INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team REVES

Rendering and Virtual Environments with Sound

Sophia Antipolis

THEME 3B

Activity Report

2003

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]

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

External advisor
Alexandre Olivier [3D Modelling]

Research scientist (partner)
Fredo Durand [Assistant Professor at MIT]
Maria Roussou [FHW and UCL]

Invited Researcher
Ignacio Martin [Researcher University of Girona, November-December]

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

Project technical staff
Emmanuel Gallo [January-July 2003]

Technical staff
Gaël Braconier

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

Interns
Matthieu Cunzi [MOO, ARC ARCHEOS]
Charles-Felix Chabert [DEA, UNSA]
Manuel Asselot [Centrale, March-December 2003]
Roman Bayon [ENPS, March-September 2003]

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

Key words: *rendering*, *image rendering*, *sound rendering*, *plausible rendering*, *high-quality 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" [31], 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 [19], but also simplier methods such as accessibility maps [28]. 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 [18], 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 [29] 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 [23] [22][21] 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 [20]. 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 [24] [29]), 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 [25] 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 [30].

3.2. Virtual and Augmented Environments with Sound

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

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

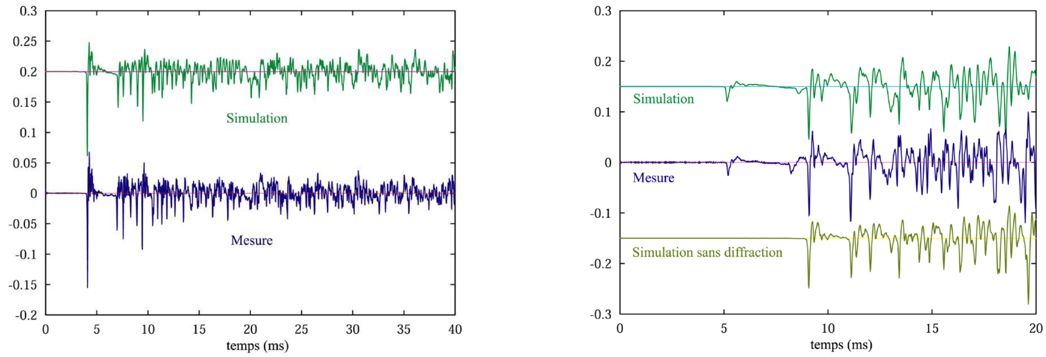

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

real and synthetic illumination, for shadows and reflections, enriching virtual and augmented environments with sound, in a consistent manner and finally appropriate interaction and visual paradigms for virtual and augmented environments.

3.2.1. Efficient and Simple Relighting

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

Figure 3. (a) Original conditions (b) The door has been removed virtually, and a virtual object and light have been added (method of [27])

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 [26][27]. 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 [26][27] to high quality approaches [32], 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

Key words: *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 is in the context of the ARC ARCHEOS (see also Section 8.2.1) and the third is with the Virtual Reality group of EDF and a programme funded by the "Fondation EDF".

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 ARCHEOS, we have worked on the Agora site of ancient Argos, and in particular on the Tholos monument. We have 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 is 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 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.

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

Key words: *evaluation*, *training*, *architecture*, *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. 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 work in the context of a transportation project with the INRETS.

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

Key words: *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 hope to be able to convince industrial partners of the importance of these techniques and of their maturity for technology transfer.

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

Key words: *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. Global Illumination

Participants: Gaël Braconier, George Drettakis, Ignacio Martin, Alex Reche.

We are developing a global illumination system, ECLAIRES, based on the work carried out by Xavier Granier during his Ph.D. thesis. ECLAIRES is a radiosity system based on hierarchical clustering methods. ECLAIRES is a multi-platform system running under Linux, Windows and Irix and is based on the REVESAPI common library (see 5.4).

The aim of this developement, supported by a national Inria funding for a software engineer (G. Braconier), is to provide a robust library for computing global illumination for realistic environments. The two applications concerned are architecture and archeology. For urban planning and architectural applications, we collaborate closely with the CSTB, which provides us models of real buildings allowing us to validate our algorithms.

We are investigating the use of semantic information for architectural scenes, based on the industry standard IFC (Industry Foundation Classes), and the unified approach which combined particle tracing and radiosity, for non-diffuse and diffuse transfer respectively.

5.2. AURELI: AUdio REndering LIbrary

Participants: Nicolas Tsingos, Emmanuel Gallo, Charles-Félix Chabert.

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.

5.3. 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.

5.4. RevesAPI: A common API for 3D graphics

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

We started last year 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.

