INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team REVES

Rendering and Virtual Environments with Sound

Sophia Antipolis

THEME COG

Activity Report

2004

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]

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

External advisor
Alexandre Olivier [3D Modelling]

Research scientist (partner)
Maria Roussou [FHW and UCL]
Ignacio Martin [Researcher University of Girona]

Invited Researcher
Simon Premoze [Columbia, December]

Technical staff
David Geldreich [part time IR, SEMIR, as part of the Workbench activities]

Project technical staff
Manuel Asselot [CREATE]

Technical staff
Gaël Braconier

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]

Interns
Gregory Bensa [DEA, ESSI]
Japinder S. Chawla [International Institute of Technology, Dehli, India]

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-immersif (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" [26], 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 [14], but also simplier methods such as accessibility maps [23]. The approaches we intend to investiage 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 and auditory consistency if required.

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [13], 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 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 type 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 [24] 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 [18] [17][16] typically have robustness or memory consumption problem 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 [15]. 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 [19] [24]), 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 [20] have 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 [25].

3.2. Virtual and Augmented Environments with Sound

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

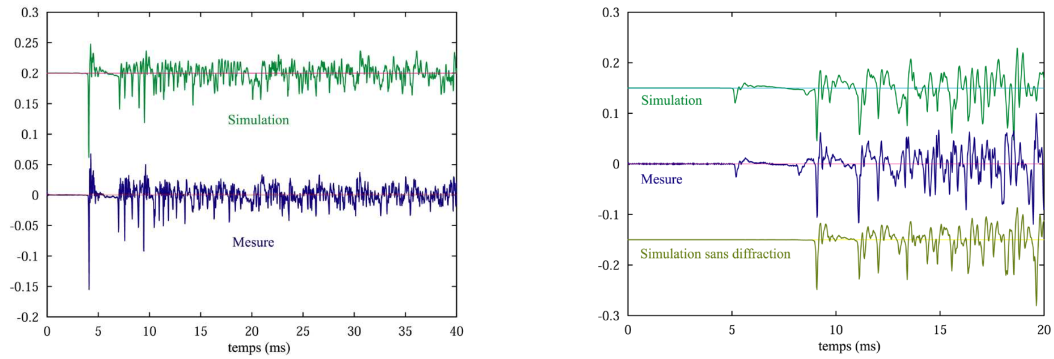

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

The second major research direction of our group is on virtual, augmented or mixed environments, which include both visual and sound representations. We are mainly interested in interactive environments, permitting the user to create and maniuplate scenes consisting of both real and synthetic objects. As a first step, we consider *real* objects to be digitised representations of reality, rather than the real world.

Our first goal is to apply and adapt our rendering expertise, presented in the previous paragraphs to virtual and augmented realy. There are three areas in which we concentrate our efforts: consistent lighting between real and synthetic illumination, for shadows and reflections, enriching virtual and augmented environments with sound, in a consistent manner and finally appropriate interaction and visual paradigms for virtual and augmented environments.

3.2.1. Efficient and Simple Relighting

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

Our previous work on interior relighting has given satisfactory solutions, allowing us to add virtual object with consistent lighting, but implies severe restrictions on the lighting conditions of input images [21][22]. Such approaches are based on the creation of "shadow free" base textures using heuristics, and a relatively precise reconstruction of the geometry. For outdoors scenes, geometric complexity and the fact that lighting conditions cannot be easily manipulated render such approaches less appropriate. However, some of the techniques developed can be applied, and we believe that the key is to combine automated techniques with user interaction at the various stages of the process.

The long-term goal is to turn on a video camera in a scene (potentially with partially pre-reconstructed geometry), and be able to add virtual objects or light sources interactively in a consistent manner into the video stream. Relighting could also be achieved in this manner, or using semi-transparent glasses or headsets. Applications of such an approach are numerous, for archeology, architecture and urban planning, special effects, manufacturing, design, training, computer games etc.

Figure 3. (a) Original conditions (b) The door has been removed virtually, and a virtual object and light have been added (method of [22])

This long term vision will require a way to smoothly vary from low-quality methods [21][22] to high quality approaches [27], 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 completent 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 both in images and in sound, artefacts, monuments building of lost civilisations.

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 is with the Foundation of the Hellenic World (FHW) http://www.fhw.gr, 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 is in the context of the EU IST project CREATE (see also Section 8.4.1). The combination of educational theories and field expertise allows a novel approach to the usage of Virtual Environments for both experts and novices. The first application of the project is in cultural heritage, and REVES contributes 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 takes 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, but 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 Danser 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. Results of this work were presented at the VAST, Computer Graphics and Archeology conference this year (see Section 6.1.4).

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 (see Section 8.4.1. Our collaboration with the CSTB has concentrated on three axes: application of global illumination for buildings, urban planning in the context of our common IST project CREATE (see also Section 8.4.1) 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. Despite initial encouraging contacts, the current economic downturn in the industry has limited our collaboration with industrial partners in this domain.

Integration of sound spatialisation, either geometric or statistical is clearly promising, and is currently almost entirely missing in exisiting games. We believe tjat the importance of these techniques and of their maturity for technology transfer is significant; we are currently finalizing a technology transfer agreement with EdenGames in this domain (see also Section 7.7).

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 fonctionalities include audio i/o plug-ins, audio buffer handling, 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. Point-Based Rendering and Shadows

Participants: George Drettakis, Emmanuel Gallo.

The work developed in 2001-2002 by Marc Stamminger at REVES has been integrated into the common software platform of the CREATE project. As a precursor to this step, two separate libraries have been developed, one for shadows and another for point-based rendering, which have been tested on real-world data for evaluation purposes with local companies with visualization needs.

