INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team REVES

Rendering and Virtual Environments with Sound

Sophia Antipolis

THEME COG

Activity Report

2005

Table of contents

1. Team

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

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

Postdoctoral Fellow
Guillaume Lemaitre [Post-doctoral fellowship/RNTL OPERA]

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

Ph.D. students
Florent Duguet [AMX Scholarship]
Marie-Claude Frasson [FQRNT Quebec Fellowship]
Alex Reche [CIFRE Scholarship with the CSTB]
Emmanuel Gallo [CIFRE Scholarchip with the CSTB]

Research scientist (external collaborators)
Michael Ashikhmin [Assistant Professor, Stonybrook University]
Pau Estalella [Ph.D. student University of Girona]
Ignacio Martin [Researcher University of Girona]
Simon Premoze [Postdoctoral fellow, Columbia]
Maria Roussou [University College London and ICCS, Athens]

Administrative assistant
Marie-Line Ramfos [Secretary (SAR) Inria]

Interns
Julien Etienne [DEA, MPRI]

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" [41], 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 [29], but also simpler methods such as accessibility maps [38]. 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.

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [28], 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 [39] 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 [33] [32], [31] 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 [30]. 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 (see also the paragraph 4.2 on applications). 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 [34] [39]), 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 [35] 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 [40].

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 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

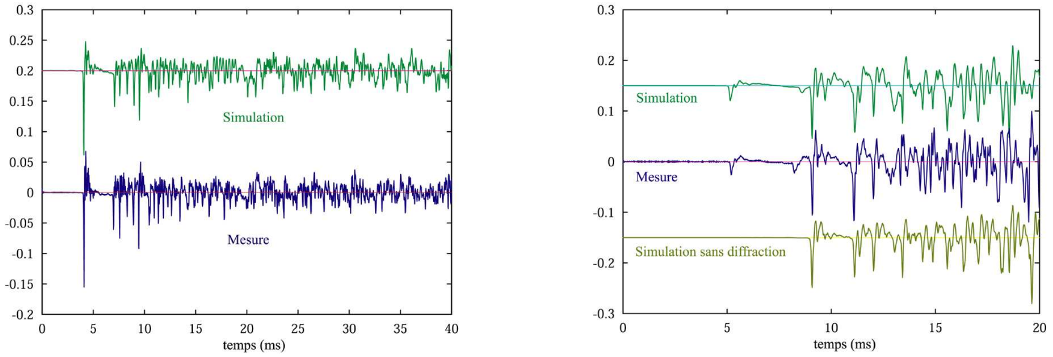

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

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.

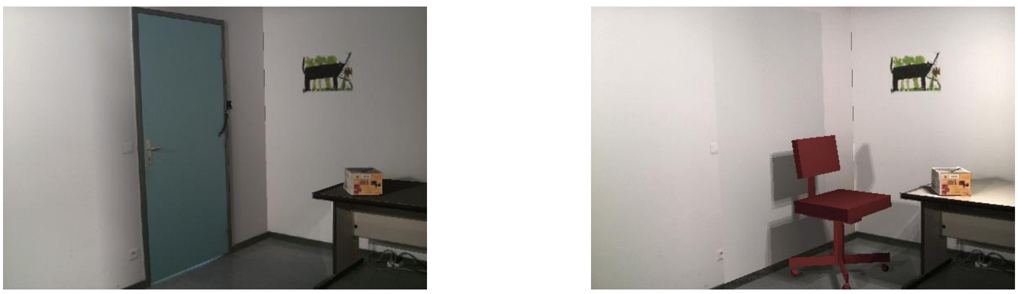

Figure 3. (a) Original conditions (b) The door has been removed virtually, and a virtual object and light have been added (method of [37])

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 [36], [37]. 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 [36], [37] to high quality approaches [42], 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 is entirely possible that the lack of usable interfaces is part of the reason that 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. Application Domains

4.1. Virtual Heritage

Keywords: *Virtual heritage, virtual archeology*.

Virtual heritage is a recent area which has seen spectacular growth over the past few years. Archeology and heritage exhibits are natural application areas for virtual environments and computer graphics and sound, since our research can reconstruct artefacts, monuments or buildings of lost civilisations, both in images and sound.

We are interested both in "general public" presentations of cultural heritage artefacts and sites, and in the actual archeological research of the experts.

Our cultural heritage work has concentrated around three main collaborations. The first was with the Foundation of the Hellenic World (FHW) http://www.fhw.gr/fhw/en/home/, the second was in the context of past project the ARC ARCHEOS (http://www.inria.fr/reves/Archeos and the third is with the Virtual Reality group of EDF and a programme funded by the "Fondation EDF" (see also Section 7.4).

Our collaboration with the FHW was in the context of the EU IST project CREATE which ran from March 2002 to March 2005 (however, our contribution was completed at the end of 2004). In CREATE the combination of educational theories and field expertise permitted a novel approach to the usage of Virtual Environments for both experts and novices. The first application of the project was in cultural heritage, and REVES contributed significantly to the development of higher quality virtual environments based on real world sources as well as virtual environment enhancements (view-dependent texturing, efficient vegetation rendering, lighting, shadows and spatialised sound). Work took place both with archeologists who are interested in virtual reconstruction of monuments and with children in the context of the educational mission of the FHW.

