment, and the various interactions which are necessary. An important feature of this environment is its extensibility, which allows the programmers to add new features as they are requested for the experimental work. These extensions are automatically available from the scripting language, permitting the continued usage by non-programmers.

OgreAL and OgreVRScript are available on the CROSSMOD SVN.

4.3. LibSL - Simple Library for Graphics

Participant: Sylvain Lefebvre.

Computer Graphics research requires a variety of specialized tools: image and mesh manipulation, interaction with several graphics API on several platforms, graphics processor (GPU) programming.

To ease programming in this context, Sylvain Lefebvre has developed a graphics-programming toolbox, as a C++ library. It simplifies several programming tasks and ease sharing and compilation of code under various environments. This development was started within the EVASION / INRIA Grenoble team and is now continuing within the REVES team. The library is available as a private project on the INRIA Forge.

5. New Results

5.1. Plausible Image Rendering

5.1.1. Interactive Diffuse and Glossy Indirect Illumination

Participants: Carsten Dachsbacher, Marc Stamminger.

We developed a novel method for plausible real-time rendering of indirect illumination effects for diffuse and non-diffuse surfaces. The scene geometry causing indirect illumination is captured by an extended shadow map, as proposed in previous work, and secondary light sources are distributed on directly lit surfaces. The rendering of indirect illumination is done using a GPU-friendly deferred shading method. We use different ways to obtain secondary light sources, responsible for the indirect lighting, and show how their contribution can be rendered efficiently. An importance sampling strategy, implemented entirely on the GPU, allows efficient selection of secondary light sources. Adapting the light's splat shape to surface glossiness also allows efficient rendering of caustics. Unlike previous approaches the approximated indirect lighting does barely exhibit coarse artifacts - even under unfavorable viewing and lighting conditions. We describe an implementation on contemporary graphics hardware and present adaptation to and results in game-typical applications.

This work was published as a chapter [28] in the book ShaderX5, published in 2006.

5.1.2. Procedural Reproduction of Terrain Textures with Geographic Data

Participants: Carsten Dachsbacher, Tobias Bolch, Marc Stamminger.

Surface textures of high resolution and quality, either acquired from aerial or satellite imagery or computed using procedural models, are crucial for photorealistic terrain rendering. Procedural models provide a compact representation and can be evaluated at run-time. We developed an extension to an existing, GPU-friendly procedural texturing model, such that it can be fitted semiautomatically to real-world data. In order to increase realism and to account for geographic conditions, we also include temperature, solar radiation and rainfall distributions - simulated or modeled using measured data from gaging stations - into the reproduction process. The original surface texture is no longer required for rendering: instead a new texture of arbitrary resolution is synthesized at runtime. In addition to the compact procedural model we store elevation data, which is required anyway for the terrain rendering, and low-resolution geographic data. We show results of our method applied to a comparatively little cultivated region in Central Asia.

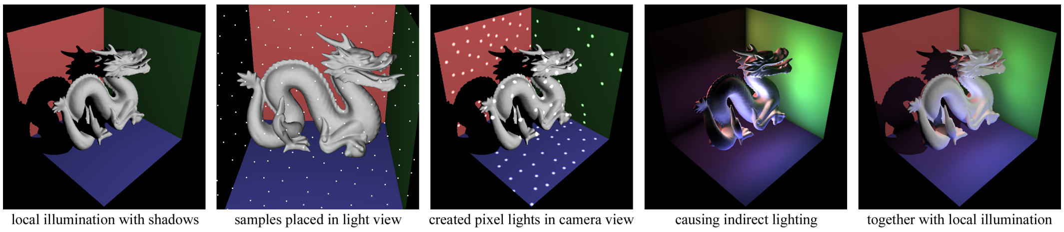

Figure 4. The splatting of indirect illumination: pixel lights are distributed on surfaces captured by the RSM. The region of significant contribution is computed and the indirect light is computed using deferred shading.

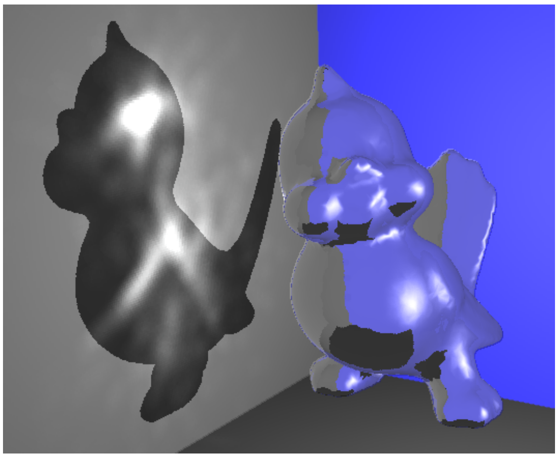

Figure 5. Caustics, obtained from single-sided refraction, can be rendered at real-time frame rates.