5.3. RevesAPI: A common API for 3D graphics

Participants: Gael Braconier, George Drettakis, Florent Duguet, Marie-Claude Frasson, Alex Reche.

We started in 2002 the development of a common platform for research algorithm implementation. The platform is now operational and a significant amount of code has been developed in common by interns, PhD students, and researchers. This development platform exposes an API for applications which involve 3D graphics. This API is implemented on multiple platforms and can be compiled using different compilers. The platform is independent of the widget toolkit, and two different toolkits are currently available: Qt and the native Win32 API. The system runs on Windows using Visual C++ compiler, or the MinGW compiler (http://www.mingw.org), on linux using gcc-2.95 and gcc-3.x, and on IRIX using gcc or CC. Some parts of the system has been ported successfully to a Compaq iPAQ PDA running PocketPC 2002.

Tools for 3D graphics have been developed and are shared amongst the team members including a scene graph structure, an OpenGL renderer for this scene graph, and a (partial) VRML 2.0 parser. Application templates have been provided as well; using these templates a new user can have a customizable VRML viewer running with 4 lines of C++ code. The main motivation of this work was to gather and share as much code as possible so that current as well as new users can reuse code from another contributor with a minimal integration work. For example, a new user can write a specific renderer without having to care about the scene graph construction, the parsing of the geometry file (VRML 2.0 is partially supported), and the specificities of the platform the user works on; the only code (s)he has to write is the pure rendering code (see image for dynamic renderer selection).

This approach has recently been tested in practice, with complex data structures developed for one group member being used by another, thus greatly accelerated development time. Interface modules are also shared amongst group members.

6. New Results

6.1. Plausible Rendering

6.1.1. Flexible Point-Based Rendering on Mobile Devices

Participants: Florent Duguet, George Drettakis.

We have developed a flexible point-based rendering algorithm that works on PDAs. A 3D object is first encoded into a hierarchical recursive grid data structure, and then efficiently rendered using a dedicated point-based rendering algorithm. We presented a study of the type of hierarchical structure which should be used. We chose the 3x3x3 recursive grid (which we call the tri-grid), for its good compromise on rendering speed and storage compactness. We extended a rendering algorithm presented previously and made it flexible to allow rendering of the structure at different levels of the hierarchy in the same rendering pass. With this algorithm, we can achieve efficient view-frustum culling, and automatic, view-dependent level of detail. We also presented a one-pass shadow-map algorithm, avoiding expensive matrix transformations.

These algorithms have been implemented on an iPAQ HP3850, with an ARM processor at 200MHz, running PocketPC 2002. Objects sampled at 1.3 Million points can be rendered with shadows at interactive framerates, i.e, 2.11 fps. The paper has been published in IEEE CG&A, special issue on Point-Based Graphics, July 2004 [1] (see figure 4).

Figure 4. Rendering a sampling of 55 million points of the Stanford Lucy model (28 million faces), on a 4096x4096 image; computed in 4 seconds

6.1.2. Interactive audio rendering

Participants: Emmanuel Gallo, Nicolas Tsingos, Guillaume Lemaitre, George Drettakis.

In the context of the european IST project CREATE, we developed several techniques aimed at accelerating the audio rendering of complex scenes, comprising a large number of point sources. The first techniques groups sound sources based on geometrical and psycho-acoustical information updated in real-time. This technique dynamically best-fits a specified number of clusters to the sound scene and, as such, can be used for dynamic allocation of computational resources. We also explored how programmable graphics hardware can be used in this context to speed up audio calculations. A second technique uses psycho-acoustics to evaluate auditory masking occuring in the sound scene in real-time. This information is then used to cull (i.e., not render) inaudible sound sources. This work was published as an INRIA technical report and also appeared as a technical paper in the 2003 Online Audio Resource Guide of Gamasutra.com, a web site specializing in computer gaming. We further improved the techniques and

validated our approach by conducting psycoacoustic tests on 20 subjects. The results have shown that our approach can be used to significantly simplify a complex 3D audio environment without degradation in the perceived audio quality nor 3D audio localization abilities. The method is illustrated in Figure 5.

These results appeared at the conference SIGGRAPH 2004, and have French and US patents pending (see Section 7.6).

We are extending this approach to more complex auditory environment segmentation methods, more accurate masking models and more efficient cluster processing in the context of the RNTL project OPERA. In

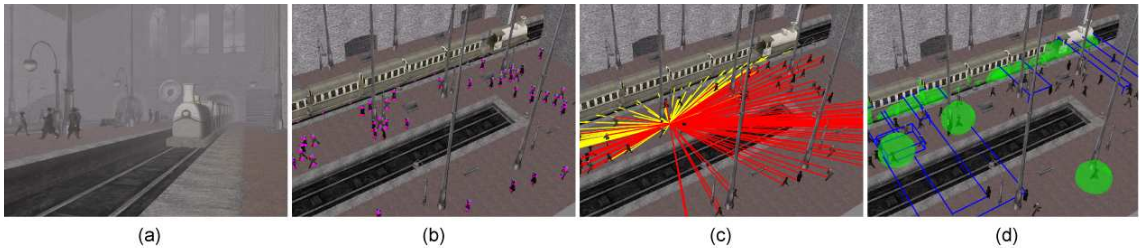

Figure 5. (a) An application of the perceptual rendering pipeline to a complex train-station environment. (b) Each pedestrian acts as two sound sources (voice and footsteps). Each wheel of the train is also modeled as a point sound source to get the proper spatial rendering for this extended source. Overall, 160 sound sources must be rendered (magenta dots). (c) Colored lines represent direct sound paths from the sources to the listener. All lines in red represent perceptually masked sound sources while yellow lines represent audible sources. Note how the train noise masks the conversations and footsteps of the pedestrians. (d) Clusters are dynamically constructed to spatialize the audio. Green spheres indicate representative location of the clusters. Blue boxes are bounding boxes of all sources in each cluster.

particular, we have been experimenting with scalable signal processing approaches and "impostor" representation where the contribution of a group of sound sources is replaced by a "texture" directly retrieved from an atlas of pre-computed sounds or recordings. We are also developing a more accurate sound priority metric accounting not only for the loudness of the sound signals but also on higher level cognitive effects, such as the perceived "urgency" of a signal.