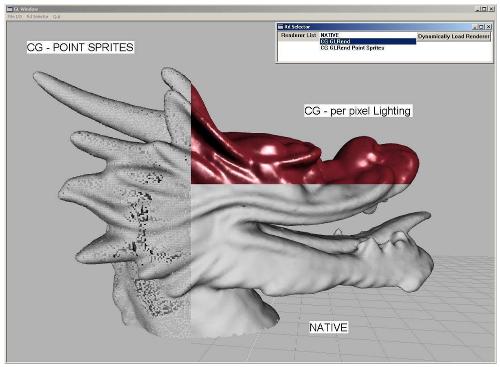

Figure 4. An image showing three different types of rendering possible under REVESAPI. Native is the standard REVESAPI renderer, the other two are CG (pixel/vertex shader) based renderers.

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.

This work is described in the research report RR-4833 [15], and has been submitted for publication. Photographs of the algorithm running on the PDA are shown in Fig. 5.

6.1.2. Perspective Shadow Maps for Games

Participant: George Drettakis.

Collaboration with Marc Stamminger at the University of Erlangen.

Our work on perspective shadows maps, presented last year at SIGGRAPH 2002 was extended and adapted to a Computer Games context. This work was presented in a more "wider public" context as a chapter part Computer Games Gems IV [3].

6.1.3. Interactive Image-Based Modelling using Points

Participants: Marie-Claude Frasson, George Drettakis.

In collaboration with Pierre Poulin at the University of Montréal, and Marc Stamminger at the University of Erlangen.

We have developed a novel approach which permits point-based modeling from images. As described previously (Sec. 6.1.1), complex objects can be displayed rapidly when represented by points rather than polygons. Based on this idea, we have built an interactive system which permits automatic reconstruction of complex geometric objects such as plants, organic matter and generally irregular geometry. The main emphasis of our solution is the fact that the user can interactively control the process at every step of the reconstruction, thus significantly improving the quality of the geometry generated and rendered.

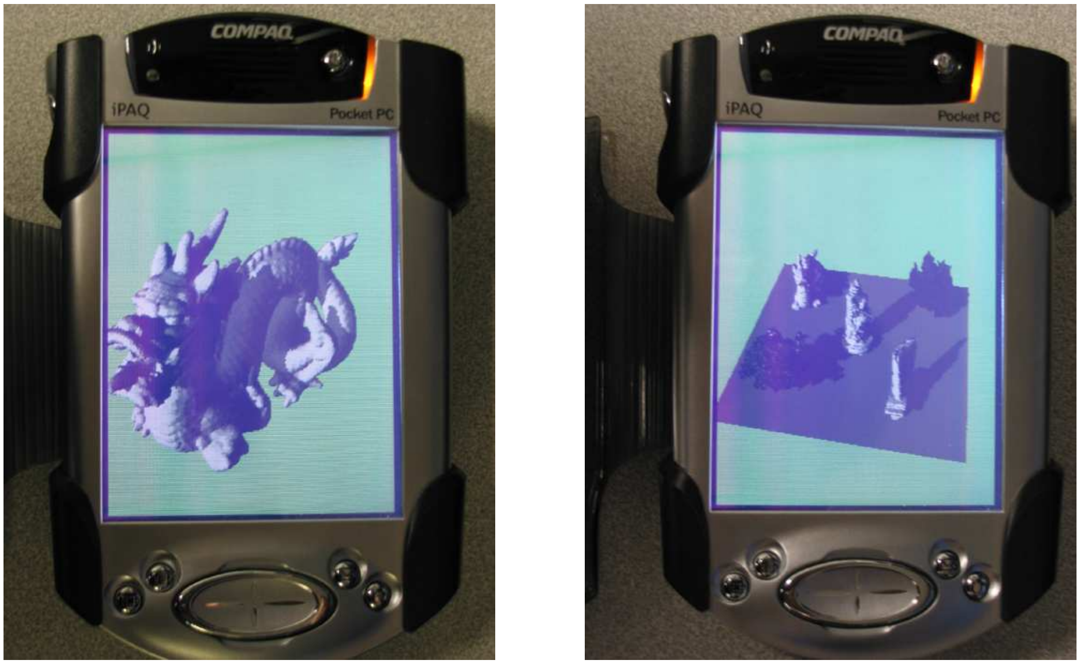

Figure 5. Images of flexible point-based rendering.

The system starts with a small number of images, without any restrictions on camera pose. After the calibration phase of the camera, and the interactive definition of an approximate volume containing the object to be reconstructed, the system generates a large number of points inside this volume. A set of known or improved techniques are them combined to test the quality of the points generated, for example colour or depth comparisons. As mentioned above, interactivity and user intervetion are central to our approach; the user can designate regions in which the system should add more points, regions where a mask can be applied to guide point generation, or regions where a curve filling approach (for example tree branches) is more appropriate. At all intermediate steps, and of course once the model is complete, the system can generate images from any viewpoint using the point-based representation and a view-dependent texturing approach. The results of this research (see Fig. 6), are very satisfactory for objects which are hard to reconstruct using other methods because of small, irregular details.

This work was presented at the international conference Graphics Interface 2003 [8] in Canada.

6.1.4. Improving detail visualisation of image-based reconstructed scenes

Participants: Marie-Claude Frasson, George Drettakis.

