INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team REVES

Rendering and Virtual Environments with Sound

Sophia Antipolis - Méditerranée

THEME COG

Activity Report

2008

Table of contents

1. Team

Research Scientist

George Drettakis [Team Leader, Research Director (DR) Inria, HdR]
Nicolas Tsingos [Research Associate (CR) Inria, until May 4, HdR]
Sylvain Lefebvre [Research Associate (CR) Inria]

Technical Staff

Fernanda Andrade Cabral [Modeller, CROSSMOD project, since March]
Manuel Asselot [Specialist Engineer, 2 year contract INRIA-funded]
David Grelaud [Associated Engineer, 2 year contract INRIA-funded]
Fanourios Moraitis [Expert Engineer/Modeller, until March]

PhD Student

Nicolas Bonneel [CROSSMOD project]
Marcio Cabral [INRIA CORDIS scholarship]
Cécile Picard [INRIA/EdenGames collaboration]

Post-Doctoral Fellow

Efstathios Stavrakis [INRIA Post-Doctoral Fellowship]

Visiting Scientist

Michiel van de Panne [Visiting Scientist, Associate Professor Univ. of British Columbia, Canada, since August]

Administrative Assistant

Sophie Honnorat [Administrative Assistant (AI) Inria]

Other

Monique Varlet [CROSSMOD assistant]
Georgios Koulieris [Intern, Univ. of Athens, since October]
Marcel Meili [Intern, ETH Zurich]

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 of the year

Nicolas Tsingos defended his habilitation (HDR) in April 2008 [14].

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 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" [44], 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 [32], but also simpler methods such as accessibility maps [41]. 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 [31], 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 [42] 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 [36], [35], [34] 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 [33]. 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 [37] [42]), 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 [38] 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).

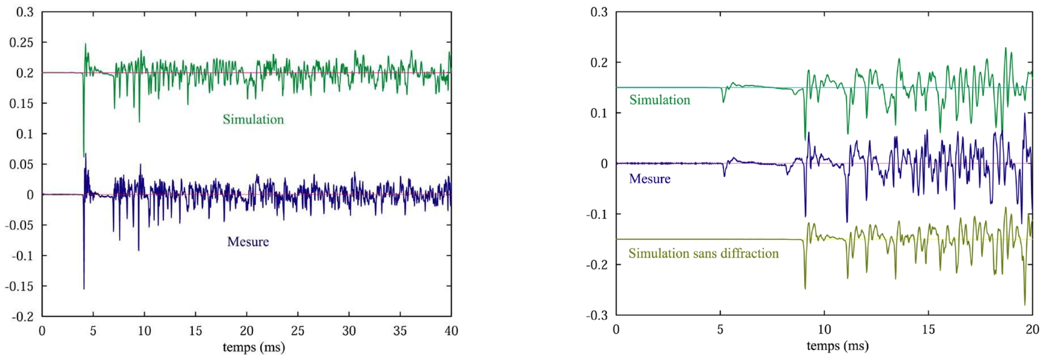

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and 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 [43].

3.2. Virtual and Augmented Environments with Sound

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

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.

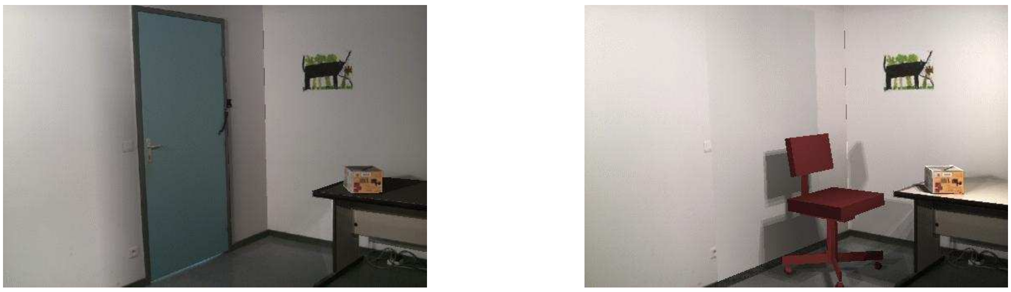

Figure 3. (a) Original conditions (b) The door has been removed virtually, and a virtual object and light have been added (method of [40])

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 [39], [40]. 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 [39], [40] to high quality approaches [45], 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. AURELI: AUdio REndering LIbrary/AUDILE

Participants: Nicolas Tsingos, Manuel Asselot, Thomas Moeck, Nicolas Bonneel.

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 functionalities or as an extension to the API itself. Core functionalities 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 the European project 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 OgreVRAL

Participants: Nicolas Bonneel, David Grelaud, George Drettakis.

In the context of his DREAM "mission" David Geldreich has written a framework (OgreVR) to use Ogre3D, an open-source graphics rendering engine, 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 availaible on the REVES project CVS and is available in SDK form in the CROSSMOD SVN. As part of a previous 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. OgreVRAL is 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. Texture Synthesis from Photographs

Participants: Christian Eisenacher, Sylvain Lefebvre.

The goal of texture synthesis is to generate an arbitrarily large high-quality texture from a small input sample. Generally, it is assumed that the input image is given as a flat, square piece of texture, thus it has to be carefully prepared from a picture taken under ideal conditions.

Instead, in this work we propose to extract the input texture from any surface from within an arbitrary photograph. This introduces several challenges: Only parts of the photograph are covered with the texture of interest, perspective and scene geometry introduce distortions, and the texture is non-uniformly sampled during the capture process. This breaks many of the assumptions used for synthesis.

To achieve our goal, we combine a simple novel user interface with a generic per-pixel synthesis algorithm to achieve high-quality synthesis from a photograph. Applications range from synthesizing textures directly from photographs to high-quality texture completion (Figure 4).