6.1.3. Perceptive multi-modal rendering

Participants: Manuel Asselot, Nicolas Tsingos, Japinder S. Chawla, George Drettakis.

The goal of this research is to efficiently render virtual scenes taking advantage of the human perception of the environment. When a series of events occurs, may them be aural, visual, or both at the same time, our attention targets only some of them. For instance, if an object is much more contrasted than all those surrounding it, or if we can hear some shrill sound at a given time, we will tend to focus our attention on that precise event. An aural event can also influence the perceptive importance of a visual event that would correspond to its source. The perceptibility of a sound would conversely be modified if its source was visually important.

We started work on this research direction by initially concentrating on the study of attention only for the visual aspect, which represents a fairly large part of the problem. The aim was to develop an algorithm capable of detecting the most important zones in a 3D scene, depending on the point of view without using a tracking device. The idea we exploited was to use two different kinds of analysis and make them converge into a single, plausible, saliency map.

The first analysis only takes the scenegraph data into account whereas the second makes use only of the information given in the pixels. Two sorts of applications could use this algorithm. On the one hand, the rendering process could be accelerated using an attention guided level of details algorithm. On the other hand, we could refine the important parts of a scene if they do not appear with the required importance, effecting attention control. Prototype scenegraph and pixel based analysis tools have already been developed. We are now working on level of detail control issues and merging of our two analysis tools.

6.1.4. A Point-Based Approach for Capture, Display and Illustration of Very Complex Archeological Artefacts

Participants: Florent Duguet, Daniel Girardeau Montaut, George Drettakis.

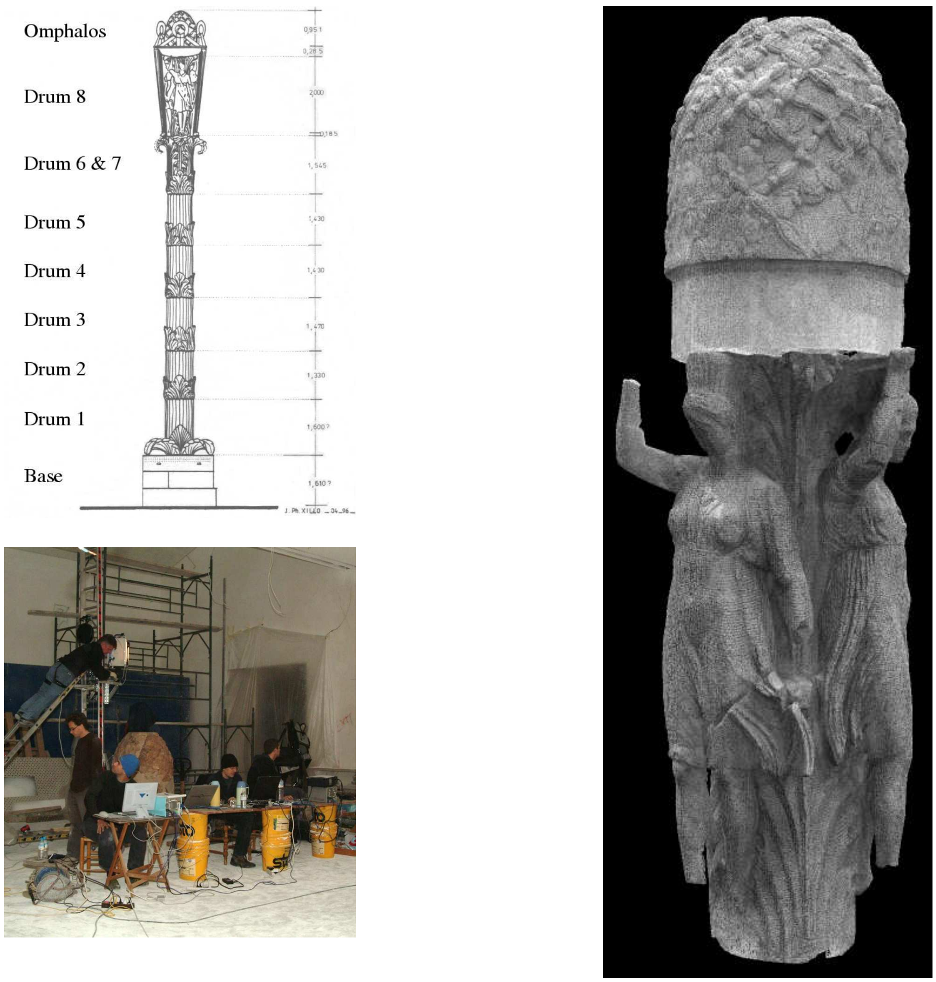

Figure 6. (Left) An artists illustration of the Dancers Column (Middle) A photograph from the capture in Greece (c) A reconstruction of part of the column using our visualization technique.

Figure 7. Skydome rendering of the omphalos developpée, artifact of the Delphe scanning project. Image 2270x1024 computed in 5 minutes - 40 directions per pixel.