In the context of the past project ARCHEOS (2001-2003), we worked on the Agora site of ancient Argos, and in particular on the Tholos monument. We contacted and worked with two archeologists who have conflicting hypotheses on the utility of the monument, and reconstructed the two different possibilities in a virtual world. In addition, we have investigated the relative effectiveness or utility of realistic vs. non-photo realistic (or expressive) renderings for archeological reconstruction. Although this was very experimental work, we have had very positive feedback from the archeologists.

Finally, we are participating in a cultural heritage project with the VR department of EDF (see also Section 7.4). The goal here is the reconstruction of a complex monument at Delphes (the Dancer column and the Omphalos), and to test a number of hypotheses (structural, aesthetic etc.). Our contribution is the development of efficient and high-quality reconstruction and rendering techniques, in the context of the thesis of F. Duguet, using a point-based approach.

We are also interested in the use of sound in a cultural heritage context. The acoustic properties of many sites could be particularly interesting to study, for example the famous acoustics of ancient Hellenic theatres.

4.2. Urban planning, Architecture, Evaluation and Training

Keywords: *architecture, evaluation, training, urban planning*.

Urban planning and architecture are very appropriate application domains for computer graphics and virtual environments, since they often deal with future projects requiring visual or auditive representations. In addition, evaluation, design review and training can also greatly benefit from digital or virtual mock-ups. Our expertise in rendering and its application to VE's can greatly benefit the process.

Our work in this domain has been with two main partners, the CSTB (the French National Scientific and Technical Center for Construction http://www.cstb.fr), and in particular their Virtual Reality department at Sophia-Antipolis, and the architectural company AURA in Marseille, as part of our work on the CREATE project. Our collaboration with the CSTB has concentrated on three axes: urban planning in the context of our common IST project CREATE and sound simulation in the context of the OPERA RNTL project (see also Section 8.2.2).

Training and evaluation are domains in which VE's can be applied naturally. The use of coherent sound and image renderings in VE's can be particularly important for training in risk-critical environments (for example nuclear power plants). We have several contacts with companies working in these domains (notably EDF).

4.3. Computer Games

Keywords: *computer games*.

Computer Games have been the driving force in the development of rendering and computer graphics research, especially in terms of low-cost hardware over the past few years. Interactive rendering for ever-more complex scenes, both in terms of geometry and lighting effects is clearly of great interest for games companies.

Integration of sound spatialisation, either geometric or statistical is clearly promising, and is currently almost entirely missing in exisiting games. We believe that the importance of these techniques and of their maturity for technology transfer is significant; we currently working closely with the company EdenGames in this domain (see also Section 7.1).

A more long-term goal for this application domain is the use of virtual environments in low-cost immersive or semi-immersive contexts. Their augmentation with both synthetic images and sound should have great potential. As a first step, such applications would be limited to more "theme-park" style environments. However, the emergence of low-cost "wall-projection" stereo and tracking systems could result in the development of installations that have cost equivalent to that of a home-cinema setup today, making them potentially viable in a mass-market context.

4.4. Audiovisual and Post-production

Keywords: *post-production*.

Although our emphasis is on interactive applications, our high-quality rendering research, both in sound and images could be of interest in post-production or the film industry. We are also interested in combined interactive/offline approaches, such as the previsualisation tools in certain modelling/animation packages.

Integrating vision-based match-moving techniques for placement of sound tracks could also be useful, but care must be taken to preserve artistic control for sound engineers and technical directors.

5. Software

5.1. AURELI: AUdio REndering LIbrary

Participants: Nicolas Tsingos, Emmanuel Gallo.

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 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.

5.2. Matlab SPL (Sound Description Library) toolbox and ".sig" audio description files

Participants: Emmanuel Gallo, Guillaume Lemaitre.

Fast audio processing uses compact sound descriptions attached to each file. We have developed two libraries of functions (in C and in Matlab) which aims at computing such a description for some audio file. The Matlab SPL (Sound Processing Library) toolbox has been defined in an INRIA technical report which will be soon published. A accompanying report [25] mainly reviews the concept of sound description and proposes a state-of-the-art of the different kinds of sound descriptions already used. It also reminds the theoretical and physical definitions of energy and power (often confused), and defines our own normalization conventions. These definitions are implemented in a Matlab library (SPL) which aims at computing a sound description in a specific file (.sig). Finally the report describes the Matlab SPL functions and the .sig file structure.

5.3. Ogre3D-based graphics testbed

Participants: David Geldreich, George Drettakis, Frank Firsching.

In the context of his DREAM "mission" David Geldreich has set-up a customized development environment based on the Ogre3D graphics engine. This setup is currently being used in the perceptual rendering research project, and Frank Firsching is integrating an audio-library interface between OpenAL (an open-source audio library standard) and Ogre3D. We expect this setup to be the software basis for new interns in the group.