This work is the result of the master thesis of Christian Eisenacher, who stayed within the REVES team for six months. The paper was a collaboration with Marc Stamminger from the University of Erlangen and was published at the Eurographics conference in 2008 [22].

Figure 4. Top left: Image used as the source of texture. Bottom left: Image in which to synthesize. Right: Result.

5.1.2. Lazy Solid Texture Synthesis

Participants: Sylvain Lefebvre, George Drettakis.

A typical way to give an appearance to a surface is to define colors in a volume surrounding the object. The volume is then sampled in every point of the surface to determine its color. Algorithms have been proposed to automatically generate such volumes from a set of 2D example images: They *synthesize* a volume texture from 2D examples. However, existing approaches generate a **full** volume of color content, hence requiring a large amount of memory.

In this work, which is a collaboration with Microsoft Research Asia (Yue Dong and Xin Tong), we introduce a new algorithm with the unique ability to restrict synthesis to a subset of the voxels, while enforcing spatial determinism. This is especially useful when texturing objects, since only a thick layer around the surface needs to be synthesized. A major difficulty lies in reducing the dependency chain of neighborhood matching, so that each voxel only depends on a small number of other voxels.

The result is a new parallel, spatially deterministic solid texture synthesis algorithm which runs efficiently on the GPU. Our approach generates high resolution solid textures on surfaces within seconds. Memory usage and synthesis time only depend on the output textured surface area. The GPU implementation of our method rapidly synthesizes new textures for the surfaces appearing when interactively breaking or cutting objects (Figure 5).

This work has been published in the Computer Graphics Forum journal in 2008 [21] (special issue, proceedings of the Eurographics Symposium on Rendering).

5.1.3. Structure-Preserving Reshape for Textured Architectural Scenes

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

In this work, we present an approach for modeling architectural scenes by reshaping and combining existing textured models. Our reshape takes into account both geometry and texture, which are tightly coupled for architectural meshes. The problem is cast as a linear system which is solved for in real-time. 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.

Figure 5. A surface textured with our approach. The object is broken interactively. A texture is synthesized on-demand for the newly appearing surfaces.

This work has been accepted for publication at the Eurographics 2009 international conference [17].

5.1.4. Image Based Tree Relighting

Participants: Marcio Cabral, Sylvain Lefebvre, George Drettakis, Michiel van de Panne.

In this ongoing work we are investigating how to achieve image based relighting of photographs of trees. Our approach for relighting is based on an estimation of several parameters of the original lighting conditions. Based on these estimated values, a texture synthesis like technique will be used to produce the target relit photograph.

5.1.5. Efficient and Practical TileTrees

Participant: Sylvain Lefebvre.

The TileTree approach [28] combines benefits from both texture atlases and volumetric texturing approaches. TileTrees use a spatial data structure, e.g. an octree with low depth, to place square texture tiles around the surface to be textured. At rendering time the surface is projected onto these tiles and the color is retrieved by a simple 2D texture fetch from a tile map. This avoids the difficulties of global planar parameterizations while still mapping large pieces of surface to regular 2D textures. Beside a simple shader implementation for runtime lookups, the texture representation is compact, can be seamlessly interpolated and - this is interesting when used in a painting application - natively supports adaptive resolution.

In this follow-up work we explored several optimizations to the original TileTree data structure. We explain how to improve lookup efficiency with a simpler volume data structure and a packing strategy optimized for faster lookups. The result is a much faster access into the data structure with almost no additional memory overhead (Figure 6).

This work (in collaboration with Carsten Dachsbacher from the University of Stuttgart), was published as a chapter in the ShaderX 6 book, published in 2008 [28]. Each book in the ShaderX series is a collection of advanced rendering techniques aimed at Computer Graphics professionals.

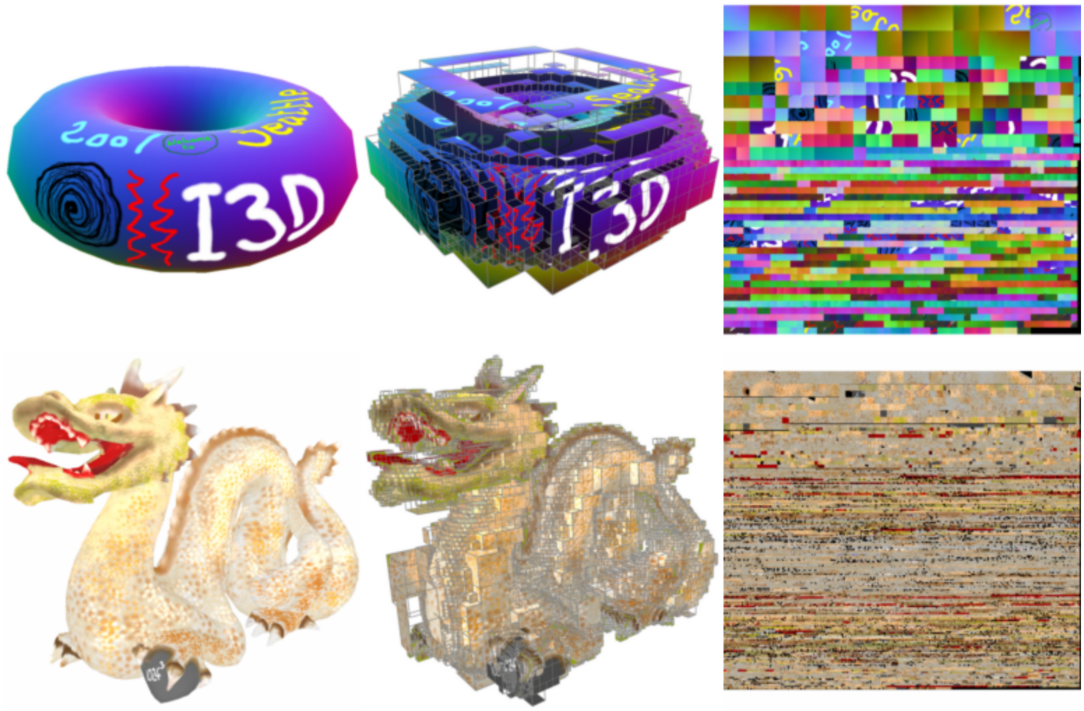

Figure 6. Two objects textured with a TileTree

5.1.6. Filtered Tilemaps

Participant: Sylvain Lefebvre.

