Activity Report 2011

Project-Team REVES

Rendering and virtual environments with sound

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Interaction and Visualization

Table of contents

Project-Team REVES

Keywords: Audio, Computer Graphics, Interaction, Visualization, Virtual Reality

Beginning of the Team: 01/07/2002.

1. Members

Research Scientists
George Drettakis [Team Leader, Senior Researcher (DR) Inria, HdR]
Adrien Bousseau [Junior Researcher (CR) Inria]

External Collaborator
Ares Lagae

Engineer
Adrien David [Engineer ADT, 2 year contract INRIA-funded]

PhD Students
Emmanuelle Chapoulie [BDO scholarship]
Gaurav Chaurasia [MESR scholarship]
Sylvain Duchêne [VERVE European project]
Pierre-Yves Laffont [INRIA CORDIS scholarship]

Post-Doctoral Fellows
Carles Bosch [ERCIM-funded, until April]
Peter Vangorp [INRIA-funded, until November]
Charles Verron [INRIA AAR, since September]

Administrative Assistant
Sophie Honnorat [Administrative Assistant (AI) Inria, part-time]

Others
Aurélien Bouteille [Intern, Polytech'Nice-Sophia, May-August]
Evanthia Dimara [Intern, Unviersity of Athens, ARC NIEVE, October-December]
Marion Ferrara [Intern, Polytech'Nice-SOphia, ARC NIEVE, June-August]
Sandra Scognamillo [Intern, Polytech'Nice-SOphia, ARC NIEVE, June-August]
Yunfei Cloud Shao [Visiting PhD, University of Toronto, November-December]

2. Overall Objectives

2.1. General Presentation

Images, often accompanied by sound effects, have become increasingly present in our everyday lives; this has resulted in greater needs for content creation. Despite the fact that many traditional means exist, such as photography, artistic graphic design, audio mixing, they typically still remain the reserve of the expert, and require significant investment in time and expertise. Our main interest is computer image and sound synthesis, with an emphasis on automated methods. Our main goals include the simplification of the tasks required for the production of sound and images, as well as the development of new techniques for their generation.

The application domain is vast. It ranges from audiovisual production, which typically requires long, offline computation to obtain high quality results, all the way to real-time applications such as computer games or virtual reality, for which the main consideration is to guarantee 60 frames per second frame rates, or, in general the reduction of latency to user reaction. The process of generation of images and sound, generally called *rendering* is our primary interest; our second main interest are virtual environments (VE's) as well as augmented (AE's) or mixed environments (ME's), that is scenes containing both real objects (often digitized) as well as purely synthetic objects. We are interested in both the generation and the interaction with these environments. We use the term virtual environments for scenes with a certain degree of interactivity, potentially in a semi-immersive (stereo and tracking, workbench) or immersive (CAVE, RealityCenter) context.

2.2. Highlights of the Year

Adrien Bousseau won the Eurographisc Ph.D. award in 2011 (three awards are given each year). He also co-authored the paper [23], which won the "Best Paper Award" at the ACM Symposium on Interactive 3D Graphics and Games. Finally, A. Bousseau is a co-author of the paper on Diffusion Curves which was selected for publication in the "Communications of the ACM".

Marcio Cabral won the best PhD thesis in CG Award at SIBGRAPI 2011 (Brasilian conference of Computer Graphics), and Pierre-Yves Laffont won the best paper award at AFIG 2011 for his papers [21]. Charles Verron received the 2011 Rocard Price from the Société Francaise d'Acoustique (SFA) in November 2011.

rendering,image Rendering,sound Rendering,plausible Rendering,high-quality Rendering

3. Scientific Foundations

3.1. Rendering

We consider plausible rendering to be a first promising research direction, both for images and for sound. Recent developments, such as point rendering, image-based modeling and rendering, and work on the simulation of aging indicate high potential for the development of techniques which render *plausible* rather than extremely accurate images. In particular, such approaches can result in more efficient renderings of very complex scenes (such as outdoors environments). This is true both for visual (image) and sound rendering. In the case of images, such techniques are naturally related to image- or point-based methods. It is important to note that these models are becoming more and more important in the context of network or heterogeneous rendering, where the traditional polygon-based approach is rapidly reaching its limits. Another research direction of interest is realistic rendering using simulation methods, both for images and sound. In some cases, research in these domains has reached a certain level of maturity, for example in the case of lighting and global illumination. For some of these domains, we investigate the possibility of technology transfer with appropriate partners. Nonetheless, certain aspects of these research domains, such as visibility or high-quality sound still have numerous and interesting remaining research challenges.