In this work we developed a complete point-based pipeline for the capture, display and illustration of very large scans of archeological artifacts. This approach was developed as part of a project involving archeologists and computer graphics researchers, working on the Delphi “Dancers Column". We first determined the archeologists requirements for interactive viewing and documentary illustration. To satisfy these needs we used a compact point-based structure of the very large data, permitting interactive viewing in 3D. This helps the archeologists to examine and position the fragments. We introduce efficient construction algorithms for this structure, allowing it to be built on limited-memory platforms, such as those available on the field. We also propose a new stylized rendering approach based on an inverse cylindrical projection and 2D skydome rendering. This illustrative style has been used as a planning tool for fragment docking and as a substitute for traditional illustration in an archeological publication. Other uses of these tools are currently under way in the context of this project. This work was presented at the VAST 2004 conference [7], in collaboration with Jean-Luc Martinez of the Louvre and the Ecole Française d'Athènes, and Francis Schmitt of the ENST.

The capture process works as follows: A set of scan sheets, calibrated into a single coordinate system, is a sampling of a real 3D object surface using points. In our project, we have been working with cylindrically shaped objects (i.e., a column). We thus unfolded those point clouds on a plane and build an image of heights. This image of heights is considered as the sampling of a terrain (2D1/2). We can then compute for each pixel of this image, the portion of visible sky. This lighting technique provides the archeologists with a very informative image of the object with a good feeling of underlying 3D form, Figure 7.

This technique has been extended into 3D with a light tracing kind of algorithm: for each direction of the light, a fine grid of cells is lit with a sweeping algorithm. A very small amount of operations are needed to compute the lighting from one sweeping plane to the next.

A detailed description, in particular of the 3D algorithm was presented at the French Computer Graphics Conference AFIG in Poitiers [8].

6.1.5. Recovering structure and aspect detail from example

Participants: Marie-Clause 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 analysized 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. We are currently working on recovering the aspect (interior and borders) of the regions in order to generate procedurally completely new textured structures. Current results are promising.

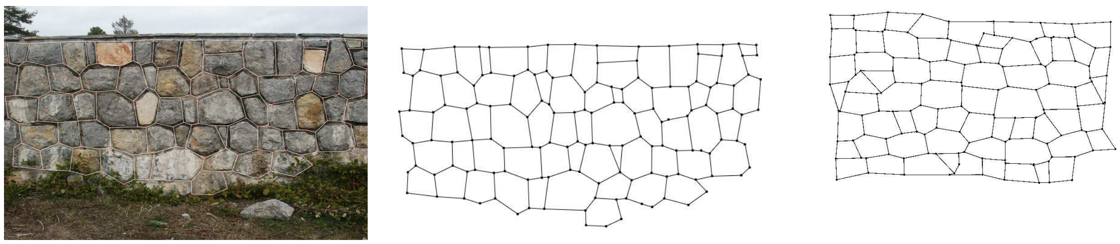

Figure 8.

6.2. High-Quality Rendering

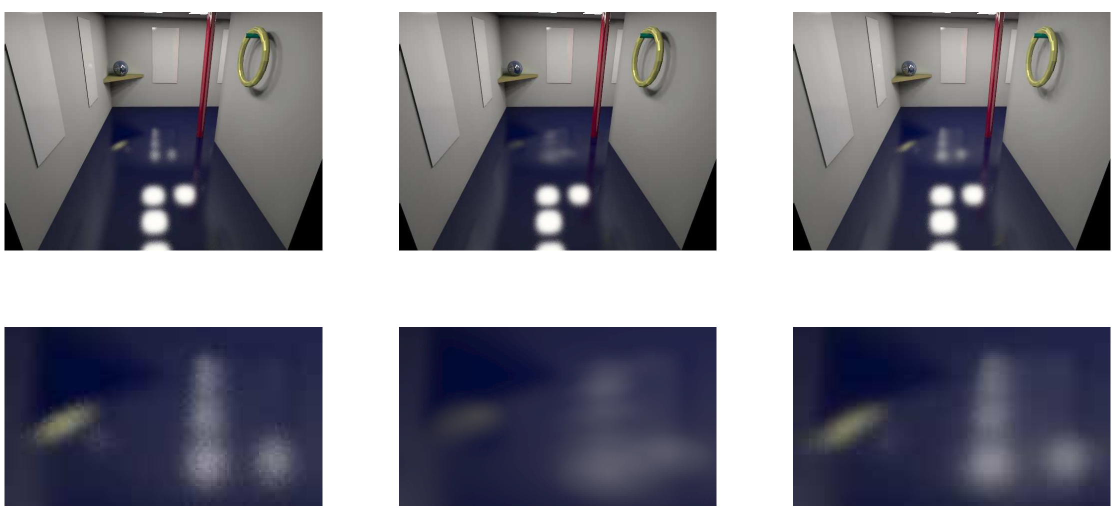

Figure 9. Images obtained using the non-diffuse final reconstruction method. Lower part, zoom into the rear part of the glossy corridor floor. >From left to right: Rays only (too noisy), particles only (detail washed out) and the new method.

6.2.1. A Final Reconstruction Approach for an Unified Global Illumination Algorithm

Participants: Xavier Granier, George Drettakis.

We have developed a novel final reconstruction step based on the viewpoint in the context of a uniform approach for global illumination, combining the advantages of radiosity calculations for the diffuse component and particle and ray-tracing for glossy/specular effects. With our approach, we can deal with difficult lighting configurations such as indirect non-diffuse illumination. The first step of this algorithm consists in a view-independent solution based on hierarchical radiosity with clustering, integrated with particle tracing. This first pass results in solutions containing directional effects such as caustics, which can be interactively rendered. The second step consists in a view-dependent final reconstruction that uses all existing information to compute higher quality, ray-traced images. This approach combines ray-tracing and the particles stored in object space to greatly improve image quality (see Figure 9). This work appeared this year in ACM Transactions on Graphics [2].