This project aims at helping artists build depth maps typically used to add details to their scenes. Most of the time, artists have to build their displacement or depth maps by hand; this can be tedious when the textures are large. We have developed a method to automate this process. To add detail to a texture, the artist only has to supply a few detail extracts (typically 2 or 3), along with the associated hand-drawn depth. The system will use these patches to produce a complete depth map for the entire input texture according to the information in the patches. The details on these objects are then generated according to the depth maps and junctions between two displaced textures are handled correctly to produce a seamless blend (see Fig. 7).

6.1.5. Interactive audio rendering

Participants: Emmanuel Gallo, Nicolas Tsingos, Charles-Félix Chabert, George Drettakis.

Figure 6. Results of Interactive Point-Based Modelling. Top, the input images used for reconstruction, and bottom, the resulting point-based model (with and without texture reprojection).

Figure 7. Result of the method on an image-based reconstructed block. The junction has been generated and corrected so that the two displaced textures match seamlessly.

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 began a series of perceptive tests to validate our approaches.

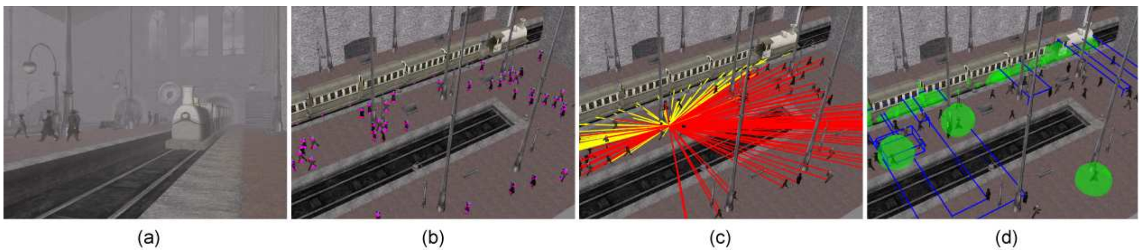

Figure 8. (a) An application of the perceptual rendering pipeline to a complex train-station environment. (b) Each pedestrian acts as two sound sources (voice and footsteps). Each wheel of the train is also modeled as a point sound source to get the proper spatial rendering for this extended source. 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.

In the context of Charles-Felix Chabert's master work [14], we also explored modeling and rendering of extended sound sources for multi-media applications. Rendering spatially extended sound sources from a single monophonic (recorded or synthesized) signal has many applications in the sound post-production/video-game industry. We developed a novel approach to combine a small number of decorrelated copies of the original signal in order to preserve its timbre and quality in the synthesized "spatially extended" version. The spatially extended version is generated by rendering a distinct number of point sources in 3D space using standard spatialization stategies. Each of these sources is associated with a decorrelated copy of the original signal. Decorrelation is done in a pre-processing stage hence reducing the amount of required on-line calculation [14].

Finally, a state-of-the-art paper on audio rendering techniques for virtual reality applications, co-authored by Tom Funkhouser at Princeton and Jean-Marc Jot at Creative Labs, has also been accepted for publication in Presence, the reference journal of the virtual reality community [2].

6.1.6. Perceptive multi-modal rendering

Participants: Manuel Asselot, Nicolas Tsingos.

The goal of this research is to efficiently render virtual scenes taking advantage of the human perception of the environment [12]. 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.

6.2. High-Quality Rendering

6.2.1. High quality audio rendering

Participant: Nicolas Tsingos.

Collaboration with Bell Laboratories and Princeton University

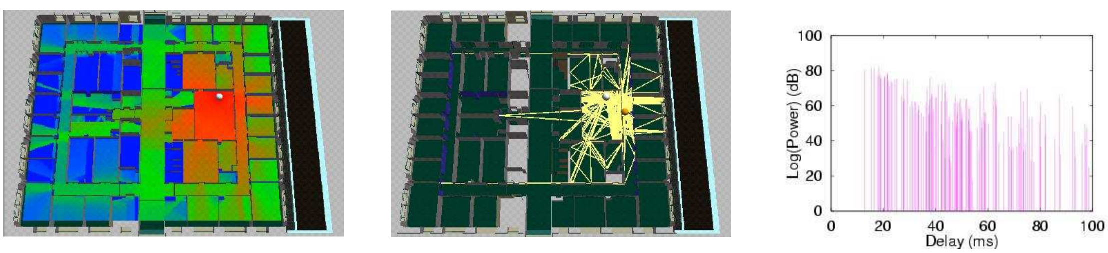

Figure 9. Sound level map, propagation paths and corresponding impulse response generated with our beam-tracing engine in an architectural environment.

A paper compiling all the research work, started at Bell Laboratories, concerning efficient high-quality simulation of sound propagation has been published in JASA (Journal of the Acoustical Society of America), the reference journal in the field. Our techniques, based on beam tracing, allow for efficient and aliasing free construction of reflection and diffraction paths in 3D architectural environments.