3.1.1. Plausible Rendering

3.1.1.1. Alternative representations for complex geometry

The key elements required to obtain visually rich simulations, are sufficient geometric detail, textures and lighting effects. A variety of algorithms exist to achieve these goals, for example displacement mapping, that is the displacement of a surface by a function or a series of functions, which are often generated stochastically. With such methods, it is possible to generate convincing representations of terrains or mountains, or of non-smooth objects such as rocks. Traditional approaches used to represent such objects require a very large number of polygons, resulting in slow rendering rates. Much more efficient rendering can be achieved by using point or image based rendering, where the number of elements used for display is view- or image resolution-dependent, resulting in a significant decrease in geometric complexity. Such approaches have very high potential. For example, if all object can be rendered by points, it could be possible to achieve much higher quality local illumination or shading, using more sophisticated and expensive algorithms, since geometric complexity will be reduced. Such novel techniques could lead to a complete replacement of polygon-based rendering for complex scenes. A number of significant technical challenges remain to achieve such a goal, including sampling techniques which adapt well to shading and shadowing algorithms, the development of algorithms and data structures which are both fast and compact, and which can allow interactive or real-time rendering. The type of rendering platforms used, varying from the high-performance graphics workstation all the way to the PDA or mobile phone, is an additional consideration in the development of these structures and algorithms. Such approaches are clearly a suitable choice for network rendering, for games or the modelling of certain natural object or phenomena (such as vegetation, e.g. Figure 1, or clouds). Other representations merit further research, such as image or video based rendering algorithms, or structures/algorithms such as the "render cache" [35], which we have developed in the past, or even volumetric methods. We will take into account considerations related to heterogeneous rendering platforms, network rendering, and the appropriate choices depending on bandwith or application. Point- or image-based representations can also lead to novel solutions for capturing and representing real objects. By combining real images, sampling techniques and borrowing techniques from other domains (e.g., computer vision, volumetric imaging, tomography etc.) we hope to develop representations of complex natural objects which will allow rapid rendering. Such approaches are closely related to texture synthesis and image-based modeling. We believe that such methods will not replace 3D (laser or range-finger) scans, but could be complementary, and represent a simpler and lower cost alternative for certain applications (architecture, archeology etc.). We are also investigating methods for adding "natural appearance" to synthetic objects. Such approaches include *weathering* or *aging* techniques, based on physical simulations [25], but also simpler methods such as accessibility maps [32]. The approaches we intend to investigate will attempt to both combine and simplify existing techniques, or develop novel approaches founded on generative models based on observation of the real world.

3.1.1.2. Plausible audio rendering

Similar to image rendering, plausible approaches can be designed for audio rendering. For instance, the complexity of rendering high order reflections of sound waves makes current geometrical approaches inappropriate. However, such high order reflections drive our auditory perception of "reverberation" in a virtual environment and are thus a key aspect of a plausible audio rendering approach. In complex environments, such as cities, with a high geometrical complexity, hundreds or thousands of pedestrians and vehicles, the acoustic field is extremely rich. Here again, current geometrical approaches cannot be used due to the overwhelming number of sound sources to process. We study approaches for statistical modeling of sound scenes to efficiently deal with such complex environments. We also study perceptual approaches to audio rendering which can result in high efficiency rendering algorithms while preserving visual-auditory consistency if required.