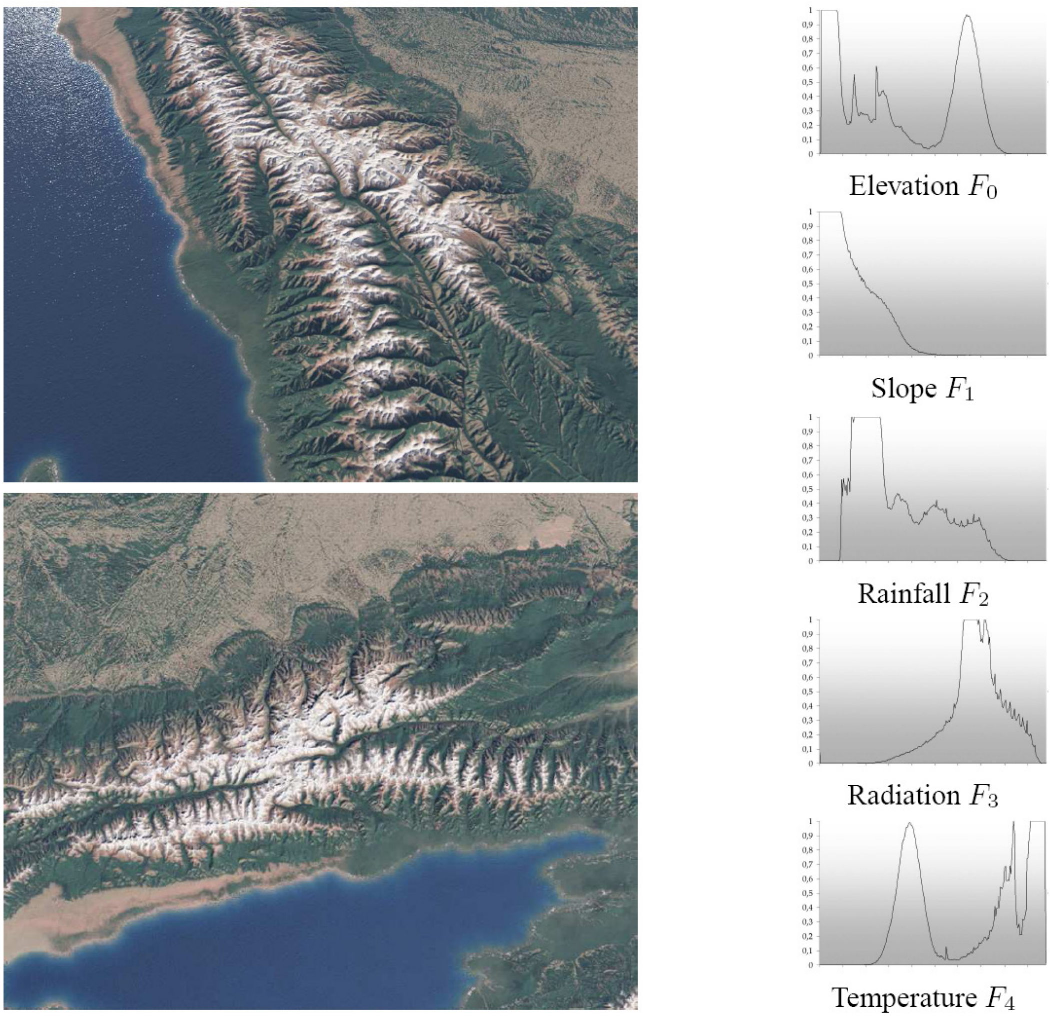

Figure 6. A procedural texture generated by our algorithm.

This work has been published at the VMV conference in 2006 [32].

5.1.3. Occludance

Participants: Carsten Dachsbacher, Marc Stamminger, George Drettakis.

A reformulation of the rendering equation removes visibility computations out of the main global illumination loop. This is achieved by introducing a new quantity called occludance. Our new approach results in a new global illumination algorithm, in which the bulk of the computation is shifted into simple local iterations. As a result, this computation is easy to parallelize, and is GPU friendly. A direct consequence of our method is that we can treat moving or dynamic objects with much less additional computational overhead. We have implemented a diffuse "radiosity-style" global illumination algorithm based on these principles, which runs mainly on the GPU. Our results demonstrate interactive computation of indirect illumination, including moving objects or characters for visually interesting scenes.

5.1.4. A GPU-driven Algorithm for Accurate Interactive Reflections on Curved Objects

Participants: Pau Estalella, Ignacio Martin, George Drettakis, Dani Tost.

This work is a GPU-driven method for the fast computation of specular reflections on curved objects. For every reflector of the scene, the method computes a virtual object for every other object reflected in it. This virtual reflected object is then rendered and blended with the scene. For each vertex of each virtual object, a reflection point is found on the reflectors surface. This point is used to find the reflected virtual vertex, enabling the reflected virtual scene to be rendered. This method is based on a previous result we published at VMV 2005, which was however not amenable to hardware implementation and thus was much slower.

The method renders the 3D points and normals of the reflector into textures, and uses a local search in a fragment program on the GPU to find the reflection points. By reorganizing the data and the computation in this manner, and correctly treating special cases, we make excellent use of the parallelism and stream-processing power of the GPU. The results of the method show that, we can display high-quality reflections of nearby objects interactively. For an example see Figure 7.

This work has been published at the Eurographics Symposium on Rendering [33].

Figure 7. Left, a frame of an interactive sequence using our method, running at 82fps. Right, a ray-traced image (5 seconds to compute one frame)

5.1.5. Approximate Ambient Occlusion For Trees

Participants: Kyle Hegeman, Simon Premoze, Michael Ashikhmin, George Drettakis.