6.3. Virtual Environments with Sound

6.3.1. View Dependent Layered Textures

Participants: Alex Reche, George Drettakis.

Capturing and rendering of real scenes in an immersive virtual environment is still a challenging task. We have developed a new technique to generate view dependent textures using image layers that produces high quality images. In addition the new approach reduces texture memory consumption. The algorithm has two steps: the first part builds the layered textures and the second part uses these textures to render the scene.

To generate the layered textures the algorithm takes high quality images as input, the corresponding calibrated cameras and the 3D model of the objects in the photographs. Then we use the camera information to sort the geometry by visibility, grouping the objects with the same level of visibility into *visibility layers*.

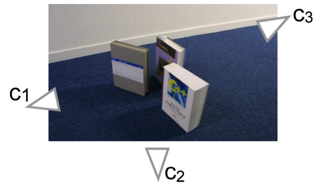

Figure 10. Scene Configuration

Figure 11. Generated layers for C1 before and after image editing

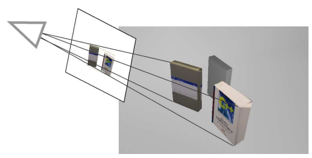

Figure 12. Projection of the first layer

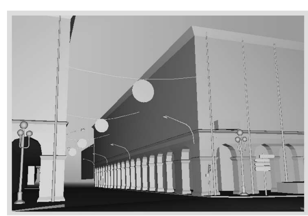

Figure 13. 3D Model of a more complex scene

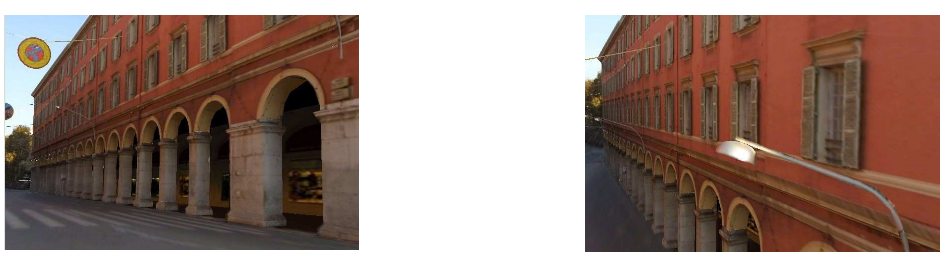

Figure 14. Results with the complex scene

These layers are projected into the image space to generate *image Layers*. The image layers are then split into smallest layers using a criteria based on OpenGL, to optimise the texture memory used. At the end of the process all the images have been segmented into a sequence of image layers, each one of them corresponding to a level of visibility.

For rendering, we extended the traditional projective texture approach to be able to manage layered textures. The main difference with standard projective textures is that the matrix projection for a layer is an asymmetric projective matrix and we need to keep the relation between each object and the associated layer into the image. We used this new projective texture technique in a view dependent projective texture algorithm. Compared with precedent algorithms, our new technique does not require the subdivision of geometry (imposed by the visibility problems during the texture projection in previous techniques) and allows the artist to edit the texture using standard tools such as clone brushing to fill holes due to occlusion. This step was impossible using previous approaches; our method gives complete control of the final result to the artists.

A short version of this research was published in the Proceedings of the Pacific Graphics Conference 2003 [10], in Canmore (Canada), and presented as a poster and it will be presented at the AFIG conference in Paris for the French Graphics Community [9]. A longer version of it is going to be Technical Report[16].

6.3.2. *Realistic Virtual Environments for Urban Planning*

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 started an official collaboration with the architects of the Tramway project of the city of Nice using our tools. An initial presentation of these enhancements in the context of the CREATE project was presented as a poster at the international symposium on Mixed and Augmented Reality (ISMAR) 2003 [7]. A specific description of the VE enhancements developed by REVES will also be presented as a poster at IEEE VR2004.

6.4. Interaction and Visual Paradigms for Virtual and Augmented Environments

6.4.1. *Realistic and Expressive (Non-Photorealistic) Rendering for Virtual Environments and Archeology*

Participants: Roman Bayon, Matthieu Cunzi, George Drettakis, Maria Roussou, Alexandre Olivier-Mangon.

Research in collaboration with ARTIS and MIT, in the context of ARC ARCHEOS.

In the context of the ARCHEOS national research action (see 8.2.1), we have investigated questions related to expressive (better known as *non-photorealistic*) renderings in the context of archeological applications. The research developed this year concentrated on two main axes, that of rendering styles for walkthroughs and the use of expressive rendering in semi-immersive or immersive systems.

Matthieu Cunzi, in collaboration with his DEA stage supervisor J. Thollot and other colleagues at ARTIS and MIT, developed a novel approach for rendering the effect of canvas granularity for expressive rendering. The solution consisted of developing a solution for infinite zoom based on repetitive textures, and appropriate texutre warping for rotation. This work was published at the international conference Graphics Interface 2003 in Canada [5].