3.1.2. High Quality Rendering Using Simulation

3.1.2.1. Non-diffuse lighting

A large body of global illumination research has concentrated on finite element methods for the simulation of the diffuse component and stochastic methods for the non-diffuse component. Mesh-based finite element approaches have a number of limitations, in terms of finding appropriate meshing strategies and form-factor calculations. Error analysis methodologies for finite element and stochastic methods have been very different in the past, and a unified approach would clearly be interesting. Efficient rendering, which is a major advantage of finite element approaches, remains an overall goal for all general global illumination research. For certain cases, stochastic methods can be efficient for all types of light transfers, in particular if we require a view-dependent solution. We are also interested both in *pure* stochastic methods, which do not use finite element techniques. Interesting future directions include filtering for improvement of final image quality as well as beam tracing type approaches [33] 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 [29], [28], [27] typically have

Figure 1. Plausible rendering of an outdoors scene containing points, lines and polygons [24], 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.

robustness or memory consumption problems which make them difficult to apply to scenes of realistic size. Our solutions to date are based on general and flexible formalisms which describe all visibility event in terms of generators (vertices and edges); this approach has been published in the past [26]. Lazy evaluation, as well as hierarchical solutions, are clearly interesting avenues of research, although are probably quite application dependent.

3.1.2.3. Radiosity

For purely diffuse scenes, the radiosity algorithm remains one of the most well-adapted solutions. This area has reached a certain level of maturity, and many of the remaining problems are more technology-transfer oriented We are interested in interactive or real-time renderings of global illumination simulations for very complex scenes, the "cleanup" of input data, the use of application-dependent semantic information and mixed representations and their management. Hierarchical radiosity can also be applied to sound, and the ideas used in clustering methods for lighting can be applied to sound.

3.1.2.4. High-quality audio rendering

Our research on high quality audio rendering is focused on developing efficient algorithms for simulations of geometrical acoustics. It is necessary to develop techniques that can deal with complex scenes, introducing efficient algorithms and data structures (for instance, beam-trees [30] [33]), 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 [31] has shown that geometrical approaches can lead to high quality modeling of sound reflection and diffraction in a virtual environment (Figure 2). We will pursue this research further, for instance by dealing with more complex geometry (e.g., concert hall, entire building floors).

Finally, several signal processing issues remain in order to properly and efficiently restitute a 3D soundfield to the ears of the listener over a variety of systems (headphones, speakers). We would like to develop an open and general-purpose API for audio rendering applications. We already completed a preliminary version of a software library: AURELI [34].

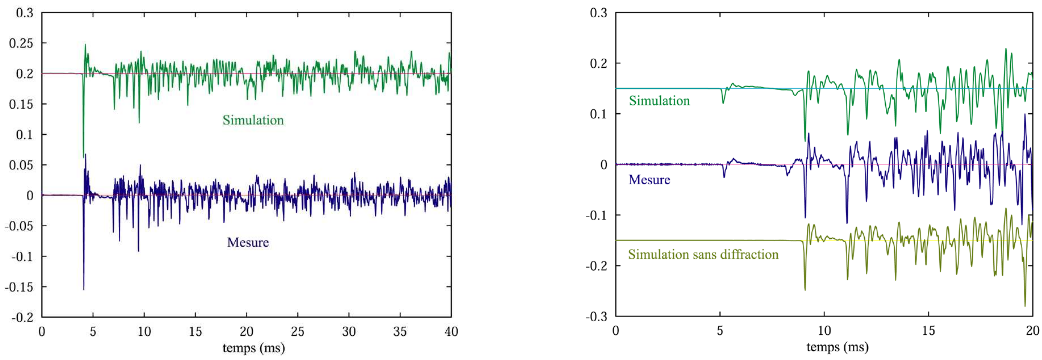

Figure 2. A comparison between a measurement (left) of the sound pressure in a given location of the "Bell Labs Box", a simple test environment built at Bell Laboratories, and a high-quality simulation based on a beam-tracing engine (right). Simulations include effects of reflections off the walls and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