Natural scenes contain large amounts of geometry, such as hundreds of thousands or even millions of tree leaves and grass blades. Subtle lighting effects present in such environments usually include a significant amount of occlusion effects and lighting variation. These effects are important for realistic renderings of such natural environments; however, plausible lighting and full global illumination computation come at prohibitive costs especially for interactive viewing. As a solution to this problem, we present a simple approximation to integrated visibility over a hemisphere (ambient occlusion) that allows interactive rendering of complex and dynamic scenes. Based on a set of simple assumptions, we show that our method allows the rendering of plausible variation in lighting at modest additional computation and little or no precomputation, for complex and dynamic scenes.

This work has been published at the 2006 I3D conference [34].

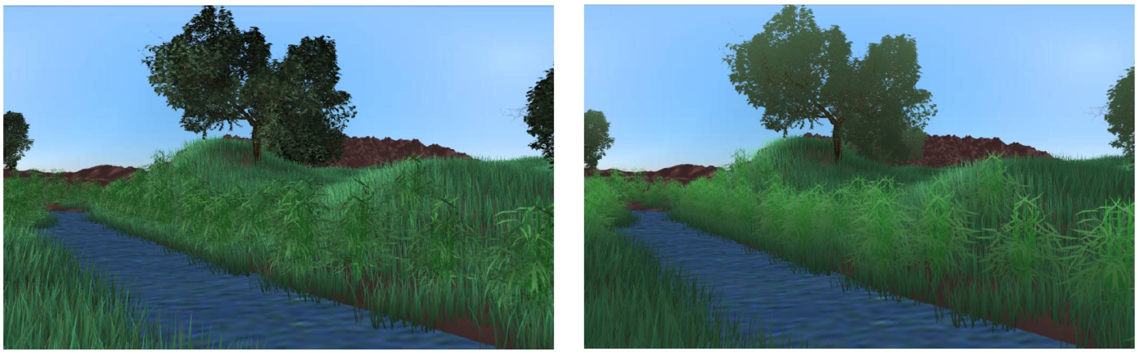

Figure 8. (a) Image of an ecosystem rendered with our approximation. (b) The Reeves-Blau approximation. We can clearly see the plausible variation in contrast in our approximation compared to that of the previous method.

5.1.6. Effective Multi-resolution Rendering and Texture Compression for Captured Volumetric Trees

Participants: Christian Linz, Alex Reche-Martinez, George Drettakis, Marcus Magnor.

Trees can be realistically rendered in synthetic environments by creating volumetric representations from photographs. However, volumetric tree representations created with previous methods are expensive to render due to the high number of primitives, and have very high texture memory requirements.

This project is a direct extension of the work developed at REVES in 2004 in the context of Alex Reche's thesis. This work addresses the two shortcomings of that method, that is the rendering speed and excessive texture memory requirements. We resolve these problems with an efficient multi-resolution rendering method and an effective texture compression solution.

The method uses an octree with appropriate textures at intermediate hierarchy levels and applies an effective pruning strategy, replacing the previous spatial structure. For texture compression, we adapt a vector quantization approach in a perceptually accurate color space, and modify the codebook generation of the Generalized Lloyd Algorithm to further improve texture quality.

In combination with several hardware acceleration techniques, this approach achieves a reduction in texture memory requirements by one order of magnitude; in addition, it is now possible to render tens or even hundreds of captured trees at interactive rates. For an example see Figure 9.

This work has been published at the Eurographics Workshop on Natural Phenomena [35].

5.1.7. Tile-Trees

Participants: Sylvain Lefebvre, Carsten Dachsbacher.

Figure 9. A scene with 290 trees running at 12 fps, and requiring 2.9 MB of texture memory for 3 different types of trees. Using the best previous approach, several seconds are required to render frame and 641 MB texture memory are needed.

Texture mapping with atlases suffer from several drawbacks: Wasted memory, seams, uniform resolution and no support of implicit surfaces. Texture mapping in a volume solves most of these issues, but unfortunately it induces an important space and time overhead.

To address this problem, we introduce the Tile-Tree: A novel data structure for texture mapping surfaces. Tile-Trees store square texture tiles into the leaves of an octree surrounding the surface. At rendering time the surface is projected onto the tiles, and the color is retrieved by a simple 2D texture fetch into a tile map. This avoids the difficulties of global planar parameterizations while still mapping large pieces of surface to regular 2D textures. Our method is simple to implement, does not require long pre-processing time, nor any modification of the textured geometry. It is not limited to triangle meshes. The resulting texture has little distortion and is seamlessly interpolated over smooth surfaces. Our method natively supports adaptive resolution.

We show that Tile-Trees are more compact than other volume approaches, while providing fast access to the data. We also describe an interactive painting application, enabling to create, edit and render objects without having to convert between texture representations.

This work was submitted in November 2006 and accepted for publication at the I3D 2007 conference.

5.2. Plausible Audio Rendering

5.2.1. Audio matting

Participants: Nicolas Tsingos, Emmanuel Gallo.

We proposed a novel approach to real-time spatial rendering of realistic auditory environments and sound sources recorded live, in the field. Using a set of standard microphones distributed throughout a real-world environment we record the sound-field simultaneously from several locations. After spatial calibration, we segment from this set of recordings a number of auditory components, together with their location. We compared existing time-delay of arrival estimations techniques between pairs of widely-spaced microphones and introduced a novel efficient hierarchical localization algorithm. Using the high-level representation thus obtained, we can edit and re-render the acquired auditory scene over a variety of listening setups. In particular, we can move or alter the different sound sources and arbitrarily choose the listening position. We can also composite elements of different scenes together in a spatially consistent way.