In the context of his 6-month internship Roman Bayon developed a number of solutions for expressive rendering on our workbench platform. He investigated a number of software solutions, including the use of our common API. The final solution adopted involved the use of standard tools i.e., Performer and CAVElib.

Figure 15. Above, partially textured reconstruction of Garibladi square, developed for the CREATE project, and below an image of the CAVElib simulator showing our virtual environment enhancements (point-based rendering for vegetation, shadows and view-dependent texture mapping), integrated into a VE platform. The simulation also include 3D sound.

In his internship, the idea of transposing the canvas solution to stereo was also partially investigated; this issue proved to be more complex than initially expected, and further research is required. An internship report describing this work is available [13].

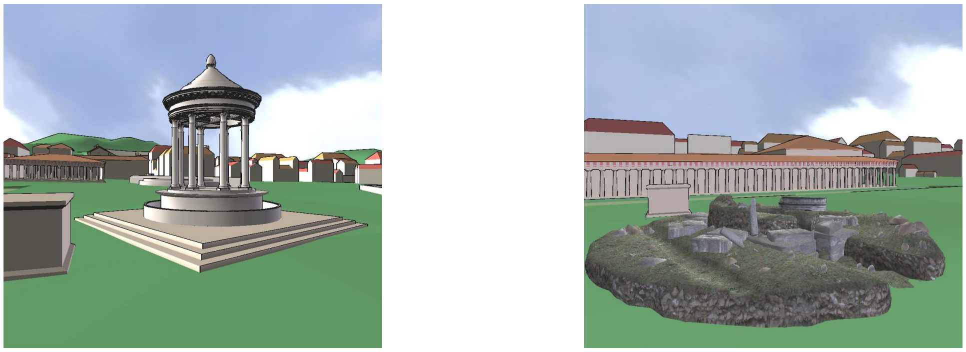

Figure 16. Left, a pen-and-ink image of the Tholo (Piérart hypothesis), and right, a mixed image of realistic rendering for the actual stones on the site today and the NPR rendering of the hypothesis of the Agora in hellenistic Argos (450 AD).

In collaboration with M. Roussou, and in the context of both the EU CREATE project and ARC ARCHEOS we presented a position paper on the relative advantages and disadvantages of realistic vs. non-photorealistic renderings for cultural heritage applications. This paper was presented at the first EG workshop on Graphics and Heritage, which was common event with VAST 2003 [11]. The presentation resulted in an interesting discussion concerning expressive rendering, validating the idea that expressive rendering for VE's is a very promising research direction.

7. Contracts and Grants with Industry

7.1. CSTB

Participants: George Drettakis, Alex Reche.

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.

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 ARCHEOS 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 Superieure, 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.

7.5. Patent Application

Participants: Nicolas Tsingos, Emmanuel Gallo, George Drettakis.

We are in the process of filing 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).

8. Other Grants and Activities

8.1. Regional/Local Projects

8.1.1. Collaboration with CSTB Sophia-Antipolis

Participants: George Drettakis, Nicolas Tsingos, Alex Reche, Gaël Braconier.

We collaborate with CSTB in the context of the European IST project CREATE. We are also conducting tests for global illumination simulation for architectural environments in the context of the ECLAIRES project, developed by G. Braconier. Such tests allow us to validate our lighting algorithms.

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.

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 8.1-surround sound restitution system. The final system should be operational by early 2004.

8.2. National Projects

8.2.1. Cooperative Research Action ARCHEOS

Participants: George Drettakis, Matthieu Cunzi, Alexandre Olivier-Mangon.

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

REVES is the coordinator of the national research action ARCHEOS, whose initial goal was to investigate the use of expressive (non-photorealistic) rendering for archeology. The partners of the action are ARTIS in Grenoble, the Homeric Studies group ERGA in Grenoble, the architecture School of Lyon, the Ecole Normale Supérieure in Paris and the Foundation of the Hellenic World in Greece. One of the main goals of this project is to initiate a dialogue between the computer graphics and VE researchers on the one hand, and the historians and archeologists on the other.

We have concentrated our work on a specific archeological example, which is the Agora of Argos, an ancient city in Greece. One monument of particular interest is the Tholos (see Fig. 17 and 16), where there is some doubt of its exact utility. Different archeological hypotheses have been proposed concerning its form and its utility.

In its second year, the ARC ARCHEOS funded the internship of Roman Bayon (see also Section 6.4.1) as well as the summer work of Pascal Barla at ARTIS in Grenoble after his DEA.

Two main meetings took place this year. The first meeting, in Grenoble in May, was particularly interesting, since two of the main archeologists who have excavated the Argos site were present (A. Pariente and M. Pierart). The interest of the archeologists was evident, both for the general aspect of VE's and for expressive renderings. The second meeting was a research brainstorming on Computer Graphics aspect of the work in July in Sophia-Antipolis.