4. Software

4.1. RID: Rich Intrinsic Decomposer

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We developed a software platform to perform rich intrinsic decomposition methods from photographs of outdoor scenes, as described in [21] and in an article currently submitted for publication. It includes main scripts and functions in Matlab for treatment of the input data, interfaces to software for multi-view reconstruction (Bundler, PMVS) and meshing from point clouds (method developed by Julie Digne, a postdoc in the Geometrica team). We then interface software for image matting using the Matting Laplacian, and User-Assisted Intrinsic Images. The system also includes an interface with Adobe Photoshop, for visualization and demonstration of our results in end-user image editing software. The method performs the computation of sun, sky and indirect lighting received at 3D points of an automatically reconstructed scene, using a modified version of the PBRT stochastic raytracer. Finally, there is a scene calibration module and an OpenGL viewer.

4.2. Imerse: Inria Multi-Environment Realistic Simulation Engine

Participants: Adrien David, George Drettakis.

In the context of the ADT Interact3D and the ARC NIEVE, we developed Imerse, a middleware to be used as a VR engine, helping in the implementation of realistic simulations for immersive installations. Imerse provides a wrapper to OSG's (OpenSceneGraph) deep scene graph and its traversals abilities into an abstracted collection of high level objects which directly represent realistic entities (such as indoor elements, machines and realistic characters). It provides capacities such as skeletal animations or spatialized audio by interfacing with APF, while its clear composite pattern allows implementing more behaviors easily.

Finally, a generic design based on triggers and functors lets the final user implement complex scenarios of VR applications with the feeling of writing a script in C++. Applications developed on top of Imerse plug transparently into osgVR developed in the DREAM group (i.e., the research support development group of our INRIA center). We are using osgVR to render OSG's scene graph in a distributed manner, since rendering clusters are available in an increasing number of installations. osgVR is a software layer developed by the DREAM research support group, ensuring synchronization and events/inputs distribution among a list of rendering slaves. These two libraries are available on GForge.

4.3. APF: state-of-the-art 3D audio library

Participants: Adrien David, George Drettakis.

This work was performed in collaboration with Jean-Christophe Lombardo of the DREAM research engineer service at INRIA Sophia-Antipolis Méditerranée. REVES has several audio research publications over the last 10 years, which correspond to a class of functionalities. The first component is the masking or culling algorithm, which aims at removing all the inaudible audio sources from a virtual scene based on perceptual metrics. The second component, called clustering, aims at grouping audio sources that are spatially close to each other and premix them to a representative cluster source, so that all spatialization related processing can be applied only on the representative premixed source [9]. Other audio topics were also considered and developed, like progressive and scalable frequency domain mixing, sound propagation, scalable reverberation, modal sound synthesis and contact sounds generation [1].

In order to maintain all the knowledge in the group and re-use these technologies in the Immersive Space, a previous young engineer, a previous engineer (David Grelaud) wrote a fully documented audio library (APF) which gathers about 10 audio publications and 1 US patent. APF is a cross-platform, object oriented C++ API available on GForge. All the code has been re-implemented and a completely new software architecture resulted in a twofold increase in the speed of our algorithms. APF runs in the Immersive Space and uses the tracking system to spatialize virtual audio sources around the listener. It can also exploit personal Head Related Transfer Functions (HRTF).

We have implemented a network communications layer to create an audio rendering server on a separate machine, and the library is fully integrated into the osgVR platform.

APF has also been critical in establishing collaborations in the context of various grant proposals (EU and national).

4.4. GaborNoise Software

Participants: Ares Lagae, George Drettakis.

We proposed a new procedural noise function last year, Gabor noise [6]. In the context of this project, we have developed a software package, which includes a CPU reference implementation of the 2D noise, and a complete GPU implementation of the 2D noise, surface noise, and 3D noise. This software package has been filed for APP protection and is in the process of being transferred to industrial partners.

This work is a collaboration with Sylvain Lefebvre, former member of the team, now at INRIA Nancy.

5. New Results

5.1. Plausible Image Rendering

5.1.1. Filtering Solid Gabor Noise

Participants: Ares Lagae, George Drettakis.