Our approach provides efficient rendering of complex soundscapes which would be challenging to model using discrete point sources and traditional virtual acoustics techniques.

Example results are available at http://www-sop.inria.fr/reves/projects/audioMatting.

5.3. Virtual Environments with Sound

5.3.1. Progressive Perceptual Audio Rendering of Complex Scenes

Participants: Nicolas Tsingos, Thomas Moeck, Nicolas Bonneel.

Despite recent advances, including sound source clustering and perceptual auditory masking, high quality rendering of complex virtual scenes with thousands of sound sources remains a challenge. Two major bottlenecks appear as the scene complexity increases: the cost of clustering itself, and the cost of pre-mixing source signals within each cluster.

In this work, we first propose an improved hierarchical clustering algorithm that remains efficient for large numbers of sources and clusters while providing progressive refinement capabilities. We have also developed a lossy pre-mixing method based on a progressive representation of the input audio signals and the perceptual importance of each sound source. We performed quality evaluation user tests which indicate that the recently introduced audio saliency map is inappropriate for this task. Consequently we propose a “pinnacle", loudness-based metric, which gives the best results for a variety of target computing budgets. We also performed a perceptual pilot study which indicates that in audio-visual environments, it is better to allocate more clusters to visible sound sources. We propose a new clustering metric using this result. As a result of these three solutions, our system can provide high quality rendering of thousands of 3D-sound sources on a “gamer-style" PC.

6. Contracts and Grants with Industry

6.1. 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.

6.2. 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 new agreement will be used to fund a PhD thesis in order co-develop novel real-time synthesis algorithms for interactive audio rendering. The PhD will begin early January 2007.

6.3. CSTB

Participants: George Drettakis, Nicolas Tsingos, Emmanuel Gallo.

Another CIFRE contract with the CSTB for the thesis of Emmanuel Gallo. He works on sound rendering for urban simulation, in the context of the OPERA project. In this context, Emmanuel Gallo spends a significant amount of time at the CSTB and strives to apply his work directly on urban planning and architectural applications.

6.4. Renault

We have signed a collaboration agreement with Renault for potential evaluation of the audiomatting technique.

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 Olivier Warusfel of IRCAM. The collaboration with CNRS has helped us in the development of experimental protocols for perceptual validation work, and has resulted in a common paper submssion. Similarly, our collaboration with IRCAM has been very beneficial for our work in acoustics.

7.1.2. Collaboration with CSTB Sophia-Antipolis

Participants: George Drettakis, Nicolas Tsingos, Emmanuel Gallo.

The collaboration with CSTB continues through the PhD research of Emmanuel Gallo (CIFRE fellowship program).

7.1.3. The workbench platform

Participants: David Geldreich, George Drettakis, Nicolas Tsingos.

The regional Provence-Alpes-Cote d'Azur government has co-funded (with INRIA) the acquisition of semi-immersive platform for virtual environments, also known as "workbench". David Geldreich initially setup the system, and now continues the support.

The platform is composed of a Barco Baron screen (1.5 m diagonal) which can be tilted from near horizontal (table) to near vertical position. The screen is equiped with a BarcoReality 908 CRT projector driven by an off-the-shelf PC (2 Intel Xeon 2.8GHz + GeForce 6800 Ultra AGP 8x graphics) running under Linux and Windows XP. Stereo display is achieved through a frequency-doubler StereoGraphics EPC-2 and active LCD shutter-glasses (StereoGraphics CrystalEyes and NuVision/60GX). Finally, we also use a 6-DOF Polhemus Fastrak 3D tracking system interfaced with a stylus and a gamepad for interaction, and an additional captor for view-point tracking and view-dependent rendering.

For the audio, we use six Yamaha MSP3 speakers driven by a Motu 896 firewire audio interface.

D. Geldreich installed the system and developed a suite of APIs and tools allowing for easy integration of the stereo display and tracking system in any Performer/VTK- based application. This year, we upgraded our tracking facility to an ART, vision-based system.

7.2. National Projects

7.2.1. RNTL project OPERA: PErceptual Optimizations for Audio Rendering

Participants: Nicolas Tsingos, Guillaume Lemaitre, Emmanuel Gallo.

The OPERA project successfully ended in May 2006 with the presentation of the two demo applications "EAR" and "CHAT". The project lead to 4 publications and 3 others currently submitted to publication. (see http://www-sop.inria.fr/reves/OPERA for additional info).

7.3. European Projects

7.3.1. Open FET IST CROSSMOD

Participants: George Drettakis, Monique Varlet, David Geldreich, Nicolas Tsingos, Nicolas Bonneel, Thomas Moeck.

Figure 10. Screenshots of the CHAT demo. The video game "Flower Power Shooter" from Virtools was augmented with on-line 3D chat capabilities. The voice over IP servers used for streaming the speech data implements on-the-fly auditory masking evaluation from the point of view of each partipant in order to reduce the bandwidth.

Figure 11. Screenshots of the EAR demo featuring interactive road traffic simulation with pre-computed reflections and diffractions off the surrounding buildings.