The work of ARCHEOS was also presented at two archeological research meetings, one in Boston as a poster and one in Athens at the Ecole Française d'Athènes, which is responsible for the Argos site. In both cases, the archeologists present expressed great interest in the work, and were particularly interested to see its continuation.

The closing event of the ARC ARCHEOS is planned for early December; the morning will include a presentation of the research results and demos (see Fig. 17), and the afternoon will include a round-table of archeologists, computer graphics and virtual environment researchers.

8.2.2. ACI MD SHOW

Participant: George Drettakis, Florent Duguet.

A national project, coordinated by ARTIS in Grenoble, started this autumn. Our paticipation in this project will be on point-based or alternative rendering of very large data sets, such as scans of statues or alternative representations of trees. The other participants of this project are IPARLA in Bordeaux and ISA in Nancy.

8.2.3. RNTL project OPERA: PErceptual Optimizations for Audio Rendering

Participant: Nicolas Tsingos.

REVES is also coordinator of a new 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 will start early 2004.

Figure 17. Left, the ARCHEOS demo prepared for the archeologists. This is a snapshot of the VE CAVElib simulator running on the workbench. We see here the Marchetti hypothesis of the Tholos. Right, the stones actually existing on the site, placed in the positions of the hypothetical reconstruction of the monument.

8.3. Visiting Researchers

This year the following researchers visited the group Ken Perlin (New York University), Mat Kaplan (University of Utah), Alexandre Meyer (University of London) and Xavier Decoret (ARTIS, Grenoble). Several other researchers visited the group in the context of ARCHEOS and CREATE.

8.4. European Projects

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

8.4.1. CREATE

As metioned previously in various sections, a significant part of our activities this 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. In March, we hosted the 6-monthly project meeting in Sophia-Antipolis. The other partners of the project are UCL (UK), FHW (GR), REALVIZ, CSTB, PERCRO (I) and UCY (CY).

The capture and display pipeline starts using improved versions of the REALVIZ Stitcher and ImageModeler tools. We developed an improved capture pipeline which starts with the creation of panoramic images using Stitcher at various locations in the environment. These are then passed to ImageModeler for calibration, and then are the basis for Image-Based modeling. We then use the VE enhancement tools to display the environments created in an immersive or semi-immersive context.

Two application demonstrators are under development as part of the contract. The first concerns the site of ancient Messene, which is being built by the Foundation of the Hellenic World. Alex Reche participated in the original capture of the Messene site in Greece this May, which involved 3 days on site and another 2 days at the FHW to calibrate the model. The goal is to build a demonstrator which will be used by children 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 have started this autumn. After a series of demonstrations to officials and engineers of the real project (see 9.4), we have established an official convention of collaboration with the "Mission Tramway" which is the organisation in charge of the implementation of the project.

Initially we created a pilot captured model of Place Massena, which was enhanced with point-based vegetation, view-dependent texturing, perspective shadows and Initially a series of photographs were taken from various viewpoints. Then 3 photographs were selected to represent the point of view and we extracted textures from them. Virtual elements, such as the Tramway or the trees, were also integrated into the scene (see Fig. 15 and 18(a)). The system has been successfully ported to the CSTB RealityCenter (see Fig. 18(b)).

Figure 18. Left, the CREATE demo running on the INRIA workbench. Right the CREATE demo running in the CSTB RealityCenter.

After establishing contact with the real Tramway project and discussing choices in detail with the engineers and the architects, we decided to develop Place Garibaldi as the final demostrator of the urban planning CREATE scenario. A series of 5 panoramas were taken to capture and represent different points of view and to build a complete 3D model. This project is one of the most detailed and most complex to be developed using the REALVIZ tools, which have been successively adapted and extended to meet the needs of this project. The initial capture and modelling of this site is complete (see Fig. 15).

Current discussions with the architects are centered on the different choices for the development of the square, their presentation and the alternatives which are possible. We expect this activity to result in interesting applied VE research for a real application.

8.5. Bilateral Collaborations

8.5.1. France-Québec

Our collaboration with the University of Montréal and Pierre Poulin continued successfully this year with the presentation of our common work at Graphics Interface 2003 [8].

8.5.2. France-Greece (Hellas)

Our work in common with the FHW has continued in the context of the CREATE project, and with the presentation of two publications (ISMAR [7] and VAST [11]) this year.

8.5.3. France-Germany

Our work with M. Stamminger at the university of Erlangen has continued with the presentation of the "wider public" presentation of the perspective shadow maps in Game Programming Gems IV [3].

8.5.4. France-Spain

We have re-established contact with the research group of Girona, with the visit of Ignacio Martin from this group for a month and a half. We intend to continue this collaboration next year.

8.5.5. France-United States of America

During the summer, Florent Duguet visited Fredo Durand in the Graphics lab at MIT funded by the MIT-France project for two months. A new research project has been started during this visit putting together several techniques from both parties. The focus of the project is mesh smoothing and surface reconstruction. The collaboration of MIT and the REVES team is continuing with a publication project.