Solid noise is a fundamental tool in computer graphics. Surprisingly, no existing noise function supports both high-quality anti-aliasing and continuity across sharp edges. Existing noise functions either introduce discontinuities of the solid noise at sharp edges, which is the case for wavelet noise and Gabor noise, or result in detail loss when anti-aliased, which is the case for Perlin noise and wavelet noise. In this project, we therefore present a new noise function that preserves continuity over sharp edges and supports high-quality anti-aliasing. We show that a slicing approach is required to preserve continuity across sharp edges, and we present a new noise function that supports anisotropic filtering of sliced solid noise. This is made possible by individually filtering the slices of Gabor kernels, which requires the proper treatment of phase. This in turn leads to the introduction of the phase-augmented Gabor kernel and random-phase Gabor noise, our new noise function. We demonstrate that our new noise function supports both high-quality anti-aliasing and continuity across sharp edges, as well as anisotropy. Fig. 3 shows several solid procedural textures generated with our new random-phase solid Gabor noise.

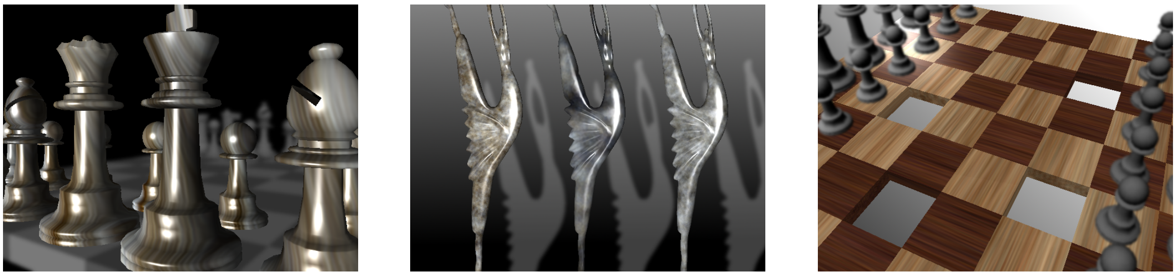

Figure 3. Filtering Solid Gabor Noise. Solid procedural textures generated with our new random-phase solid Gabor noise: (a) Granite chess pieces; (b) Granite dancer statuettes; (c) Wooden chess board. Our new noise function enables a variety of isotropic and anisotropic solid noise textures.

This work was presented at ACM SIGGRAPH 2011 in Vancouver and published in ACM Transactions on Graphics [18] (SIGGRAPH paper).

5.1.2. Image-Guided Weathering for Flow Phenomena

Participants: Carles Bosch, Pierre-Yves Laffont, George Drettakis.

The simulation of weathered appearance is essential in the realistic modeling of urban environments. With digital photography and Internet image collections, visual examples of weathering effects are readily available. These images, however, mix the appearance of the weathering phenomena with the specific local context. In [12], we have introduced a new methodology to estimate parameters of a phenomenological weathering simulation from existing imagery, in a form that allows new target-specific weathering effects to be produced. In addition to driving the simulation from images, we complement the visual result with details and colors extracted from the images. This methodology has been illustrated using flow stains as a representative case, demostrating how a rich collection of flow patterns can be generated from a small set of exemplars. In Fig. 4, we show the major components required for this approach.

This work was published in ACM Transactions on Graphics [12] and also presented at ACM SIGGRAPH 2011 in Vancouver (TOG paper).

5.1.3. Relighting Photographs of Tree Canopies

Participants: Marcio Cabral, George Drettakis.

We present an image-based approach to relighting photographs of tree canopies. Our goal is to minimize capture overhead; thus the only input required is a set of photographs of the tree taken at a time of day, while allowing relighting at any other time. We first analyze lighting in a tree canopy, both theoretically and using simulations. From this analysis, we observe that tree canopy lighting is similar to volumetric illumination. We assume a single-scattering volumetric lighting model for tree canopies, and diffuse leaf reflectance. To validate our assumptions, we apply our method on several synthetic renderings of tree models, for which we are all able to compute all quantities involved.