Figure 10. Left: the original image. Right: the computed alpha-mask (note that we used black for opaque for clarity of the figure).

6.2.2. Volumetric Reconstruction and Interactive Rendering of Trees from Photographs

Participants: Alex Reche, Ignacio Martîn, George Drettakis.

The goal of this research is to efficiently capture very realistic trees that can be use for interactive rendering. All the reconstruction is done using simply a series of pictures around a tree that are consequently pre-processed to first calibrate the cameras to be able to know where the pictures are taken with respect the tree and then alpha matte the images to separate the tree from the background (see Figure 10).

The reconstruction technique used to reconstruct the opacities of the tree is highly inspired on tomography techniques used in medical imaging. One difference is in the input images, in tomography the input images are pictures of density, in our case the input images are pictures of opacities (alpha maps extracted using alpha matting techniques). The other, main difference is that the volumetric reconstruction only captures the low-frequency characteristics of the tree, and has no colour information. Clearly, this is insufficient for the purposes of visualisation. As a result, to capture the high-frequency detail, the final tree is enhanced with view dependent textures extracted from the original images (see Figure 11).

This research was published in the ACM Transactions On Graphics and presented at Siggraph'04 in Los Angeles in August 2004 .

Figure 11. Screenshots of our interactive renderings of two reconstructed trees.

6.3. Virtual Environments with Sound

6.3.1. Perceptive 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. [4], and described in paragraph 6.1.2. Ongoing work seeks to improve this pipeline in two main directions: improve the segmentation of the scene in clusters, and optimize the computation of a cluster representative. 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, whereas the second approach consists in matching the cluster with some sound textures stored in a precomputed library.

6.3.2. User-Centered Approach on Combining Realism and Interactivity in Virtual Environments

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

This work is part of the EU IST project CREATE.

As part of the EU IST project CREATE, we have developed an series of virtual environment enhancements in an integrated software platform, incorporating view-dependent texture mapping, point-based rendering for vegetation, perspective shadow maps and 3D spatialised sound. All of these components have been integrated into the common CREATE platform based on Performer, CAVElib and the high-level scripting language XP developed at EVL and FHW allowing easy integration of novel functionalities.

This system has been used in many demonstrations, and we have an official collaboration with the architects of the Tramway project of the city of Nice using our tools. A specific description of the VE enhancements developed by REVES was presented as a poster at IEEE VR2004, and included as an extended abstract in the conference proceedings [10].

Figure 12. Two faces of a panorama cubemap used for image modeling. Wireframe of the resulting model and a view with extracted textures.

The complete image-based 3D capture process for the creation and display of photorealistic virtual environments (VEs) was presented at the Eurographics Symposium on Virtual Environments 2004 [6]. The process described results in which VEs aim to realistically recreate existing real-world scenes that can be displayed in a range of immersive VR systems using a high-quality, view-dependent algorithm and further enhanced using advanced vegetation, shadow display algorithms and 3D sound. The two target scenes for CREATE, an archaeological site and an urban environment, were chosen according to real-world applications in the areas of urban planning/architecture and cultural heritage education. The users in each case are able to reconstruct or manipulate elements of the VEs according to their needs, as these have been specified through a detailed user requirements survey. The example of Place Garibladi is shown in Fig. 12; more images are shown in Section 8.4.1. Furthermore, a user task analysis and scenario-based approach was adopted for the design of the virtual prototypes and the evaluation.

In particular we studied the workflow of the architects and developed an interface in a Virtual Environment which could help them in their work. A set of user experiments were designed and performed with the architects of the real Tramway project, investigating the influence of interactivity and realism on the Virtual Environment, as well as its utility and applicability in urban design. The results of this experimental study are submitted for publication.

6.4. Interaction and Visual Paradigms for Virtual and Augmented Environments

6.4.1. A modeling by drawing tool using a relief metaphor

Participants: David Bourguignon, Raphaelle Chaine, Marie-Paule Cani, George Drettakis.

This project aims at providing a new modeling system which takes advantage of two-dimensional drawing knowledge to design three-dimensional free-form shapes. A set of mouse or tablet strokes is interpreted by the system as defining both a two-dimensional shape boundary and a displacement map. This information is used for pushing or pulling vertices of existing surfaces, or for creating vertices of new surface patches. Moreover, to relieve the burden of 3D manipulation from the user, patches are automatically positioned in space. (See figure 13.)

The iterative design process alternates a modeling by drawing sequence and a viewpoint change. In order to stay as close as possible to the traditional drawing experience, the system imposes the minimum number of constraints on the topology of either the strokes set or the resulting surface. Overall, the system has been successful at sketching rough 3D shapes.

The current state of the project is very promising and it has been published in the Proceedings of the First Eurographics Workshop on Sketch-Based Interfaces and Modeling 2004. However, several interesting venues remain to be explored which will certainly improve both the usability of the modeling tool and the quality of the models generated.