Nicolas Tsingos is actively collaborating with Bell Laboratories, Princeton University and Creative Labs.

We currently have active research contacts with Columbia (NY) and Cornell (NY).

9. Dissemination

9.1. Participation in the Community

9.1.1. Programme Committees

G. Drettakis was a member of the programme committees for the Eurographics Symposium on Rendering 2003, and VAST2003/1st Eurographics workshop on Graphics and Heritage. G. Drettakis and N. Tsingos reviewed papers for the following conferences: SIGGRAPH, Eurographics, ACM SIGGRAPH Symposium on Interactive 3D Graphics, as well as several journals (JASA, IEEE CG & A, IEEE TVCG etc.).

9.1.2. Thesis Committees

G. Drettakis was an external examiner for the thesis of P. Reuter at the University of Bordeaux.

9.1.3. Web server

Participants: George Drettakis, Maria Roussou.

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

The web server was completely redesigned and put online in Spring this year. 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 (École des mines) in January 2003 (15 hours). He also taught three hours of the Medical Robotics course at the Graphics DEA at the UNSA and one hour in Computer Graphics for second-year engineering students of the ESSI. Alex Reche taught 12 hours of the Image-Based reconstruction and OpenGL course and laboratories for the second year students of the LPMI at the Nice IUT. Florent Duguet taught three hours of the Computer Graphics course and four hours of associated laboratories for third-year engineering students of the ESSI. He is also a regular 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. Marie-Claude Frasson gave two hours of a Ray-Tracing course to second-year engineering students at the ESSI, plus eight hours of associated laboratories. She also supervised the 4 hours OpenGL laboratories of the second-year LPMI students of the IUT of Nice. Nicolas Tsingos was responsible for the “advanced 3D rendering" module of second-year LPMI of the IUT of Nice (20 hours).

9.2.2. Ongoing PhDs

Marie-Claude Frasson and Alex Reche registered at the University of Nice for their third PhD year, while Florent Duguet is beginning is second year, also at the University of Nice Sophia-Antipolis (UNSA).

9.3. Participation at conferences

9.3.1. Presentations at Conferences

Alex Reche presented his short paper as a poster at the Pacific Graphics 2003 conference in Canmore, Canada and its french version at the Journées de l'AFIG 2003 in Paris. Matthieu Cunzi presented his paper at Graphics Interface, and Pierre Poulin presented the paper on point reconstruction. The ISMAR 2003 poster was presented by Celine Loscos of UCL and the VAST 2003 paper was presented by Maria Roussou of FHW/UCL. Almost all members of the team held a booth and participated at the Imagina 2003 conference in Monaco.

9.3.2. Participation at Conferences and Workshops

In addition to the above-mentioned conference presentations, the researchers, Ph.D. students and engineers went to the 2003 Eurographics Symposium on Rendering. G. Drettakis, F. Duguet and N. Tsingos were at the San Diego SIGGRAPH 2003 conference. M-C. Frasson and Alex Reche went to the Eurographics 2003 conference in Granada.

9.4. Demonstrations and Press

Since the workbench has become operational, our group is often solicited for demonstrations.

In the context of CREATE, we established contact with the real Tramway project, by first presented the prototype CREATE demonstrator (Massena) to M. Bonis chief engineer of the Tram Project and two public relations officers in April. We then presented the work to M. Sanz municipal councillor in charge of Transportation of Nice and M. Cherifi PR delegate for the CANCA in May, which was a necessary step in establishing contact.

More recently, we have demonstrated the work to the the architect company MOER (June and August), which is in charge of implementing the Tramway project.

We also presented research demonstrator on the workbench to the companies OKTAL SA and VIRNET (April/May), INRIA directors of national units (May), Toyota (July), the Mayor of Antibes and municipal delegates (August 2003) and the vice-president of the Regional local government, as well as other local officials (December 2003). We also presented the workbench during "La fete de la Science" in October (150 people over two days).

We also contributed to the INRIA press release on virtual reality, http://www.inria.fr/inria/rapportannuel/ra2002/virtualreality.en.pdf the INRIA 2002 annual report http://www.inria.fr/inria/rapportannuel/ran.en.html and the booklet for the 20 years of the INRIA Sophia-Antipolis site.

10. Bibliography

Articles in referred journals and book chapters

[1] T. FUNKHOUSER, N. TSINGOS, I. CARLBOM, G. ELKO, M. SONDHI, J. WEST, G. PINGALI, P. MIN, A. NGAN. *A Beam Tracing Method for Interactive Architectural Acoustics.* in « The Journal of the Acoustical Society of America (JASA) », 2003, http://www-sop.inria.fr/reves/publications/data/2003/FTCESWPMN03.

[2] T. FUNKHOUSER, N. TSINGOS, J.-M. JOT. *Survey of Methods for Modeling Sound Propagation in Interactive Virtual Environment Systems.* in « Presence and Teleoperation », 2003, http://www-sop.inria.fr/reves/publications/data/2003/FTJ03.