Tilings have several applications in Computer Graphics. They are used to progressively load large textures, only loading those tiles that are used by the current viewpoint. They are also used to generate large, non repeating textures from a small set of precomputed square tiles.

Unfortunately, tilings present difficulties when it comes to filtering (bi-linear interpolation and MIP-mapping). The discontinuities at tile boundaries hinder direct filtering by the graphics hardware. Often, proper filtering implies the reimplementation of tri-linear filtering in the fragment program. This results in long shaders and possibly rendering artifacts. Moreover this hinders the use of anisotropic filtering which is crucial for terrain rendering.

In this article we analyze the difficulties associated with rendering tilings and present a simple yet generic approach to display properly filtered tilings (Figure 7). This work was published as a chapter in the ShaderX 6 book, published in 2008 [29]. Each book in the ShaderX series is a collection of advanced rendering techniques aimed at Computer Graphics professionals.

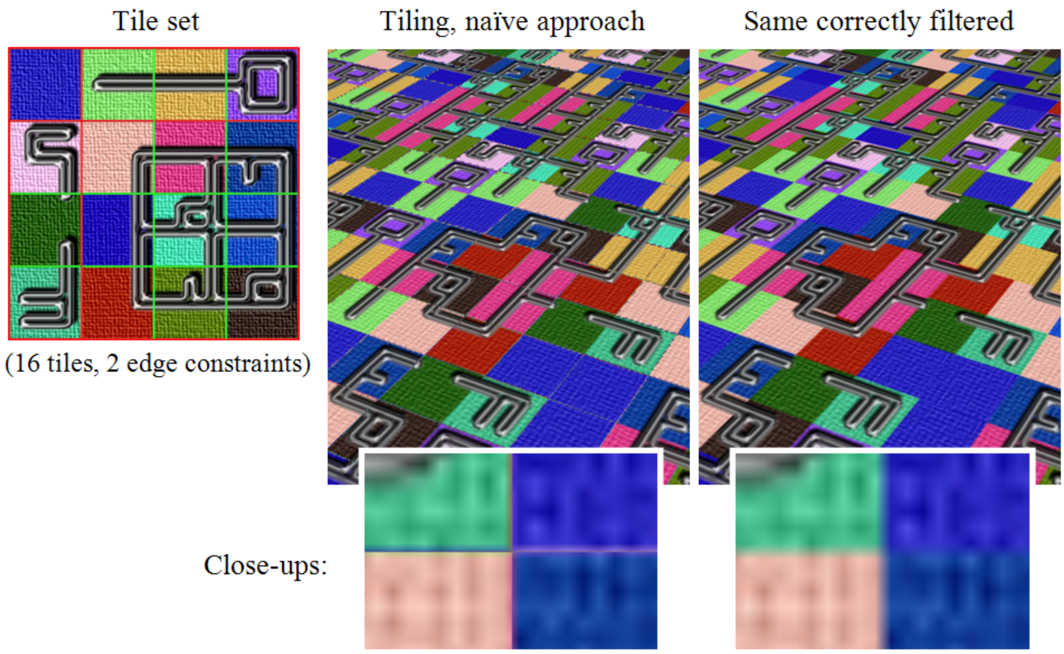

Figure 7. A constrained tile set is used to generate a large planar tiling. Left closeup: The default hardware filtering does not work properly across tiles, resulting in undesirable color bleeding. Right closeup: Using our method, interpolation and filtering works as expected.

5.1.7. Lightweight Analysis of Hair

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

This work in progress focuses on the acquisition of hair reflectance parameters from photographs. With current algorithms, we are able to quickly generate full globally illuminated hairs with a physical reflectance, requiring hand tweaking of reflectance parameters. We now use a machine learning approach to learn parameters from image based features, and apply this learning to real photographs in order to capture those parameters.

This work is currently in progress and is in collaboration with Fredo Durand (MIT) and Sylvain Paris (Adobe).

5.1.8. Hierarchical Hashing

Participant: Sylvain Lefebvre.

Perfect spatial hashing is a great tool to quickly locate sparse data in space. Unfortunately it does not allow for hierarchical data structures to be efficiently defined. On the other hand, octrees are widely used to capture hierarchical or multi-resolution data, but they lack the compactness of spatial hashing.

In this ongoing project we explore ways to define hierarchical data structures offering both the advantages of perfect spatial hashes and the versatility of octrees. We focus on fast and efficient construction - possibly incremental - and on efficient access from data parallel architectures such as GPUs. This work is a collaboration with Ares Lagae from the University of Leuven.

5.1.9. Progressive Texture Mapping of Dynamic Objects

Participants: Marcel Meili, Sylvain Lefebvre.

Texturing objects changing shape is a difficult problem. For instance, if an object is broken into several pieces, new surfaces appear. Applying an image to these new surfaces is uneasy, since most texture mapping approaches require a global optimization to compute the mapping between the surface and an image. The TileTree [28] data structure is a particularly interesting candidate for this case: It does not require any global optimization process, the surface is implicitly parameterized, and it only manipulates square 2D tiles of color data.