CROSSMOD, "Cross-modal Perceptual Interaction and Rendering", is an Open FET (Future and Emerging Technologies) STREPS of the 6th FP IST. REVES coordinates this project; the other partners are CNRS, IRCAM (FR), University of Bristol (UK), Vienna University of Technology (A) University of Erlangen (D) and the CNR in Pisa (I).

CROSSMOD has three main objectives: the advancement of fundamental understanding of cross-modal effects for the display of virtual environments (VEs), the development of novel algorithms based on this understanding, and their evaluation using target applications. We concurrently develop the fundamental and foundational aspects related to audiovisual perception and display, as well as the novel and concrete algorithmic solutions. The scope of CROSSMOD is thus quite fundamental; however, its success will be measured by the efficiency/quality improvements of the new algorithmic solutions developed, and the validation of our hypotheses in realistic VE settings.

Work in the first year of CROSSMOD was structured in four main phases: a cross-disciplinary bibliography reading, brainstorming informed by the reading, specification of experiments, setups and software and finally pilot studies, setups and code development.

REVES participated in the first phase by reading a number of papers including many in neurosciences; REVES also had the responsibility of setting up and maintaining the wiki used as a collaborative tool in this process. In the subsequent phase, REVES led the work in tasks concerning audio quality and level-of-detail, notably performing pilot studies to evaluate the audio quality metrics, including the recently introduced audio saliency map. An initial attempt at the development of an audio-visual clustering metric, based on ideas taken from ventriloquism research was also performed at REVES (see also 5.3.1).

REVES participates in three other research tasks.

- The first is the work on eye-tracking for automatic task map creation based on eye-tracking for which Nicolas Tsingos has the overall scientific responsibility.
- The second direction REVES participated in is how to adapt VR for neuroscience experiments. REVES hosted CNR, CNRS and IRCAM researchers and provided the system framework on the workbench for the researchers involved,
- The third direction involves perceptual evaluation methodologies to define appropriate and efficient visual differences predictors. These will be used for perceptual graphics quality assessment and level-of-detail control.

In terms of systems, INRIA developed the first set of basic software components, permitting the development of the initial audio-visual experiments. In particular REVES developed a VR software platform which can be used for experiments, and some basic audio spatialisation software. The software platform is based on Ogre3D for graphics, OpenAL for audio and VRPN for basic VR device support. Stereo and tracker-based rendering was built in-house in a portable manner, and is currently being ported to other partners sites as requested. REVES also designed and wrote OgreVRScript which is a scripting language to facilitate the development of virtual environments by the neuroscience/psychoacoustics partners of the project. Other systems work involves the development of prototypes for graphics level-of-detail control using Threshold Maps in Ogre3D, and a prototype implementation of the audio saliency map.

As project coordinator, we were responsible for the publication of all the internal wiki material and continued updates, the design and implementation of the initial website published in May and the coordination of the outsourcing for the final version of the website. Finally INRIA setup the SIUG questionnaire on the website, and participated in the writing of the first newsletter. REVES also is responsible for the overall coordination of the scientific and management activity of CROSSMOD. Given its exploratory and multidisciplinary nature, this is a challenging, but enriching, task.

7.4. Visiting Researchers

This year we had three major visits, involving Ph.D. students and Masters internships.

Frank Firsching from the University of Erlangen visited us for 3 months as a PhD exchange visitor. Christian Linz from the Max-Plank-Institut Informatik (Saarbruecken), now at TU Braunschweig visited us in April for 3 days to work on the effective tree rendering project (see Section 5.1.5).

Thomas Moeck and Christian Eisenacher, both from the University of Erlangen, are doing their master internship at REVES (May-December for Thomas Moeck and November 2006-April 2007 for Christian Eisenacher).

We also had visitors for short periods in the context of the CROSSMOD project: Matteo Dellepiane from CNR (Pisa) visited us for a week in August and a week in October, Isabelle Viaud-Delmon from CNRS, Olivier Warusful from IRCAM and Etienne Cortell from IRCAM visited us for 2 days in October.

We hosted several talks in the course of the year: Katerina Mania from University of Crete in February. Argiro Vatakis from Cambridge University, Crossmodal Research Laboratory, in September. Sylvain Paris from MIT in April. Paul Kry from EVASION / INRIA Grenoble in September. Fabio Pelaccini from Dartmouth University in June. Adrian Bousseau from ARTIS / INRIA Grenoble in May.

7.5. Bilateral Collaborations

7.5.1. France-Germany

Two students from the graphics group of Marc Stamminger at the University of Erlangen are doing their master thesis at REVES: Thomas Moeck and Christian Eisenacher. We also had a collaboration with Christian Linz from Max-Plank-Institut and Marcus Magnor from TU Braunschweig on the effective tree rendering project.

7.5.2. France-United States of America

The collaboration with Prof. Ashikhmin at Sunnybrook (NY) on ambient occlusion approximation has resulted in a publication this year. We also have an on-going collaboration with Hugues Hoppe from Microsoft Research.

7.5.3. France-Italy

In the context of the CROSSMOD project, we have developed a working relationship with CNR Pisa. This includes the visits of the Ph.D. student Matteo Dellepiane, working on head-scanning for HRTF simulation as well as issues in audio-visual synchrony.

7.5.4. France-Spain

We have continued to be in contact with the research groups of Girona and also UPC Barcelona and in particular Ignacio Martin. This year we worked with his Ph.D. student Pau Estaella, co-supervised with Dani Tost from the UPC Barcelona, leading to the publication [33].