[3] M. STAMMINGER, G. DRETTAKIS, C. DACHSBACHER. *Perspective Shadow Maps.* Charles River Media, 2003, chapter ??, http://www-sop.inria.fr/reves/publications/data/2003/SDD03.

[4] N. TSINGOS, E. GALLO, G. DRETTAKIS. *Breaking the 64 spatialized sources barrier.* in « Gamasutra Audio Resource Guide 2003 », May, 2003, http://www-sop.inria.fr/reves/publications/data/2003/TGD03b.

Publications in Conferences and Workshops

[5] M. CUNZI, J. THOLLOT, S. PARIS, G. DEBUNNE, J.-D. GASCUEL, F. DURAND. *Dynamic Canvas for Immersive Non-Photorealistic Walkthroughs.* in « Proc. Graphics Interface 2003 », 2003, http://www-sop.inria.fr/reves/publications/data/2003/CTPDGD03.

[6] 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.

[7] C. LOSCOS, H. R. WIDENFELD, M. ROUSSOU, A. MEYER, F. TECCHIA, E. GALLO, G. DRETTAKIS, A. RECHE, N. TSINGOS, Y. CHRYSANTHOU, L. ROBERT, M. BERGAMASCO, A. DETTORI, S. SOUBRA. *The CREATE Project: Mixed Reality for Design, Education, and Cultural Heritage with a Constructivist Approach.* in « ISMAR 03, The Second IEEE and ACM International Symposium on Mixed and Augmented Reality », The National Center of Sciences, Tokyo, Japan, IEEE/ACM, October, 2003, http://www-sop.inria.fr/reves/publications/data/2003/LRRMTGDRTCRBDS0.

[8] P. POULIN, M. STAMMINGER, F. DURANLEAU, M.-C. FRASSON, G. DRETTAKIS. *Interactive Point-based Modeling of Complex Objects from Images.* in « Proc. Graphics Interface 2003 », June, 2003, http://www-sop.inria.fr/reves/publications/data/2003/PSDFD03.

[9] A. RECHE, G. DRETTAKIS. *Textures projectives à calques dépendantes du point de vue.* in « Actes de AFIG 2003 », AFIG, December, 2003, http://www-sop.inria.fr/reves/publications/data/2003/RD03d.

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

[11] M. ROUSSOU, G. DRETTAKIS. *Photorealism and Non-Photorealism in Virtual Heritage Representation.* in « VAST 2003 and First Eurographics Workshop on Graphics and Cultural Heritage », Eurographics, A. CHALMERS, D. ARNOLD, F. NICCOLUCCI, editors, November, 2003, http://www-sop.inria.fr/reves/publications/data/2003/RD03c.

Internal Reports

[12] M. ASSELOT. *Rendu perceptif multi-modal pour les environnements virtuels immersifs.* Technical report, REVES/INRIA, November, 2003.

[13] R. BAYON. *Développement de techniques d'affichage non-photoréalistes efficaces pour un environnement de réalité virtuelle appliquée à l'archéologie.* Technical report, REVES/INRIA, August, 2003, http://www-sop.inria.fr/reves/publications/data/2003/Bay03.

[14] C.-F. CHABERT. *Rendu de sources sonores étendues pour le multimédia.* Technical report, REVES/INRIA, September, 2003.

[15] F. DUGUET, G. DRETTAKIS. *Flexible Point-Based Rendering on Mobile Devices.* Technical report, number RR-4833, INRIA, REVES/INRIA Sophia-Antipolis, May, 2003, http://www.inria.fr/rrrt/rr-4833.html.

[16] A. RECHE, G. DRETTAKIS. *View-Dependent Layered Projective Texture Maps.* Technical report, INRIA, Sophia-Antipolis, July, 2003, http://www.inria.fr/rrrt/rr-5016.html.

[17] N. TSINGOS, E. GALLO, G. DRETTAKIS. *Perceptual Audio Rendering of Complex Virtual Environments.* Technical report, number RR-4734, INRIA, REVES/INRIA Sophia-Antipolis, February, 2003, http://www.inria.fr/rrrt/rr-4734.html.

Bibliography in notes

[18] 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.

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

[20] 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.

[21] 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, pages 245–257, Juin, 1996.

[22] 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.

[23] F. DURAND, G. DRETTAKIS, C. PUECH. *Fast and Accurate Hierarchical Radiosity Using Global Visibility.* in « ACM Transactions on Graphics », volume 18, Avril, 1999, pages 128–170.

[24] 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.

[25] 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.

[26] C. LOSCOS, G. DRETTAKIS, L. ROBERT. *Interactive Virtual Relighting of Real Scenes.* in « IEEE Transaction of Visualization and Computer Graphics », volume 6, Octobre, 2000, pages 289–305.

[27] 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, pages 329–340, Juin, 1999.

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

[29] 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.

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

[31] 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.

[32] 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.