In this project we explored how the TileTree could be built dynamically, using latest parallel languages such as CUDA. Instead of building the TileTree for an entire object, we seek to progressively construct it, as the user visualizes the shape. Thus, the system makes no difference between a surface becoming visible for the first time and a newly created surface. Using our approach an object can be loaded and the user can start painting it with no set up time. The object could be sculpted at any time and painting could resume without interruption. The project focused on the first step of the system, that is the detection of the parts of the TileTree requiring an update.

This research was started in the context of the master internship of Marcel Meili, student at ETH Zurich, and supervised by Sylvain Lefebvre.

Figure 8. Oriental scene, showing the ability to handle large numbers of impact sounds.

5.2. Plausible Audio Rendering

5.2.1. Fast Modal Sounds with Scalable Frequency-Domain Synthesis

Participants: Nicolas Bonneel, George Drettakis, Nicolas Tsingos.

We developed a fast solution for contact sound rendering, exploiting the sparsity of modal sounds in the frequency domain, and taking into account asynchrony tolerance between the visual of an impact and its produced sounds to delay expensive computations. Our fast frequency domain modal synthesis performs 5 to 8 times faster than previous time domain modal synthesis. We proposed a fast energy estimation of contact sounds based on scalar products to allow integrating realtime generated contact sounds in a perceptual pipeline which previously only handled pre-recorded sounds. This pipeline allows masking between loud recorded sounds and dim impact sounds. It also allows handling clustering for fast spatialization of contact sounds. It finally allows scalable processing of impact sounds giving a global frequency bin budget to both impact and recorded sounds (Figure 8).

This work is a collaboration with Isabelle Viaud-Delmon from CNRS/IRCAM and Doug James from Cornell University. It was presented at SIGGRAPH 2008 (LA, USA, August) [15].

We also performed a pilot study indicating that better budgets indeed gave better quality, that high budgets gave acceptable quality compared to the reference rendering, and that a 200ms audio-visual delay for our scheduling strategy was not perceived as asynchronous more than 75% of the times.

5.2.2. Audio Texture Synthesis for Complex Contact Interactions

Participants: Cécile Picard, Nicolas Tsingos.

This project presents a new synthesis approach for generating contact sounds for interactive simulations. To address complex contact sounds, surface texturing is introduced. Visual textures of objects in the environment are reused as a discontinuity map to create audible position-dependent variations during continuous contacts. The resulting synthetic profiles are then used in real time to provide an excitation force to a modal resonance model of the sounding objects. Compared to previous sound synthesis for virtual environments, our approach has three major advantages: (1) complex contact interactions are addressed and a large variety of sounding events can be rendered, (2) it is fast due to the compact form of the solution which allows for synthesizing at interactive rates, (3) it provides several levels of detail which can be used depending on the desired precision (see Figure 9).

This work was a collaboration with François Faure from the Evasion/INRIA project-team in Grenoble, and has been accepted for publication at AES 35th International Conference [24].

5.2.3. Topological Sound Propagation with Reverberation Graphs

Participants: Efstathios Stavrakis, Nicolas Tsingos.

Reverberation effects due to sound scattering off wall surfaces carry major cues related to the size of the environment and distance to sound sources. Therefore, reverberation can help users establish a better sense of presence in virtual environments and thus is arguably one of the most important audio effects to simulate. In this work (in collaboration with Paul Calamia from Rensselaer Polytechnic Institute), we devised *Reverberation Graphs* [18]; a novel method to estimate global sound-pressure decay and auralize corresponding reverberation effects in 3D interactive virtual environments.

We use a 3D model composed of cells connected via portals (Figure 10(a)). First, we precompute pressure decay characteristics within each cell and between their coupling interfaces. Then at run-time, we identify sound propagation routes from sources to the listener (Figure 10(b)) and compute pressure decay envelopes (Figure 10(c)). Finally, our proposed scalable artificial reverberator uses these decay envelopes to auralize reverberation effects, including room coupling.

Our approach compares well with off-line geometrical techniques, but computes reverberation decay envelopes at interactive rates, ranging from 12 to 100 Hz. Our complete system can render as many as 30 simultaneous sources in large dynamic virtual environments.

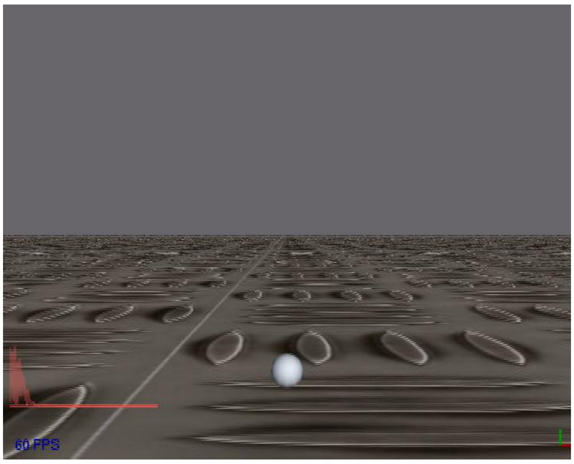

Figure 9. Illustration of the complex contact sound interaction approach.

Videos demonstrating the use of our technique within a game engine (Figure 10(d)) can be found online at: http://www-sop.inria.fr/reves/projects/revGraphs/.