6. New Results

6.1. Plausible Image Rendering

6.1.1. Accurate Interactive Specular Reflections on Curved Objects

Participants: Pau Estalella, Ignacio Martin, George Drettakis.

Rendering interactive reflection from curved specular objects has been a long standing problem in computer graphics. In this work, which is a collaboration with the University of Girona in Spain, and the GEOMETRICA project of INRIA Sophia-Antipolis, a new solution was proposed. The difficulty has always been finding the point of reflection for a given object and view in the traditional stream-processor graphics pipeline, as used in the Graphics Processing Unit (GPU).

Our new approach creates virtual reflected objects which are blended into the scene. We use a property of the reflection geometry which allows us to efficiently and accurately find the point of reflection for every reflected vertex, using only reflector geometry and normal information. This reflector information is stored in a pair of appropriate cubemaps, thus making it available during rendering. The implementation developed achieves interactive rates on reasonably-sized scenes. In addition, we introduce an interpolation method to control the accuracy of our solution depending on the required frame rate. This work was published at the VMV conference [18].

6.1.2. Perceptual rendering

Participants: Julien Etienne, Nicolas Tsingos, George Drettakis.

The goal of the M.Sc. project of Julien Etienne, was to develop a percetually-based strategy for level-of-detail, or selective, rendering. To achieve this goal we studied the two main Visual Difference Predictors, notably that proposed by Daly and that of Sarnoff/Lubin. We chose the Sarnoff model for its computational efficiency; this model was implemented in our platform. We are currently investigating new algorithms which can take advantage of the visual difference predictor to improve real-time level of detail management, and in particular for multiple representations (points, polygons etc.).

6.1.3. Treating very large point-based data sets

Participants: Florent Duguet, George Drettakis.

In this research we proposed a unified solution for the creation of levels of detail on very large input data, which are typically the result of detailed laser-scans. We build a hierarchical signed distance function in an octree around the data and use this hierarchy to generate a continuum of levels of detail. Our distance function

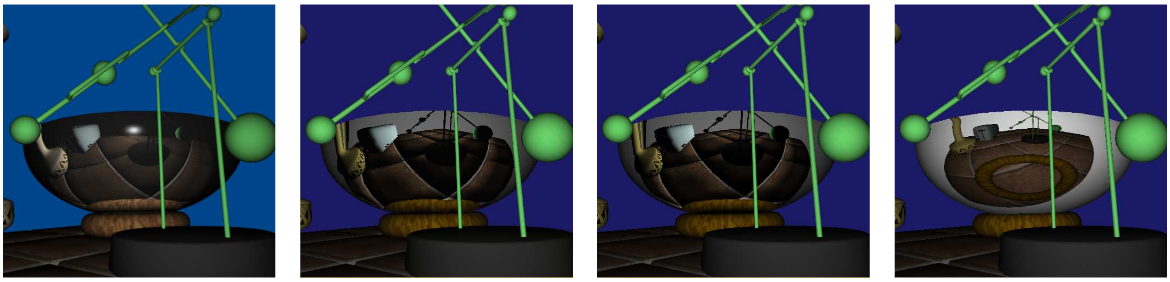

Figure 4. The Kitchen scene. From left to right. The reference ray-traced image, the scene rendered using our approach searching for all the vertices, the scene rendered using our approach, searching for only 10% of the vertices and approximating the others, the scene using dynamic environment maps (i.e., recomputed at each frame). Our approach for full search as well as with the approximation is clearly much closer to the ray-traced solution than dynamic environment maps.

construction, based on the Gradient Vector Flow and the Poisson equation, builds on multigrid resolution algorithms. Using an appropriate interpolation scheme within the octree we obtain a continuous hierarchical distance function, which allows us to define a continuum of levels of detail for huge geometries. During this process, holes and undersampling issues in the input data are automatically corrected. We developed three applications of our hierarchy: a novel hierarchical deformable model scheme that can automatically reconstruct closed Eulerian meshes of up to a million faces in a few minutes (see Fig. 5), an alternate distance-driven contouring approach, and raytracing of huge data models (see Fig. 9). This work is presented in detail in [24] and also in [27], and is part of Florent Duguet's Ph.D. thesis [14].

2.6 k (0.7 s) 11 k (1.3 s) 43 k (3.4 s) 176 k (19 s) 704 k (2:10) 2,500 k (27m)

Figure 5. Levels of detail on the entire St Matthew model, generated using the proposed hierarchical continuous distance, and reconstructed using the proposed hierarchical deformable mesh reconstruction approach. Input model is 370,000,000 points. Below each image, the number of triangles of the rendered mesh, and reconstruction time in parentheses (after an initial construction time of 1h45 for the distance function on an octree). Shading is done using per vertex ambient occlusion.

Point-Based Surfaces can be directly generated by 3D scanners and avoid the generation and storage of an explicit topology for a sampled geometry, which saves time and storage space for very dense and large objects, such as scanned statues and other archaeological artefacts. In a collaborative project with the IPARLA INRIA Project in Bordeaux we propose a fast processing pipeline of large point-based surfaces for real-time, appearance preserving polygonal rendering. The goal is to reduce the time needed between a point set made of hundreds of millions of samples and a high resolution visualization using modern graphics hardware, tuned for normal mapping of polygons.