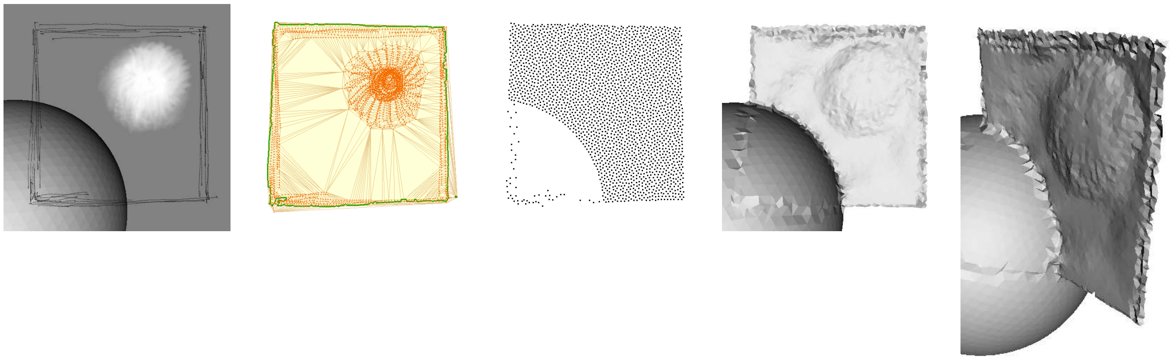

Figure 13. Tool workflow. In (a), the user has drawn both on an existing surface (an imported sphere) and on empty space, with a pencil (hard black strokes) and a brush (soft white strokes). In (b), reconstruction of the boundaries (green lines) of the drawn shape. In (c), adaptive sampling of the displacement regions. In (d) and (e), new surface obtained after reconstruction.

7. Contracts and Grants with Industry

7.1. CSTB

Participants: George Drettakis, 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.

7.2. REALVIZ

Participants: George Drettakis, Alex Reche, Alexandre Olivier-Mangon.

We have an ongoing research contract with REALVIZ which is related to the development of novel algorithms and procedures for reconstruction and from images and view-dependent rendering, in the context of the CREATE project.

7.3. Alias|Wavefront

We are part of the Alias|Wavefront software donation program. Our artist Alexandre Olivier uses 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 creates 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).

The data set will include up to 300 million points, which is a real challenge to display. Most recent graphics cards can display 60 million points per second, but the memory required for such a point set is 3.6 Giga bytes. Our work will concentrate on post-processing of the data, for better reconstruction and for efficient interactive rendering, as part of the F. Duguets thesis.

This project started last year continues and produces interresting scientific challenges, for example those described in the VAST publication [7], and AFIG publication [8]. Florent Duguet has been working on technilogy transfert to EDF in february for two weeks, in the EDF R&D center, Clamart.

7.5. Collaboration with ENST

Florent Duguet will be from june 2004 to september 2005 a visiting student at ENST Paris (http://www.enst.fr). During this period, many scientific collaborations and exchanges will take place, such as the one with Daniel Girardeau-Montaut described in paragraph [8].

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 perfectionneés de spatialisation du son" (Advanced System and Method for Sound Spatialisation).

7.7. Technology transfer

Participants: Nicolas Tsingos, Emmanuel Gallo, George Drettakis.

We are also in the process of transfering the 3D sound technology to the french game company Eden Games, a subsidiary of ATARI, in the context of two upcoming games for next-generation console hardware.

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, Gaël Braconier.

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 PIII 1GHz + GeForce 4 Ti 4600 graphics) running under Linux. 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 for interaction and an additional captor for view-point tracking and view-dependent rendering.

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.

This year D. Geldreich and N. Tsingos also worked on the final set-up of the workbench in an acoustically-treated room where the immersive display will be combined with an immersive binaural or 6.1-surround sound restitution system. The system has been operational since December 2004.

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 [7] or alternative representations of trees [3]. The other participants of this project are IPARLA in Bordeaux and ISA in Nancy.

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 has started early 2004. First results and orientations are described in section 6.3.1.

8.3. Visiting Researchers

This year we invited the following visitors: Ignacio Martin of the University of Girona (in January and in November), Michael Gleicher of the University of Wisconsin and Karri Pulli of Nokia, as part of their visits to IMAGINA in February, Georg Essl, of the Ireland Media Lab (April 2004), Ken Perlin from New York University in May, Jim Hanan of University of Queensland (Australia) in June, Pierre Poulin of the University of Montreal in July, Maria Roussou of the Foundation of the Hellenic World and University College London (July-August), Nathan Litke of Caltech in September, Simon Premoze of Columbia University (November), Dani Tost and Pau Esteaella from UPC Barcelona and Girona (November), and Sylvain Lefebvre of the EVASION/GRAVIR group at INRIA Rhône-Alpes in December.

8.4. European Projects

Participants: George Drettakis, Alex Reche, Emmanual Gallo, Alexandre Olivier-Mangon, Nicolas Tsingos.

8.4.1. CREATE

As mentioned previously in various sections, a significant part of our activities this year (as in last year) has concentrated around the EU IST CREATE project“Constructivist Mixed Reality for Design, Education, and Cultural Heritage” http://www.cs.ucl.ac.uk/create This project is coordinated by the University College London, and has two main axes of research: the application of educational constructivist theories to improve interactive learning, and the use of more realistic virtual environments. Our intervention is mainly on the provision of more realistic rendering, both for image and sound in virtual environments. The other partners of the project are UCL (UK), FHW (GR), REALVIZ, CSTB, PERCRO (I) and UCY (CY).

This was the final year of the project, in which the two application demonstrators finalized. The first concerns the site of ancient Messene, built by the Foundation of the Hellenic World and used the haptic device built by PERCRO. The demonstrator was used to improve learning in the architectural and archeological domain.

The second demonstrator is the Tramway project of the city of Nice, for which public works started in autumn 2003. After a series of demonstrations to officials and engineers of the real project (see 9.4), we established an official convention of collaboration with the "Mission Tramway" which is the organisation in charge of the implementation of the project. This collaboration started last year and continued throught this year.