8. Dissemination

8.1. Participation in the Community

8.1.1. Program Committees

George Drettakis was program committee member of SIGGRAPH 2006, EGSR 06, I3D 07. He also was reviewer for GI06. Nicolas Tsingos was program committee member of EGSR 2006, ACM MM art program 2006, SPM 2007. In 2006 he was also reviewer for SIGGRAPH, EGSR, EURASIP JASP, AES 30th conf, and SPM. Sylvain Lefebvre was reviewer for the TVCG and TOG journals, SIGGRAPH 2006, GI 2006 and Pacific Graphics 2006. Carsten Dachsbacher was a member of the SIBGRAPI 2006, GRAPP'07 and Shader X5 committees. He was reviewer for I3D 2007, SIGGRAPH 2006, Eurographics 2006, EGSR 2006, SIBGRAPI 2006, Computer Graphics Forum, Graphical Models, Pacific Graphics 2006, IEEE Vis 2006, and Shader X5.

8.1.2. Thesis Committees

G. Drettakis was an external examiner for the Ph.D. theses of C. Daschsbacher (Erlangen, February) and G. Papagiannakis at the University of Geneva (May).

8.1.3. COST and CUMIR

George Drettakis is the coordinator of the Working Group on International Relations of the INRIA national level COST (Scientific and Technological Orientation Coucil/Conseil d'Orientation Scientifique et Technologique). The specific working group is responsible for the evaluation of all INRIA-supported international activities, including the International Associated Team programme, and the various bi-lateral and regional international cooperations. The group also makes proposals to the INRIA national directorate on issues relating to general international policy. This activity involves the coordination and organisation of these evaluations, the coordination of a group of 6-8 researchers, one annual meeting of the group, and participation in the bi-monthly meetings of the COST.

Nicolas Tsingos is a member of the CUMIR, the researcher users group, serving as an interface between the researchers of the Sophia-Antipolis INRIA Unit and the computer services. This year, he has been responsible for several activities, including a poster for the CUMIR and the project to determine the renewal of hardware on the site.

8.1.4. COGD Evaluation

George Drettakis was responsible for the organisation of the COGD 4-year evaluation. This involved aiding in the selection on reviewers, and overall scientific responsibility for the evaluation. The evaluation took place May 17-19th in Paris.

8.1.5. Web server

Participant: George Drettakis.

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/publications/index.php3?LANG=gb

8.2. Teaching

8.2.1. University teaching

George Drettakis was responsible for the Computer Graphics course at ISIA (Ecole des mines) in 2005 / 2006 (15 hours), with the participation of Nicolas Tsingos (10 hours). He also taught a Computer Graphics course at Ecole Centrale de Paris with Sylvain Lefebvre and Nicolas Tsingos (6 hours each for a total of 18 hours).

8.2.2. PhD Thesis Completed and Continuing

Emmanuel Gallo started his PhD early April 2004 and will be defending in March 2007. He is working on audio rendering. Nicolas Bonneel started his PhD in August 2006 as part of the CROSSMOD project and will be working on audiovisual crossmodal effects.

8.3. Participation at conferences

8.3.1. Presentations at Conferences

Nicolas Tsingos gave a talk entitled "Representations progressives et rendu perceptif de scènes sonores complexes" (progressive representation and perceptually based sound rendering in complex environments) at "'Journees d'etudes sur la spatialisation 2006"' organised by Ircam and Telecom Paris on January 24 and 25 2006 (http://www.tsi.enst.fr/jes2006/).

8.3.2. Participation at Conferences and Workshops

G. Drettakis went to the Eurographics Symposium on Rendering 2006 conference. C. Dachsbacher went to the SIGGRAPH 2006 conference. E. Gallo and N. Tsingos went to the AES Convention in San Fransicso in October.

8.4. Demonstrations and Press

8.4.1. Demonstrations

Participants: David Geldreich, George Drettakis, Emmanual Gallo, Nicolas Tsingos.

We presented demonstrations to 14-year old school children in January and November as part of their "enterprise internships", to the students of ENS Lyon in May, to representatives of the the Optitec consortium in January and the the Prefet of Alpes Maritimes in October.

8.4.2. Traité de la RV (the treaty of Virtual reality)

Participant: Nicolas Tsingos.

Le traité de la RV (the treaty of Virtual reality) http://www-caor.ensmp.fr/interlivre/ is a 4-book series dealing with all aspects of virtual reality. It is currently edited in French but an English version should appear shortly.

This series are designed both for designers and users of virtual reality environments. It comprises a thorough state of the art report of all aspects of the domain : computer science, physics, physiology, psychology, ergonomy, philosophy, ethics... It can be used as a guide to help the reader designing a virtual reality project.

Three fundamental issues are explored :

- Analysis and modeling of human activity in real and virtual environments.
- Analysis and modeling of interfaces for user immersion and interaction
- Modeling and implementation of the virtual environment.