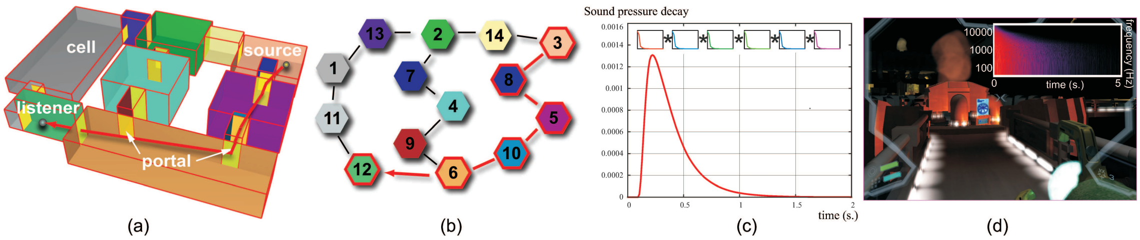

Figure 10. (a) A 3D scene, decomposed into cells and portals used to compute reverberation effects interactively. (b) Global pressure decays are estimated using the cell adjacency graph describing the 3D geometry. (c) The pressure decay along each route is computed by convolving individual decay envelopes of the route's cells and portals. An example route is highlighted in red in (b). (d) Our graph-based reverberation engine has been integrated into a 3D game with scalable reverberation processing and supports up to 30 sources interactively.

5.2.4. Reconstructing head models from photographs for individualized 3D-audio processing

Participants: Manuel Asselot, Nicolas Tsingos.

Visual fidelity and interactivity are the main goals in Computer Graphics research, but recently also audio is assuming an important role. Binaural rendering can provide extremely pleasing and realistic three-dimensional sound, but to achieve best results it is necessary either to measure or to estimate individual Head Related Transfer Function (HRTF). This function is strictly related to the individual features of ears and face of the listener. Recent sound scattering simulation techniques can calculate HRTF starting from an accurate 3D model of a human head. Hence, the use of binaural rendering on large scale (i.e. video games, entertainment) could depend on the possibility to produce a sufficiently accurate 3D model of a human head, starting from the smallest possible input.

In this paper we present a completely automatic system, which produces a 3D model of a head starting from simple input data (five photos and some key-points indicated by user). The geometry is generated by extracting information from images and accordingly deforming a 3D dummy to reproduce user head features. The system proves to be fast, automatic, robust and reliable: geometric validation and preliminary assessments show that it can be accurate enough for HRTF calculation (Figure 11).

This work was published at Pacific Graphics 2008 [20], in collaboration with the CNR Pisa (M.Dellepiane and colleagues).

Figure 11. Left: Two results of processed heads and Right: Matching simulated HRTF measurements for a scanned model and a reconstructed model

5.3. Virtual Environments with Sound

5.3.1. A Psychophysical Study of Fixation Behavior in a Computer Game

Participant: Efstathios Stavrakis.

In this work [26], we performed a psychophysical experiment that shows that in a task-oriented context, such as gaming, saliency maps may be weak predictors of fixation behavior of players, when compared to the use of eye-tracking at design time. Prediction of gaze behavior in gaming environments can be useful asset to game designers, enabling them to improve gameplay, selectively increase visual fidelity, and optimize the distribution of computing resources.

We have designed a computer game, in which the task is to control a ball from one part of a maze to reach a designated area. We have setup a worst-case scenario for saliency algorithms by strategically placing distractors with high saliency at areas of the maze that become less relevant to the task. We have recorded and statistically analyzed the fixation behavior (Figure 12(a)) of 40 participants, under different scenarios and compared our findings (Figure 12(c)) to the prediction provided by saliency (Figure 12(b)). We arrived at the conclusion that understanding where game players will potentially focus their attention can be better achieved by eye-tracking analysis at design time, rather than saliency maps alone.

This work was performed in collaboration with the University of Bristol (Veronica Sundstedt and Erik Reinhard), and the Vienna Institute of Technology(Michael Wimmer).

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

Participants: Nicolas Bonneel, George Drettakis.

In [16] 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 a bunny or 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.

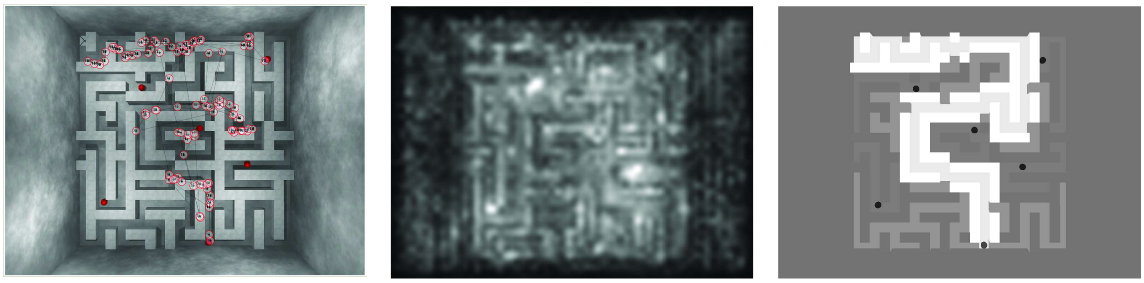

Figure 12. In Figure (a), the fixations for a representative participant of our study are superimposed over a frame from the actual maze. Figure (b) shows a saliency map of the maze and Figure (c) shows our predicted distribution map using data from 10 participants. Brighter pixels in Figures (b) and (c) can be considered to form areas where fixations are more likely to occur.

This work is a collaboration with Isabelle Viaud-Delmon and Clara Suied from CNRS/IRCAM.

This work led to a technical report [30] for citation in another submission, and is currently accepted with minor revisions to the journal ACM Transaction on Applied Perception [16].

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 pre-recorded sounds. This optimization allows computing impact sound per-frame energy more than 200 times faster than previous methods, by considering an exponentially decaying energy function. It also allows the integration of good quality attacks in the previous pipeline, thus handling recorded sounds with high quality impact sounds.

We also exploited a recent result on audio-visual crossmodal perception 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 (accepted with minor revisions), 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 PentaG, thus demonstrating the potential utility of our approach.

This work has been accepted for publication at the I3D 2009 international conference [23], 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 at 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 shown such a crossmodal effect on the semantic, and shown that incorrect spatialization did not influence reaction time.