As mentioned in section 6.3.2, we produced two publications this year on the virtual environment of Garibaldi and its application to urban planning [10][6], and a publication currently being submitted concerning the work with the architects. Examples of the resulting virtual environment we created are shown in Figures 14,15.

Figure 14. (Left) A view from a balcony from the virtual environment simulator which also runs on the workbench. (Right) A view on the ground including crowds, shadows and vegetation. 3D spatialized sound is also present, including about 150 sound sources (a voice per pedestrian).

The project ends in February 2005, but our contribution is complete. We co-organized (with the CSTB), the final review of CREATE in Sophia-Antipolis at the end of November, and the review results were very positive. Overall the project was a success. First, in terms of novel scientific production (this year four CREATE-related publications appeared [10][6][4][3]), including two SIGGRAPH papers and several other papers were published last year. Second, it gave us the opportunity and the motivation to apply our research to real-world applications, and to work hand-in-hand with end users. This was a particulary rewarding and motivating experience, and we intend to pursue such activities in the future. Finally, there is a true potential for technology transfer from this project. A technology transfer agreement has already been initiated concerning the 3D sound (see Section 7.7), and the entire CREATE system may be exploited in some form by the CSTB as a service for urban planning.

8.5. Bilateral Collaborations

8.5.1. France-Greece (Hellas)

Our work in common with Maria Roussou of the FHW has continued in the context of the CREATE project, and with the presentation of two publications (IEEE VR [10], Eurographics [6]) this year, as well as the work with the Virtual Environment evaluation with the architects this summer, which is still active.

8.5.2. France-Spain

We have continued to be in contact with the research groups of Girona and also UPC Barcelona. Ignacio Martin has visited us this year to continue working on projects started on 2003 (in January). A paper was presented on this work at the Siggraph Conference [3]. Further work is under way with his Ph.D. student Pau Estaella, co-supervised with Dani Tost from the UPC Barcelona.

8.5.3. France-United States of America

We currently have active research contacts with Simon Premoze of Columbia University (NY) and Prof. Ashikhmin at Sunnybrook (NY). Dr. Premoze visited REVES for a week in December as part of this project, on illumination for complex scenes with vegetation.

9. Dissemination

9.1. Participation in the Community

9.1.1. Programme Committees

G. Drettakis was a member of the programme committees for the VAST2004/2nd Eurographics workshop on Graphics and Heritage, the Eurographics Symposium on Point-Based Graphics (2004), the ACM Siggraph Symposium on Interactive 3D Graphics and Games 04, and the Graphics Interface 2004 conference, and reviewer papers for a number of journals and conferences (ACM SIGGRAPH 2004, EGSR 04, Visual Computer, Image Communication, Journal of Tree Physiology, CHI'04). Nicolas Tsingos reviewed papers for the Journal of the Acoustical Society of America, Eurographics Symposium on Rendering (2004), Eurographics Symposium on Point-Based Graphics (2004), Journal of Discrete Algorithms. He also took part to ACM Multimedia Art Program Comittee and the commitee for the best paper in journées de l'association francaise d'informatique graphique (AFIG) 2004.

9.1.2. Thesis Committees

G. Drettakis was an external examiner for the Habilitation thesis of M. Paulin at the University of Toulouse.

9.1.3. 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

9.2. Teaching

9.2.1. University teaching

George Drettakis was responsible for the Computer Graphics course at ISIA (Ecole des mines) in January 2004 (15 hours). Together with Nicolas Tsingos, they organized and taught the new Image and Sound Synthesis course at the newly formed Masters program at the Univeristy of Nice. This is a 15 hour, 5-session course, with a programming project and exam. Florent Duguet was teaching assistant (as part of his PhD AMX fellowship) at the University of Nice, supervising laboratories of Java beginner programming courses and general computer science laboratories. Nicolas Tsingos gave about six hours courses at ISIA and seven hours courses at the IGMMV Master. Marie-Claude Frasson gave two hours of a Ray-Tracing course to second-year engineering students at the ESSI (University of Nice), plus eight hours of associated laboratories, and G. Drettakis also taught one hour in this course.

9.2.2. Ongoing PhDs

Marie-Claude Frasson, Emmanuel Gallo and Alex Reche are registered at the University of Nice. Emmanuel Gallo started his PhD early April 2004. Marie-Claude and Alex will complete their PhD's in early 2005. Florent Duguet is beginning is fourth year, also at the University of Nice Sophia-Antipolis (UNSA). Since June, he is co-directed by Francis Schmitt and is physically located at the ENST, Paris.

9.3. Participation at conferences

9.3.1. Presentations at Conferences

Alex Reche presented the paper at EGVE [], in Grenoble in May. Maria Roussou presented the poster [10] at IEEE VR, in Chicago in March. The two SIGGRAPH papers [4][3] were presented at the SIGGRAPH conference in Los Angeles by Nicolas Tsingos and Alex Reche respectively in August. Emmanuel Gallo presented a short paper as a poster at GP2 workshop in Los-Angeles, United States. Florent Duguet gave a presentation of his paper [8] at the Journées de l'Association Francaise d'Informatique Graphique in Poitiers, France. George Drettakis presented the paper at VAST [7] in Brussels in December.

9.3.2. Participation at Conferences and Workshops

Other than the presentations of papers and posters presented above, M-C. Frasson, F. Duguet, E. Gallo and A. Reche went to the Eurographics Rendering Symposium on Rendering in Leuven, Sweden; F. Duguet and G. Drettakis went to the ACM Siggraph in Los Angeles, USA. M-C Frasson also participated at Eurographics 2004, and Afrigraph 2004; Guillaume Lemaitre took part (as an organization member) to the Sound Design Colloquium in Paris in October 2004. N. Tsingos went to the CREATIVITY 2004 event in London in November 2004, organised by the audio hardware manufacturer Creative Labs.