We contributed to 3 of the 4 books [29], [31], [30] in the series regarding audio interfaces and rendering aspects.
Book1: « l'Homme et l'environnement virtuel », sous la coordination d'Alain Berthoz et de Jean-Louis Vercher ; Chapitre 3 - Les sens de lhomme - with Philippe Fuchs, Moustapha Hafez, Mohamed Benali Koudja, Jean-Paul Papin and Olivier Warusfel.
Book2:: « Interfaçage, immersion et interaction en environnement virtuel », sous la coordination de Sabine Coquillart et de Jean-Marie Burkhardt ; Chapitre 15 - Dispositifs et interfaces de restitution sonore spatiale - with Olivier Warusfel
Book3: « Outils et modèles informatiques des environnements virtuels », sous la coordination de Jacques Tisseau et de Guillaume Moreau ; Chapitre 4 - Modèles pour le rendu sonore - with Olivier Warusfel

9. Bibliography

Major publications by the team in recent years

[1] T. BOUBEKEUR, F. DUGUET, C. SCHLICK. *Rapid Visualization of Large Point-Based Surfaces*, in "Proc. of the 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST)", M. MUDGE, N. RYAN, R. SCOPIGNO (editors). , Eurographics, Eurographics, November 2005, http://www-sop.inria.fr/reves/publications/data/2005/BDS05.

[2] 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.

[3] G. DRETTAKIS, M. ROUSSOU, M. ASSELOT, A. RECHE, A. OLIVIER-MANGON, N. TSINGOS, F. TECCHIA. *Participatory Design and Evaluation of a Real-World Virtual Environment for Architecture and Urban Planning*, Technical report, n^{o} 5479, INRIA Research Report, Sophia-Antipolis, January 2005, http://hal.inria.fr/inria-00070528.

[4] G. DRETTAKIS, M. ROUSSOU. *Real-World VR Applications and Computer Graphics/Sound Research: A Synergy for the FutureProceedings of the Intuition Workshop on Virtual Reality*, November 2005, http://www-sop.inria.fr/reves/publications/data/2005/DR05.

[5] G. DRETTAKIS, M. ROUSSOU, N. TSINGOS, A. RECHE, E. GALLO. *Image-based Techniques for the Creation and Display of Photorealistic Interactive Virtual Environments*, in "Proceedings of the Eurographics Symposium on Virtual Environments", S. COQUILLART, M. GOEBEL (editors). , June 2004, http://www-sop.inria.fr/reves/publications/data/2004/DRTRG04.

[6] 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.

[7] 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, http://www-sop.inria.fr/reves/publications/data/2004/DD04.

[8] F. DUGUET. *Rendu et reconstruction de très gros nuages de points*, Ph. D. Thesis, University of Nice Sophia-Antipolis, June 2005, http://www-sop.inria.fr/reves/publications/data/2005/Dug05.

[9] F. DUGUET, C. H. ESTEBAN, G. DRETTAKIS, F. SCHMITT. *Level of Detail Continuum for Huge Geometric Data*, Technical report, n^{o} 5552, INRIA Research Report, Sophia-Antipolis, April 2005, http://hal.inria.fr/inria-00070455.

[10] F. DUGUET, C. H. ESTEBAN, G. DRETTAKIS, F. SCHMITT. *Level of Detail Continuum for Huge Geometric DataPoster at SIGGRAPH 2005*, July 2005, http://hal.inria.fr/inria-00070455.

[11] P. ESTALELLA, I. MARTIN, G. DRETTAKIS, D. TOST, O. DEVILLERS, F. CAZALS. *Accurate Interactive Specular Reflections on Curved Objects*, in "Proceedings of VMV 2005", 2005, http://www-sop.inria.fr/reves/publications/data/2005/EMDTDC05.

[12] M.-C. FRASSON. *Decoupling Structure and Appearance for Example-Driven Detail Synthesis*, Ph. D. Thesis, Universite de Nice Sophia-Antipolis, June 2005, http://www-sop.inria.fr/reves/publications/data/2005/Fra05.

[13] P. FUCHS, M. HAFEZ, M. B. KOUDJA, J.-P. PAPIN, N. TSINGOS, O. WARUSFEL. *Traité de la réalité virtuelle, Tome I. Chapitre 3 - Les sens de l'homme.*, Presses de l'Ecole des Mines de Paris, 2005.

[14] E. GALLO, G. LEMAITRE, N. TSINGOS. *Prioritizing signals for selective real-time audio processing*, in "International Conference on Auditory Display (ICAD'05)", ICAD, July 2005, http://www-sop.inria.fr/reves/publications/data/2005/GLT05.

[15] 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, http://www-sop.inria.fr/reves/publications/data/2004/GD04.

[16] K. JACOBS, C. ANGUS, C. LOSCOS, J.-D. NAHMIAS, A. RECHE, A. STEED. *Automatic shadow detection and shadow generation for Augmented Reality*, in "Graphics Interface 2005", May 2005, http://www-sop.inria.fr/reves/publications/data/2005/JALNRS05.

[17] G. LEMAITRE, E. GALLO. *Descripteurs audio. Librairie Matlab SPL et fichiers de descriptions ".sig"*, Technical report, n^o RT-0317, INRIA, Sophia Antipolis, november 2005, http://hal.inria.fr/inria-00069862.

[18] A. RECHE, G. DRETTAKIS. *View Dependent Layered Projective Texture Maps*, in "Proceedings of Pacific Graphics 2003", J. ROKNE, R. KLEIN, W. WANG (editors). , IEEE Press, October 2003, http://www-sop.inria.fr/reves/publications/data/2003/RD03b.

[19] 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.