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

5.3.5. Auditory-visual virtual environments to treat dog phobia

Participants: Isabelle Viaud-Delmon, Nicolas Bonneel, George Drettakis.

We present an experiment aimed at treating dog phobia thanks to audiovisual virtual environments. We developed 3 different virtual environments, shown in Fig.14, in order to immerse participants. We selected a Doberman dog among 9 different dogs (Fig.13) through a preliminary experiment consisting in rating the valence of short animations of dogs. Participants were also selected based on a preliminary questionnaire, keeping those with highest scores, indicating a higher dog phobia.

Figure 13. Bull terrier, Doberman and Boxer, used in the preliminary experiment. Doberman was considered as most negatively valenced.

A pilot study was run to determine the best experimental conditions to treat dog phobia, and to retrieve participants' impression. The pilot consisted in performing specific tasks in each virtual environment. Light condition varied, and participants were immersed thanks to a large screen with passive stereo rendering, and HRTF processing for audio spatialization. We found that some tasks were too difficult, and that the hangar scene (Fig.14) was judged as the most anxiogenic. Light condition and contrast did not influence the pilot since participants focused on their task.

Figure 14. Virtual environments used in the experiment

This work [27] is a collaboration with Isabelle Viaud-Delmon and colleagues, and was presented at ICDVRAT with ArtAbilitation 2008.

5.3.6. Retargetting Example Sounds to Interactive Physics-Driven Animations

Participants: Cécile Picard, Nicolas Tsingos, François Faure.

This work proposes 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. In an off-line analysis task, 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. A technique similar to matching pursuit is used to represent each original recording as a compact series of audio grains. During interactive animations, the grains are triggered individually or in sequence according to parameters reported from the physics engine and/or user-defined procedures. A first application is simply to reduce the size of the original audio assets. Above all, our technique allows the synthesis of non-repetitive sounding events and provides extended authoring capabilities. Results of our method are provided at http://evasion.imag.fr/~Cecile.Picard/SupplementalAES. This work has been accepted for publication at the AES 2009 International Conference [25].

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 agreement funds the PhD thesis of C. Picard to co-develop novel real-time synthesis algorithms for interactive audio rendering.

6.3. Audio Technology Transfer Games

After several audio research publications at REVES, an US patent was filed in 2004 by INRIA to protect two main audio technologies developed within the group. 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 perceptive metrics. The second component, called clustering, aims at grouping audio sources that are spatially close from 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, 3D audio scene reconstruction from recordings, sound propagation and reverberation, modal sound synthesis and contact sounds generation.

After issuing 2 games on the market, Test Drive Unlimited and Alone In The Dark, Eden Game's feedback was very positive and it motivated REVES to develop its audio technologies and transfer them to a start-up company.

Manuel Asselot is in charge of this project at INRIA. In January, he applied to the national contest for the creation of innovating companies organized by the French Ministry of Research, and won a prize to finance a market research and legal expenses. In addition to this financial work, there were 4 main activities this year related to the company creation.

The first one was the transfer of the audio group knowledge from Nicolas Tsingos to Manuel Asselot. Nicolas Tsingos had planned to leave INRIA and work at Dolby in San Francisco in June 2008. As he was the main person who held all the audio knowledge in the group, such a transfer was necessary in order to keep the knowledge as much as possible at INRIA, and to be able to use it for the company creation and its products developments.

The second activity was the technical preparation of audio products. A complete audio solution had to integrate all the audio technologies that had been developed in the group. It would help demonstrating the integrated use of the various techniques, and would be used as a base platform to show what type of products the company is able to deliver to its clients.

The third activity concerned the legal aspects of the company creation. The company can use the patents filed by INRIA as soon as an agreement is negotiated between both parties. Discussions about intellectual property transfer have begun between INRIA and the attorney Manuel Asselot has contracted with. These discussions may still last for several months. There is no clear view at the moment about the exact intellectual property portfolio, as 3 additional patents may be filed in the coming months for the company use.

The fourth activity was related to the market research needed for the validation of the company project. On September 2008, Manuel Asselot contracted a market research with the company Actemis. Since then, a first phase has been completed and interviews with the actors of the targeted industries have been realized. In January 2009, when all the interviews are over, the results of the market research will provide all the necessary information to build a detailed business plan for the company.

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 [27], [30], [16], [15].

7.1.2. Workbench and Immersive space

Participants: David Grelaud, Nicolas Bonneel, George Drettakis.

Workbench 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; D. Grelaud has taken over 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 equipped 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. We have recently upgraded our tracking facility to an ART, vision-based system. This system is a shared resource used by both our group and others. We used it both as a demo platform and an immersive display used for conducting experiments and user testing.

New Immersive Space Under the programme CPER Telius, the Centre de Recherche INRIA Sophia Antipolis - Mediterranean builds an immersive space. David Grelaud is strongly involved in this project from the needs analysis to the preparation of audio-visual demonstrators which will emphasize the team's work. Before starting the public tender, we performed a state-of-the-art review in the large domain of virtual reality (VR). Among other things, we found a high quality sound system based on the binaural approach which is compatible with a CAVE-like environment. This technology allows people to hear near-to-head sources realistically. It will drastically enhance the sense of immersion in the virtual reality platform and it is crucial for us. We also published a large survey to meet researchers needs. This work considerably helped us specify the functional requirements of the project and set up the public tender.

Awaiting the installation of the new immersive space, we developed two VR demonstrators based on OgreVR which work for the moment on the workbench. See list of demonstrations below.

Lastly, we are working on a demonstrator (Fig. 15) based on PentaG (an open-source game) which gathers the last contributions of the team in the European project CROSSMOD. This implementation is also the opportunity to refactor the code and create a software toolbox for fast impact sounds synthesis. Thus, in the future, we should be able to generate contact sounds in other 3D engines.