The approach starts by an out-of-core generation of a coarse local triangulation of the original model which is then enriched by applying a set of maps capturing high frequency features of the original data set. Normal maps are used in this step; other attributes such as color or position (displacement map) can also be treated. Sampling issues of the maps are addressed using an efficient diffusion algorithm in 2D. One of the main advantages is to express most of the the features present in the original large point clouds as textures in the ample texture memory on GPUs, using only a lazy local parameterization. Direct applications are: interactive preview at high screen resolution of very detailed scanned objects such as scanned statues, inclusion of large point clouds in usual polygonal 3D engines and 3D databases browsing. This work was presented at the Symposium on VR, Archeology and Cultural Heritage (VAST 05) [17].

6.1.4. Automatic generation of consistent shadows for Augmented Reality

Participant: Alex Reche.

In the context of mixed reality, it is difficult to simulate shadow interaction between real and virtual objects when only an approximate geometry of the real scene and the light source is known.

Alex Reche participated in a research project at the University College London, where a solution was developed a for realtime rendering to simulate colour-consistent virtual shadows in a real scene. The rendering consists of a three-step mechanism: shadow detection, shadow protection and shadow generation. In the shadow detection step, the shadows due to real objects are automatically identified using the texture information and an initial estimate of the shadow region. In the next step, a protection mask is created to prevent further rendering in those shadow regions. Finally, the virtual shadows are generated using shadow volumes and a pre-defined scaling factor that adapts the intensity of the virtual shadows to the real shadow. The procedure detects and generates shadows in real time, consistent with those already present in the scene and offers an automatic and real-time solution for common illumination, suitable for augmented reality. This work was presented at the Graphics Interface conference [20].

6.1.5. Fast Ambient Occlusion for Trees and Grass

Participants: Michael Ashikhmin, Simon Premoze, George Drettakis.

We have developed a novel approach to approximate ambient occlusion, based on a simple statistical approximation of visibility for trees and grass. The approximation is based on two criteria, one which is linked to how far we are from the outermost part of the tree, and the other which takes into account average occlusion. The approximation works well for uniformly distributed leaves, which satisfy our statistical hypotheses, and delivers a plausible approximation for other, more realistic cases. This work will be published in 2006.

6.1.6. Recovering structure and aspect detail from example

Participants: Marie-Claude Frasson, George Drettakis.

This project aims to recover structure and aspect information from the analysis of examples (photographs of structured patterns like walls, animal skins, etc.). The recovery of the structure and the aspect are decoupled since synthesizing the two in the same process has never led to good results in the past. With our technique, the geometry of the structure is analyzed by fitting elliptic elements to all the regions. A new structure is synthesized by sampling new elements and populating the surface to cover using a modified version of the Lloyd's algorithm, possibly taking spatial constraints into account. The region borders are adjusted in an optimization stage leading to an editable 2D mesh that can then be applied to 3D models. To recover the aspect (interior and border textures) of the regions in order to generate procedurally completely new textured structures, we have developed a novel approach based on the analysis of the existing textures and texture

Figure 6. Left: original image. Center and right: two synthetically generated structures with synthesized aspect (texture), for both the interior of the regions and the borders.

generation techniques. Results can be seen in Fig. 6. This is the main topic of M-C. Frasson's Ph.D. thesis [15].

6.2. Plausible Audio Rendering

6.2.1. Prioritizing audio signals for selective processing

Participants: Emmanuel Gallo, Guillaume Lemaitre, Nicolas Tsingos.

This research studies various priority metrics that can be used to progressively select sub-parts of a number of audio signals for real-time processing. In particular, five level-related metrics were examined: RMS level, A-weighted level, Zwicker and Moore loudness models and a masking threshold-based model. We conducted both an objective and a pilot subjective evaluation study aimed at evaluating which metric would perform best at reconstructing mixtures of various types (speech, ambient and music) using only a budget amount of original audio data. Our results suggest that A-weighting performs the worst while results obtained with loudness metrics appear to depend on the type of signals. RMS level offers a good compromise for all cases. Finally, our results show that significant sub-parts of the original audio data can be omitted in most cases, without noticeable degradation in the generated mixtures, which validates the usability of our selective processing approach for real-time applications (Fig. 7). This work was presented at the ICAD international conference [19].

6.2.2. Scalable Perceptual Mixing and Filtering of Audio Signals using an Augmented Spectral Representation

Participant: Nicolas Tsingos.