[20] A. RECHE. *Image-Based Capture and Rendering with Applications to Urban Planning*, Ph. D. Thesis, Universite de Nice Sophia-Antipolis, April 2005, http://www-sop.inria.fr/reves/publications/data/2005/Rec05.

[21] M. ROUSSOU, G. DRETTAKIS. *Can VR be Useful and Usable in Real-World Contexts? Observations from the Application and Evaluation of VR in Realistic Usage Conditions*, in "Proc. HCI International 2005, 1st International Conference on Virtual Reality", July 2005, http://www-sop.inria.fr/reves/publications/data/2005/RD05.

[22] M. STAMMINGER, G. DRETTAKIS. *Interactive Sampling and Rendering for Complex and Procedural Geometry*, in "Rendering Techniques 2001 (Proceedings of the Eurographics Workshop on Rendering 01)", K. MYSKOWSKI, S. GORTLER (editors). , Springer Verlag, Eurographics, 2001, http://www-sop.inria.fr/reves/publications/data/2001/SD01.

[23] 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.

[24] 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.

[25] N. TSINGOS. *Scalable Perceptual Mixing and Filtering of Audio Signals using an Augmented Spectral Representation*, in "Proc. of the 8th Int. Conference on Digital Audio Effects (DAFx05)", Madrid, Spain, September 2005, http://www-sop.inria.fr/reves/publications/data/2005/Tsi05.

[26] N. TSINGOS, O. WARUSFEL. *Traité de la réalité virtuelle, Tome II. Chapitre 15 - Dispositifs et interfaces de restitution sonore spatiale*, Presses de l'Ecole des Mines de Paris, 2005.

[27] N. TSINGOS, O. WARUSFEL. *Traité de la réalité virtuelle, Tome III. Chapitre 4 - Modèles pour le rendu sonore*, Presses de l'Ecole des Mines de Paris, 2005.

Year Publications

Articles in refereed journals and book chapters

[28] C. DACHSBACHER, M. STAMMINGER. *Interactive Diffuse and Glossy Indirect Illumination*, chap. XX, Charles River Media, 2006, http://www-sop.inria.fr/reves/publications/data/2006/DS06.

[29] P. FUCHS, M. HAFEZ, M. B. KOUDJA, N. TSINGOS, O. WARUSFEL. *Traité de la réalité virtuelle, Tome I. Chapitre 3 - Les sens de lhomme.*, chap. 3, Presses de lEcole des Mines de Paris, 2006, http://www-sop.inria.fr/reves/publications/data/2006/FHBTW06.

[30] N. TSINGOS, O. WARUSFEL. *Trait´e de la r´ealit´e virtuelle, Tome II. Chapitre 15 -Dispositifs et interfaces de restitution sonore spatiale.*, chap. 15, Presses de lEcole des Mines de Paris, 2006, http://www-sop.inria.fr/reves/publications/data/2006/TW06.

[31] N. TSINGOS, O. WARUSFEL. *Trait´e de la réalité virtuelle, Tome III. Chapitre 4 -Modèles pour le rendu sonore.*, chap. III-4, Presses de lEcole des Mines de Paris, 2006, http://www-sop.inria.fr/reves/publications/data/2006/TW06b.

Publications in Conferences and Workshops

[32] C. DACHSBACHER, T. BOLCH, M. STAMMINGER. *Procedural Reproduction of Terrain Textures with Geographic Data*, in "Vision Modeling and Visualization 2006", T. AACH, R. WESTERMANN (editors). , 2006, http://www-sop.inria.fr/reves/publications/data/2006/DBS06.

[33] P. ESTALELLA, I. MARTIN, G. DRETTAKIS, D. TOST. *A GPU-driven Algorithm for Accurate Interactive Reflections on Curved Objects*, in "Rendering Techniques'06 (Proc. of the Eurographics Symposium on Rendering)", T. AKENINE-MÖLLER, W. HEIDRICH (editors). , Eurographics/ACM SIGGRAPH, June 2006, http://www-sop.inria.fr/reves/publications/data/2006/EMDT06.

[34] K. HEGEMAN, S. PREMOZE, M. ASHIKHMIN, G. DRETTAKIS. *Approximate Ambient Occlusion For Trees*, in "Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games", C. SEQUIN, M. OLANO (editors). , ACM SIGGRAPH, March 2006, http://www-sop.inria.fr/reves/publications/data/2006/HPAD06.

[35] C. LINZ, A. RECHE, G. DRETTAKIS, M. MAGNOR. *Effective Multi-resolution Rendering and Texture Compression for Captured Volumetric Trees*, in "Eurographics Workshop on Natural Phenomena (2006)", EUROGRAPHICS (editor). , September 2006, http://www-sop.inria.fr/reves/publications/data/2006/LRDM06.

Internal Reports

[36] G. LEMAITRE, E. GALLO. *Les descripteurs d'un son Librairie Matlab SPL et fichiers de descriptions '.sig'*, Technical report, n^{o} RT-0317, INRIA, January 2006, https://hal.inria.fr/inria-00069862.

References in notes

[37] 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.

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

[39] 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.

[40] 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.

[41] 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.

[42] 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.

[43] 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.

[44] 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, http://www-sop.inria.fr/reves/publications/data/2002/FTCESW02.

[45] 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.

[46] 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.

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

[48] 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.

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

[50] 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.

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