Figure 15. PentaG demonstrator: the supermarket.

7.2. European Projects

7.2.1. Open FET IST CROSSMOD

Participants: George Drettakis, Monique Varlet, David Grelaud, Manuel Asselot, Nicolas Tsingos, Nicolas Bonneel, Sylvain Lefebvre, Efstathios Stavrakis.

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 has been measured by the efficiency/quality improvements of the new algorithmic solutions developed, and the validation of our hypotheses in realistic VE settings. In the third year of CROSSMOD, REVES had a number of significant research activities. These activities are described in more detail in the appropriate sections of this report.

The published research results concern the work on contact sounds [15], visual-auditory conflicts [27], material recognition [16] and the audio pipeline for a game environment [23]. Other CROSSMOD-related work includes investigation of the effect of sound on material recognition, and the work on massive impact sounds. As project coordinator, we are responsible for the publication of all the internal wiki material and continued updates, and the external website. We also have the overall responsibility for all project Deliverables and in particular the 6-month and annual reports. As coordinator, 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.

7.3. Visiting Researchers

Michiel van de Panne, professor at the University of British Columbia, is visiting the Reves research group on a sabbatical for a year (since August 2008). Efstathios Stavrakis (originally from TU Vienna), was a post-doctoral researcher in the team for a year (until September 2008), in the context of the CROSSMOD project. Marcel Meili, from ETH Zurich, completed his masters internship (March 2008 to September 2008), on Progressive Texture Mapping of Dynamic Objects.

We hosted several visiting researchers this year: Eugene Fiume (University of Toronto, November), Wojciech Jarosz (University of California San Diego, October), Ares Lagae (from FWO Belgium, October and December), Mathieu Lagrange (from McGill University Canada, April 2008).

Durand Begault (from NASA Ames Research Center US), Karlheinz Brandenburg (from IDMT Germany), Davide Rocchesso (from University IUAV Venice) were invited last April as jury members of the habilitation defense of Nicolas Tsingos.

7.4. Bilateral Collaborations

7.4.1. France-China

We collaborate with Microsoft Research Asia (Beijing) on texture synthesis (X. Tong and Yue Dong) resulting in the EGSR paper [21].

7.4.2. France-Germany

We collaborate with the University of Stuttgart (new position of C. Dachsbacher), resulting in the publication on geometry [17].

7.4.3. France-United States of America

We have an active collaboration with Frédo Durand of MIT CSAIL on the project of hair estimation. Nicolas Bonneel visited the MIT for 3 weeks in July.

We also collaborated with Paul Calamia of Rensselaer Polytechnic Institute (reverberation graphs paper [18]) and D. James from Cornell University on the contact sounds paper [15].

7.4.4. France-Italy

In the context of the CROSSMOD project, we have developed a working relationship with CNR Pisa; this resulted in a joint publication at Pacific Graphics with Matteo Dellepiane et al. [20], in the context of CROSSMOD.

8. Dissemination

8.1. Participation in the Community

8.1.1. Program Committees

George Drettakis was co-chair of the programme committee Eurographics 2008, which took place in Crete in April 2008 and a member if the EGSR08 and I3D08/09 program committees; he reviewed papers for several other conferences and journals.

In 2008, Sylvain Lefebvre served on the Eurographics papers committee, on the Eurographics Rendering Symposium papers committee and on the Symposium on Interactive 3D Graphics and Games committee. He was reviewer for the Computer Graphics Forum journal, the Transaction on Visualization and Computer Graphics journal as well as the SIGGRAPH and SIGGRAPH Asia conferences, and the ACM Solid and Physical Modeling symposium.

Nicolas Tsingos was program committee member of EGSR 2008, and also reviewed papers for several journals and conferences.

8.1.2. Invited Talks

G. Drettakis presented a keynote presentation at the International Conference on Multimodal Interfaces (Crete, October 2008). Sylvain Lefebvre gave an invited talk at the Mini-symposium on GPGPU of the Norwegian curves and surfaces conference, as well as a seminar at ISIMA in Clermont-Ferrand.

8.1.3. Thesis Committees

G. Drettakis was a committee member for the Ph.D. thesis of V. Forest (Toulouse). Sylvain Lefebvre was examiner on the thesis committee of Luc Buatois, PhD student in the Alice team of INRIA Nancy - Grand Est.

8.1.4. Community service

G. Drettakis served on the CR2 (junior researcher) hiring committee in 2008 at Sophia-Antipolis.

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, Sylvain Lefebvre and Nicolas Tsingos taught a Computer Graphics course for the IGMMV Master (3 hours each for a total of 9 hours). They also taught a course to the ENS M1 Master (6 hours each for a total of 18 hours). Sylvain Lefebvre gave an introductory Computer Graphics course for students of "Ecole des Ponts" visiting the INRIA Sophia-Antipolis center.

8.2.2. PhD Thesis Completed and Continuing

Cécile Picard was in her 2nd year working on physically-based audio synthesis. Nicolas Bonneel was also in his 2nd year as part of the CROSSMOD project working on audiovisual crossmodal effects. 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

N. Bonneel presented the contact sound paper [15] at SIGGRAPH 2008 (LA, USA, August).

8.3.2. Participation at Conferences and Workshops

G. Drettakis attended Eurographics 2008 (Crete, April) and SIGGRAPH 2008 (L.A. USA, August) and ICMI 2008 (Crete, October). S. Lefebvre attended Eurographics 2008 and EGSR08. Marcio Cabral attended SIGGRAPH 2008. N. Bonneel attended EGSR 2008.

8.4. Demonstrations and Press

8.4.1. Demonstrations

Participants: David Grelaud, George Drettakis.