In this work, we propose a novel perceptually-based and scalable approach for efficiently filtering and mixing a large number of audio signals. Key to its efficiency is a pre-computed Fourier frequency-domain representation augmented with additional descriptors. The descriptors can be used during the real-time processing to estimate which signals are not going to contribute to the final mixture. Besides, we also proposed an importance sampling strategy allowing to tune the processing load relative to the quality of the output. We applied our approach to a variety of applications including equalization and mixing, reverberation processing and spatialization. It can also be used to optimize audio data streaming or decompression. By reducing the number of operations and limiting bus traffic, our approach yields a 3 to 15-fold improvement in overall processing rate compared to brute-force techniques, with minimal degradation of the output. Such techniques are also useful for bandwidth allocation and management when streaming audio data over the network (see Fig. 8). France Telecom, a partner of the RNTL OPERA project, is currently transfering these techniques into their voice over IP platform, *comIP*.

This work was published at the DAFx05 conference [22].

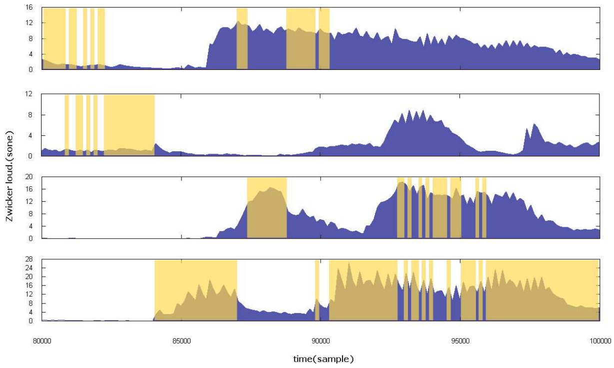

Figure 7. Four speech signals prioritized according to some loudness metric. The highlighted yellow frames are the most loud ones. Only these frames are effectively rendered at each time slot. The overall perceived impression is that the four signals are present and continuous.

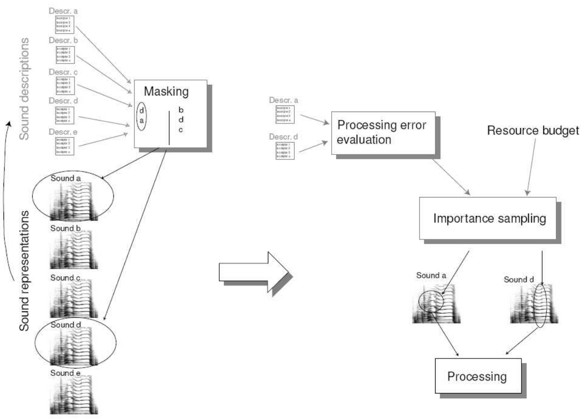

Figure 8. Overview of our scalable perceptual pipeline. All input signal frames at time t are first sorted according to their energy content and a masking estimation is performed. Audible frames are then sampled according to an importance metric so that only a subpart of their pre-computed short-Time Fourier Transform coefficients are processed to produce the output frame at t.

6.2.3. Perceptual optimization of audio rendering

Participants: Nicolas Tsingos, Emmanuel Gallo, Guillaume Lemaitre.

The RNTL project OPERA aims at building solutions to render complex auditory scenes. *Complex scenes* means here that a huge amount of sound sources have to be spatialized, what most of computing systems fail to succeed in. The project is based on the processing pipeline presented in the Siggraph paper of Tsingos et al. in 2004. Ongoing work seeks to improve this pipeline in two main directions: segmentation of the scene into clusters, and computation of a representative signal for a cluster. To group sound sources into clusters, a metric is studied, which takes into account not only geometrical criteria, but also perceptive cues such as similarities, contrast, etc. To compute a representative for the clusters, two kinds of techniques are investigated: the first one seeks to downprocess the mixing of the sound sources according to a budget number of arithmetic operations (as presented in the previous sections), whereas the second approach consists in matching the cluster with sound textures stored in a precomputed library and possibly derived from field recordings.

6.3. Virtual Environments with Sound

6.3.1. VR and Urban Planning Applications

Participants: George Drettakis, Maria Roussou.

As part of the final stage of the CREATE project we performed a detailed analysis of the usage of our VR system with architects and in the context of the real-world urban planning project for the Nice Tramway construction. This work is described in [23] and is currently submitted for publication.

Virtual Reality applications do not have very wild-spread use in fields such as urban planning and education. We have collected our observations from the use of high-end projection-based VR in different real-world settings, with practitioners but also novice users that do not normally use VR in their everyday practice, in a paper presented at the HCI International conference [21]. We discussed two applications for two different content domains (education and urban planning) and presented case studies of actual experiences with professionals and students who used these as part of their work or during their museum visit. Emphasis is given on usability issues and evaluation of effectiveness, as well as on our thoughts on the efficacy of the long term deployment of VR under realistic usage conditions, especially when the technology becomes mundane and the content takes precedence over the display medium. An overall assessment of our experience was presented, on issues relating to usability and user satisfaction with VR in real-world contexts.

The experience we had in our EU IST project "CREATE" showed us that interaction between research and real-world projects can be very beneficial. We presented our views on this topic in a paper presented at the workshop organised by the "Intuition" Network of Excellence in VR [26]. In this paper we discuss how the synergy between real-world VR applications and Computer Graphics and Sound research can be beneficial both to the advancement of research and to the actual application. To illustrate this view we present our experience in the context of the CREATE EU/IST project, which involved the close collaboration between a graphics/sound research group, VE researchers and a real-world urban planning project, involving the construction of the new Tramway in Nice. We believe that the results of this collaboration show the mutual benefit achieved and thus the significant potential of such an approach, even at the stage of exploratory research.