Our method first creates a volumetric representation of the tree using 10-12 images taken at a single time of day and use a single-scattering participating media lighting model - these photos are taken from different viewpoints, around the tree. An analytical sun and sky illumination model, namely the Preetham model, provides consistent representation of lighting for the captured input and unknown target times.

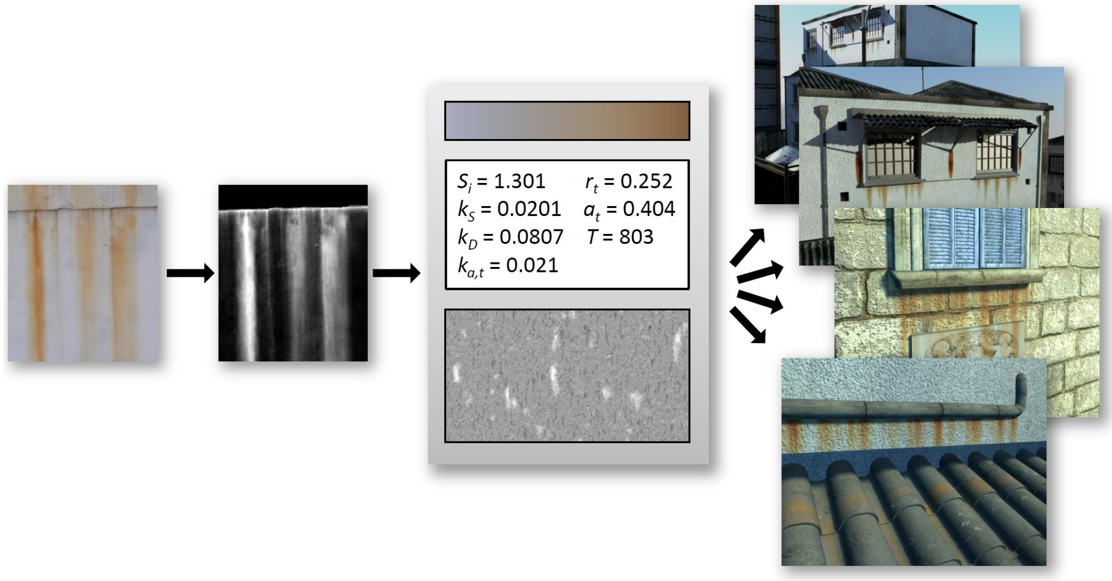

Figure 4. From left to right: given a photograph of a stain effect, we extract its stain degree map, we estimate the information required for the simulation (color, simulation parameters, and high frequency details), and run simulations on new scenes using this information, resulting in a wide variety of similar-looking stains.

We relight the input image by applying a ratio of the target and input time lighting representations. We compute this representation efficiently by simultaneously coding transmittance from the sky and to the eye in spherical harmonics. We validate our method by relighting images of synthetic trees and comparing to path-traced solutions. We also present results for photographs, validating with time-lapse ground truth sequences. An example is shown in Fig. 5. This work was published in the IEEE Transactions on Computer Graphics and Visualization [15], and was in collaboration with the past members of REVES, N. Bonneel and S. Lefebvre.

5.1.4. Silhouette-aware Warping for Image-based Rendering

Participants: Gaurav Chaurasia, George Drettakis.

We have presented a novel solution for image-based rendering (IBR) of urban scenes based on image warping. IBR techniques allow capture and display of 3D environments using photographs. Modern IBR pipelines reconstruct proxy geometry using multi-view stereo, reproject the photographs onto the proxy and blend them to create novel views. The success of these methods depends on accurate 3D proxies, which are difficult to obtain for complex objects such as trees and cars. Large number of input images do not improve reconstruction proportionally; surface extraction is challenging even from dense range scans for scenes containing such objects. Our approach does not depend on dense accurate geometric reconstruction; instead we compensate for sparse 3D information by variational warping. In particular, we formulate silhouette-aware warps that preserve salient depth discontinuities. This improves the rendering of difficult foreground objects, even when deviating from view interpolation. We use a semi-automatic step to identify depth discontinuities and extract a sparse set of depth constraints used to guide the warp. On the technnical side, our formulation breaks new ground by demonstrating how to incorporate discontinuities in variational warps. Our framework is lightweight and results in good quality IBR for previously challenging environments as shown in figure. 6.