The first demonstrator for the immersive space was used by the EPI ODYSSEE for the Ph.D defense of Maxime Descoteaux in February. We also implemented the work of Nicolas Bonneel [15] on the workbench. These “showroom” applications were used: for a virtual reality lesson for 32 students of the school Ecole National des Ponts et Chausées in September and December, a conference for the Complex Optical and Imaging Systems Competitivity Cluster POPSud on September the 18th and for a demonstration for the site visit of International Conference on Intelligent Robots and Systems (IROS’08) in September.

The REVES group also presented demos at the NEM Summit (November), for the companies EXKEE and Panasonic (Spring) and for the students of ENS Lyon in November.

8.4.2. Press

Participants: Sylvain Lefebvre, Nicolas Tsingos.

S. Lefebvre gave an interview for the young public magazine Okapi. N. Tsingos participated on the topic on 3D sound presented on the interstice site http://interstices.info/jcms/c_30410/simulation-de-scenes-sonores-spatialisees-complexes, and Inedit presented an article on the INRIA 40 year film which contained a piece on audio.

9. Bibliography

Major publications by the team in recent years

[1] 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.

[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/Vis2002Deussen.pdf.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] N. TSINGOS, I. CARLBOM, G. ELKO, R. KUBLI, T. FUNKHOUSER. *Validating acoustical simulations in the Bell Labs Box*, in "Computer Graphics and Applications, IEEE", vol. 22, n^o 4, Jul/Aug 2002, p. 28-37.

[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. TSINGOS. *Models and Algorithms for Interactive Audio Rendering*, Habilitation a Diriger les Recherches (HDR), Ph. D. Thesis, University of Nice Sophia-Antipolis, 2008, http://www-sop.inria.fr/reves/Basilic/2008/Tsi08.

Articles in International Peer-Reviewed Journal

[15] 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.

[16] 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 (accepted with minor revision)", to appear, 2009.

[17] M. CABRAL, S. LEFEBVRE, D. CARSTEN, G. DRETTAKIS. *Structure-Preserving Reshape for Textured Architectural Scenes*, in "Computer Graphics Forum (Proceedings of the Eurographics conference)", to appear, 2009.

[18] E. STAVRAKIS, N. TSINGOS, P. CALAMIA. *Topological Sound Propagation with Reverberation Graphs*, in "Acta Acustica/Acustica - the Journal of the European Acoustics Association (EAA)", 2008, http://www-sop.inria.fr/reves/Basilic/2008/STC08.

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

International Peer-Reviewed Conference/Proceedings

[20] M. DELLEPIANE, N. PIETRONI, N. TSINGOS, M. ASSELOT, R. SCOPIGNO. *Reconstructing head models from photographs for individualized 3D-audio processing*, in "Computer Graphics Forum (Special Issue - Proc. Pacific Graphics) 27(7)", 2008, http://www-sop.inria.fr/reves/Basilic/2008/DPTAS08.

[21] Y. DONG, S. LEFEBVRE, X. TONG, G. DRETTAKIS. *Lazy Solid Texture Synthesis*, in "Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)", 2008, http://www-sop.inria.fr/reves/Basilic/2008/DLTD08.

[22] C. EISENACHER, S. LEFEBVRE, M. STAMMINGER. *Texture Synthesis from Photographs*, in "Proceedings of the Eurographics conference", 2008, http://www-sop.inria.fr/reves/Basilic/2008/ELS08.

[23] 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", To appear, Feb 2009.

[24] C. PICARD, N. TSINGOS, F. FAURE. *Audio Texture Synthesis for Complex Contact Interactions*, in "VRIPHYS 08", 2008, http://www-sop.inria.fr/reves/Basilic/2008/PTF08.

[25] C. PICARD, N. TSINGOS, F. FAURE. *Retargetting Example Sounds to Interactive Physics-Driven Animations*, in "AES 35th International Conference - Audio for Games", to appear, 2009.

[26] V. SUNDSTEDT, E. STAVRAKIS, M. WIMMER, E. REINHARD. *A Psychophysical Study of Fixation Behavior in a Computer Game*, in "Proceedings of the 5th Symposium on Applied Perception in Graphics and Visualization", S. CREEM-REGEHR, K. MYSZKOWSKI (editors), ACM, 2008, http://www-sop.inria.fr/reves/Basilic/2008/SSWR08.

[27] I. VIAUD-DELMON, F. ZNAÏDI, N. BONNEEL, C. SUIED, O. WARUSFEL, K.-V. N'GUYEN, G. DRETTAKIS. *Auditory-visual virtual environments to treat dog phobia*, in "Proceedings 7th ICDVRAT with ArtAbilitation", September 2008, http://www-sop.inria.fr/reves/Basilic/2008/VZBSWND08.

Scientific Books (or Scientific Book chapters)

[28] C. DACHSBACHER, S. LEFEBVRE. *Efficient and Practical TileTrees*, Shader X6, (to appear), Charles River Media, 2008, http://www-sop.inria.fr/reves/Basilic/2008/DL08.

[29] S. LEFEBVRE. *Filtered Tilemaps*, Shader X6, (to appear), Charles River Media, 2008, http://www-sop.inria.fr/reves/Basilic/2008/Lef08.

Other Publications

[30] I. VIAUD-DELMON, F. ZNAÏDI, N. BONNEEL, D. DOUKHAN, C. SUIED, O. WARUSFEL, K.-V. N'GUYEN. *Auditory-Visual Virtual Environments to treat Dog Phobia*, International Conference Series On Disability, Virtual Reality And Associated Technologies, sep 2008.

References in notes

[31] 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.

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

[33] 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.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

[38] 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.

[39] 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.

[40] 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.

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

[42] 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.

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

[44] 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.

[45] 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.