7. Contracts and Grants with Industry

7.1. Eden Games

We have established a strong collaboration with Eden Games, an ATARI subsidiary, which is a games studio in Lyon. In particular, we have signed a technology transfer contract with Eden Games, concerning our audio technology for clustering and masking. This technology relates to the patents described in the Section 7.6.

The licensing agreement includes the use of this technology on the next-generation (XBOX 360 etc.) game consoles, in the context of two new games, scheduled for release in 2006.

This collaboration has been particularly fruitful and productive, since the integration of our technology in a real world product has given us important new insights and allowed us to improve our algorithms. We are pursuing this collaboration with common internships in the context of the Masters programme, with the prospect of co-financed (CIFRE) Ph.D. theses.

7.2. CSTB

Participants: George Drettakis, Nicolas Tsingos, Alex Reche, Emmanuel Gallo.

The thesis fellowship of Alex Reche is a CIFRE contract which is an industry-research collaboration. In this context, Alex Reche spends a significant amount of time at the CSTB and strives to apply his work directly on urban planning and architectural applications. Much of this work was performed in the context of the CREATE project.

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.

7.3. 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, and especially in the context of the CREATE project.

7.4. EDF Foundation

Participants: Florent Duguet, George Drettakis.

In the context of cultural heritage preservation and studies, EDF Foundation has funded the 3D scanning of the Omphalos column in Delphi, Greece. This column, also known as the "Dancers column" is stored in several pieces in crates in the Delphi museum. The goal of the project is to scan the different parts of the column, to test archeological hypotheses using digital data acquired with the computer, including positioning of the stones, comparison of statues etc. The project includes partners from the Louvre Musuem (J-L. Martinez, the curator who has studied this monument extensively, INSIGHT, a not-for-profit company specialising in virtual heritage and scanning and archeologists from the Ecole Normale Supérieure, Paris).

This collaboration continued this year, with the availability of the data for use in our research. The data was used notably for the following publications: [24], [27], [17]. An example of the quality of the images obtained is shown in Fig. 9.

7.5. Collaboration with ENST

Florent Duguet was co-directed by Francis Schmitt at the ENST Paris (http://www.enst.fr). During this period, many scientific collaborations and exchanges took place, notably in the context of the publications with Carlos Hernandez-Esteban [24], [27].

7.6. Patent Application

Participants: Nicolas Tsingos, Emmanuel Gallo, George Drettakis.

We have filed a French and US patent on clustering and masking algorithms for 3D sound "Dispositif et méthode perfectionnés de spatialisation du son" (Advanced System and Method for Sound Spatialisation). French patent FR n°03 13 875 (11/26/2003) and US patent N°10/748,125 (12/31/2003) which are currently pending.

Figure 9. The dancers model: 49,500,000 input points level 11 of the octree, and raytraced at image resolution 1670x3400. Rendering this image required about 4 hours. (Two light sources, and maximal reflexion depth of 2.)

8. Other Grants and Activities

8.1. Regional/Local Projects

8.1.1. Collaboration with CSTB Sophia-Antipolis

Participants: George Drettakis, Nicolas Tsingos, Emmanuel Gallo, Alex Reche.

We collaborate with CSTB in the context of the European IST project CREATE and more recently within the context of the RNTL program OPERA. Two Ph.D. students are partly funded by the CSTB in the context of the CIFRE fellowship program.

8.1.2. 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 setup, integrated and continues support for the system.

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. D. Geldreich also installed the commercial library CAVElib, used in the context of european IST project CREATE. Several members of the group adapted their applications so they can run on the platform.

8.2. National Projects

8.2.1. ACI MD SHOW

Participants: George Drettakis, Florent Duguet.

A national project, coordinated by ARTIS in Grenoble, started in autumn 2003. Our paticipation in this project is on point-based or alternative rendering of very large data sets, such as scans of statues or alternative representations of trees presented in 2004. The other participants of this project are IPARLA in Bordeaux and ISA in Nancy. This year, we developed an alternative visualization approach for large point-based surfaces [17].

8.2.2. RNTL project OPERA: PErceptual Optimizations for Audio Rendering

Participants: Nicolas Tsingos, Guillaume Lemaitre, Emmanuel Gallo.

REVES is coordinator of the RNTL project OPERA, which aims at further studying how audio rendering can be optimized using perceptual knowledge with two applications in mind: telecommunication ("chat rooms", MMOGs) and virtual reality (e.g,. urban planning). In this context, REVES collaborates with IRCAM, France Telecom R&D, LIMSI, CSTB and the company VIRTOOLS. The project started early 2004. Current results and orientations are described in section 6.2.3.