9.4. Demonstrations and Press

Participants: Alexandre Olivier-Mangon, George Drettakis, Alex Reche, Manuel Asselot, Emmanual Gallo, Marie-Claude Frasson, Nicolas Tsingos.

We demonstrated the results of our project (notably the sound and CREATE work) at IMAGINA 2004 in February in Monaco.

In the context of CREATE in collaboration with the greater Nice area council (CANCA, Communauté d'Agglomération Nice Cote d'Azur) we had several occasions to present the Garibaldi virtual reconstruction and work with officials involved in the project. We presented a version of the work of the Garibaldi square simulation to the architects of the project in March and also at the Town Hall in a meeting of the working group of the Tramway project in June.

To mark the conclusion of the CREATE project, a public event was organised at INRIA, co-located at the RealityCenter of the CSTB, the day after the final review. Public officials (A. Sans, CANCA) as well as members of the Tramway project (M. Bonis, D. Clavel) were invited and the Garibaldi virtual environment was shown (see Figure 15).

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. In particular, we presented demonstrations to representatives of Hitachi in June and Procter & Gamble in September. A Chinese delegation from the LIAMA also visited the lab in December and a workbench demonstration was presented on this occasion.

A press release was produced for the final event of CREATE and for Marie-Claude's participation in the ACM project (see below).

9.5. Miscellaneous

Participant: Marie-Claude Frasson.

Marie-Claude Frasson was selected to participate in the ACM Siggraph Pilot Project : "Women in CG Tour " aimed at encouraging women to consider a career in Computer Graphics. This tour took place in South Africa and is planned to happen every year in a different country. Five computer graphics women (amongst which Barbara Helfer, the ACM Siggraph Vice-President and Colleen McCase, the ACM Siggraph President for Education) along with a South African Professor and another known Graphics Researcher, Wolfgang Strasser, toured a number of universities (7) around South Africa to present the Computer Graphics field, its differents aspects and possibilities and why it could be a very interesting field to go into, especially for women. The

Figure 15. The four different choices for the Garibaldi square, proposed by the Mission Tramway" as part of a public hearing in the summer of 2004. We reconstructed the options and presented them as part of our public event to city and Tramway project officials.

audience was composed of university women (black and white) and schoolgirls from neighboring schools. All the events were very successful and there will be a follow-up in the next years to evaluate the success of the tour, i.e. how many women decided to go into computer science and especially computer graphics.

10. Bibliography

Articles in referred journals and book chapters

[1] 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.

[2] 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.

[3] 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.

[4] 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.

Publications in Conferences and Workshops

[5] D. BOURGUIGNON, M.-P. CANI, R. CHAINE, G. DRETTAKIS. *Relief: A Modeling by Drawing Tool*, in "Eurographics Workshop on Sketch-Based Interfaces", J. A. JORGE, J. F. HUGHES (editors)., Eurographics, August 2004, http://www-sop.inria.fr/reves/publications/data/2004/BCCD04.

[6] 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.

[7] F. DUGUET, G. DRETTAKIS, D. GIRARDEAU-MONTAUT, J.-L. MARTINEZ, F. SCHMITT. *A Point-Based Approach for Capture, Display and Illustration of Very Complex Archeological Artefacts*, in "The 5th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2004)", F. N. N. S. (. K. CAIN (editor)., Eurographics, December 2004, http://www-sop.inria.fr/reves/publications/data/2004/DDGMS04.

[8] F. DUGUET, D. GIRARDEAU-MONTAUT. *Rendu en Portion de Ciel Visible de Gros Nuages de Points 3D*, in "Actes des 17ièmes Journées de l'Association Française d'Informatique Graphique (AFIG)", 2004.

[9] E. GALLO, N. TSINGOS. *Efficient 3D Audio Processing on the GPU*, in "GP2, ACM Workshop on General Purpose Computing on Graphics Processors", ACM, August 2004, http://www-sop.inria.fr/reves/publications/data/2004/GT04.

[10] M. ROUSSOU, G. DRETTAKIS, N. TSINGOS, A. RECHE, E. GALLO. *A User-Centered Approach on Combining Realism and Interactivity in Virtual Environments*, in "Proc. of IEEE VR 2004 (ext. abstract),", IEEE, March 2004, p. 251-252, http://www-sop.inria.fr/reves/publications/data/2004/RDTRG04.

Internal Reports

[11] F. DUGUET. *Robust Higher-Order Filtering of Points*, Technical report, n° 5165, INRIA, Sophia-Antipolis, april 2004, http://www.inria.fr/rrrt/rr-5165.html.

[12] F. DUGUET. *Shadow Computations using Robust Epsilon Visibility*, Technical report, n° RR-5167, INRIA, REVES/INRIA Sophia-Antipolis, April 2004, http://www.inria.fr/rrrt/rr-5167.html.

Bibliography in notes

[13] 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.

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

[15] 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.

[16] 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, Juin 1996, p. 245–257.

[17] 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.

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

[19] 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)", Juillet 1998.

[20] 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.

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

[22] C. LOSCOS, M.-C. FRASSON, G. DRETTAKIS, B. WALTER, X. GRANIER, P. POULIN. *Interactive Virtual Relighting and Remodeling of Real Scenes*, in "Rendering Techniques'99 (10th Eurographics Workshop on

Rendering)", Springer Verlag, Juin 1999, p. 329–340.

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

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

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

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

[27] 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.