8.3. Visiting Researchers

This year we invited the following visitors: Prof. Pat Hanrahan of Stanford University (February), Isabelle Viaud-Delmon of the CNRS in March, Prof. Eugene Fiume of the University of Toronto (May), Marcus Magnor and Christian Linz of the MPI-Saarbrucken in November, and Carsten Dachsbacher also in November.

Matteo Dellepiane of the CNR Pisa spent 3 months starting in June in the context of a CNR-financed exchange programme.

Frank Firsching of the University of Erlangen started a 3-month visit in December 2005, working on audio and crossmoal audiovisual algorithms.

8.4. European Projects

8.4.1. CROSSMOD

Participants: George Drettakis, David Geldreich, Nicolas Tsingos.

A new EU-financed IST STREPS/FET project, named CROSSMOD, started December 1st. REVES is the coordinator for this 7-partner project. The goal of CROSSMOD is to take advantage of the cross-sensory interaction between images and sound, to improve selective rendering and to provide a toolbox of attention-guidance techniques for audio-visual virtual environments. To achieve this goal, a psychophysical experiment/user study approach will be undertaken, which will guide the development of the algorithmic components.

The other partners of the project are the CNRS Laboratory (UMR 7593) "Vulnerability, Adaptation et Psychopathology", who are specialists in neurosciences and VR, the Unviersity of Bristol, who have performed significant research in perceptually-based rendering, IRCAM, who are the world specialists in psychoacoustics and computer music, the Technical University of Vienna, the University of Erlangen and the CNR in Pisa.

The kickoff meeting of the project took place in Sophia-Antipolis on Dec. 5th and 6th. The initial phases of the project involve a very thorough literature search and the definition of user-studies/experiments which will determine which crossmodal effects are useful to improve audio/visual rendering.

8.5. Bilateral Collaborations

8.5.1. 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 [18].

8.5.2. France-Greece (Hellas)

Our work in common with Maria Roussou of the ICCS has continued with the presentation of two publications [26], [21] on VR and applications.

8.5.3. France-Germany

We have two ongoing collaborations with German laboratories. The first involves the group of Marcus Magnor at the Max-Plank-Institut for Informatics in Saarbruecken on improving our tree rendering and reconstruction algorithms.

The second collaboration involves the University of Erlangen and the graphics group of Marc Stamminger, in the context of the extended visit of the Ph.D. student Frank Firsching.

8.5.4. France-United States of America

We currently have active research contacts with Simon Premoze of Columbia University (NY) and Prof. Ashikhmin at Sunnybrook (NY). The work on the ambient occlusion approximation has resulted in a publication to appear in 2006 (see also Section 6.1.5).

9. Dissemination

9.1. Participation in the Community

9.1.1. Programme Committees

G. Drettakis was a member of the programme committees of ACM SIGGRAPH 2005, the Eurographics Symposium on Point-Based Graphics (2005), the ACM Siggraph Symposium on Interactive 3D Graphics and Games 05, VRST 2005 and the Eurographics Symposium on Rendering 2005, and reviewed papers for a number of journals and conferences. In 2005, Nicolas Tsingos reviewed papers for ACM Transactions on Graphics, ACM SIGGRAPH Conference, Eurographics Symposium on Rendering and IEEE VR conference. He also took part to ACM Multimedia Art Program Comittee.

9.1.2. Thesis Committees

G. Drettakis was an external examiner for the theses of Gaël Guennebaud at the University of Toulouse, and Gabriel Fournier at the University of Lyon.

9.1.3. 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 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.

9.1.4. EU FP7 Working Group and INRIA Visit on Virtual and Augmented Reality

In June, George Drettakis participated in a two-day workshop in Brussels as part of the preparation for the EU FP7 IST programme, with an estimated budget of 44 billion Euros. As a side-effect of this participation, he also organised, in collaboration with B. Raffin at the INRIA Rhône-Alpes, a one-day visit of INRIA activities in the domain of Virtual and Augmented reality. This visit is to take place in Grenoble on Dec. 14th, and involves three officers of the European Commission in the IST strategy department, who will see demonstrations and presentations of 9 INRIA projects nationwide, from all 6 research units.

9.1.4.1. 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

9.2. Teaching

9.2.1. University teaching

George Drettakis was responsible for the Computer Graphics course at ISIA (Ecole des mines) in March 2005 (15 hours), with the participation of Nicolas Tsingos (10 hours). Together with Nicolas Tsingos, they organized and taught the Image and Sound Synthesis course at the newly formed Masters program at the Univeristy of Nice. http://www-sop.inria.fr/reves/Cours/Master2004 This is a 15 hour, 5-session course, with a programming project and exam. Guillaume Lemaitre gave a 5 hour course in June 2005 and another 5 hour course in November 2005 at the Ecole Nationale Supérieure d'Ingénieur du Mans. He taught an introduction to psychoacoustics and sound design.

9.2.2. PhD Thesis Completed and Continuing

A major part of the first semester was spent on the completion of three Ph.D. students studying at REVES. In particular, Alex Reche completed his thesis [16] in April, Marie-Claude Frasson completed her thesis [15]

(UNSA) in June and Florent Duguet completed his thesis [14] in June (UNSA). F. Duguet's thesis was co-directed by Francis Schmitt, at the ENST, Paris, where he spent the last 18 months.

Emmanuel Gallo started his PhD early April 2004 and is currently continuing in his second year.

9.3. Participation at conferences

9.3.1. Presentations at Conferences

Emmanuel Gallo presented the paper [19] at ICAD'05 in Limerick, Ireland. Nicolas Tsingos presented the paper [22] at DAFX'05 in Madrid, Spain and the paper [26] at the Intuition workshop at Senlis, France in November.

9.3.2. Participation at Conferences and Workshops

Other than the paper presentations above, G. Drettakis, N. Tsingos, E. Gallo participated at the Imagina 2005 conference. G. Drettakis and M-C. Frasson participated at the EGSR'05 (Konstanz, CH), and G. Drettakis was at the ACM SIGGRAPH conference in Los Angeles 2005, where he chaired a session as a member of the program committee.

9.4. Demonstrations and Press

9.4.1. Demonstrations

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

The workbench is frequently used for demonstrations and press operations. We typically present the demonstration of the work developed in the context of the CREATE project.

This year we presented demos to high-school students in January, representatives of the Technopole Cannes in September and students of the Ecole Normale Supérieure at Cachan in October.

9.4.2. InTech Seminar

REVES was the organiser for an InTech seminar on Virtual Reality in April. The invited speakers were from Eden Games (a game company) and the SNCF Research and Development department, who use VR for training and other applications. On the academic side, Bruno Levy of INRIA Lorraine and N. Tsingos of REVES gave presentations. About 30 people registered for the presentations.

9.4.3. Traité de la RV

Participant: Nicolas Tsingos.

The "Traité de la RV" (3rd edition) is a collective resource book about VR-related issues. More than 50 authors collaborated to produce a 4-book, 2000 pages document covering all aspects of VR research and systems. We contributed three chapters on audio aspects in three of the four books corresponding to state-of-the-art reports about audio perception and hearing, spatial audio rendering systems and models for real-time audio simulation. An english translation of this work is also scheduled. For an overview of the contents of the previous 2nd edition please browse:

http://www.ensmp.fr/Fr/Services/PressesENSMP/Collections/ScMathInfor/Intro/col-mathinfo.html

Nicolas Tsingos co-authored three chapters of this book [11], [12], [13]. These chapters comprise a more than 80-page introduction on 3D audio and audio rendering in particular.

9.4.4. Press

Our audio rendering technology has been featured in a number of web sites including ITR Manager, Vunet.fr, Animasoft, etc. It also appeared in the French business newspaper "Le Nouvel Economiste" and on the French national information radio "France Info".

The January issue (no 47) of the nation-wide INRIA publication INEDIT contained an article on our work on the CREATE EU project and the application to urban planning.

An article on the same topic was also presented in the "Bleu Inria" email newsletter (issue 12, January 2005) targeted at an industrial audience.

10. Bibliography

Major publications by the team in recent years

[1] 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.

[2] 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.

[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° 4, July-August 2004, http://www-sop.inria.fr/reves/publications/data/2004/DD04.

[5] X. GRANIER, G. DRETTAKIS. *A Final Reconstruction Framework for an Unified Global Illumination Algorithm*, in "ACM Transactions on Graphics", vol. 23, n° 2, April 2004, http://www-sop.inria.fr/reves/publications/data/2004/GD04.

[6] 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.

[7] 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° 3, July 2004, http://www-sop.inria.fr/reves/publications/data/2004/RMD04.

[8] 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.

[9] 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.

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

Books and Monographs

[11] 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.

[12] 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.

[13] 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.

Doctoral dissertations and Habilitation theses

[14] 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.

[15] 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.

[16] 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.

Publications in Conferences and Workshops

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[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] 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.

Internal Reports

[23] 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° 5479, INRIA Research Report, Sophia-Antipolis, January 2005, http://www.inria.fr/rrrt/rr-5479.html.

[24] F. DUGUET, C. H. ESTEBAN, G. DRETTAKIS, F. SCHMITT. *Level of Detail Continuum for Huge Geometric Data*, Technical report, n° 5552, INRIA Research Report, Sophia-Antipolis, April 2005, http://www.inria.fr/rrrt/rr-5552.html.

[25] G. LEMAITRE, E. GALLO. *Descripteurs audio. Librairie Matlab SPL et fichiers de descriptions ".sig"*, Technical report, n° En cours de soumission, INRIA, Sophia Antipolis, november 2005, http://www.inria.fr/rrrt/rt-0317.html.

Miscellaneous

[26] 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.

[27] F. DUGUET, C. H. ESTEBAN, G. DRETTAKIS, F. SCHMITT. *Level of Detail Continuum for Huge Geometric DataPoster at SIGGRAPH 2005*, July 2005, http://www.inria.fr/rrrt/rr-5552.html.

Bibliography in notes

[28] 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.

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

[30] 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.

[31] 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.

[32] 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.

[33] 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.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

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

[39] 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.

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

[41] 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.

[42] 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.