Applications. Robust image-based rendering can be used to generate photo-realistic visual content easily which can be very useful for virtual reality applications. Commercial products like Google StreetView and

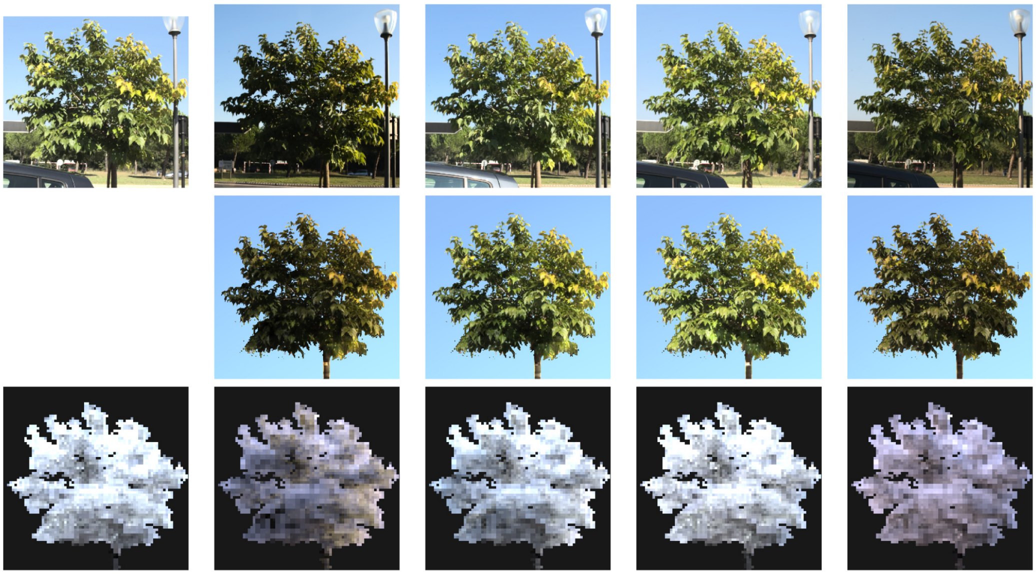

14h30 09h00 11h30 16h00 18h00

Mulberry tree Top row: input images and 4 target images with corresponding times of day. Middle row: 4 resulting relit images using our approach. Bottom row: $\tilde{E}_{in}$ and the four $\tilde{E}_{targ}$ images.

Figure 5.

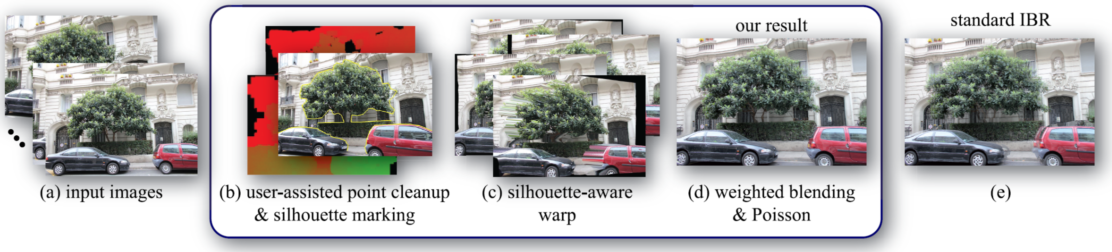

(a)

(b) (c) (d) (e)

Figure 6. Top: Overview of our approach starting from input images to rendered novel views along with novel view generated using current state of the art IBR. Bottom: Novel views rendered using our approach.

Microsoft PhotoSynth use rudimentary image-based rendering for large scale visualization of cities. Our work advances the state of the art by treating the hardest class of scenes.

This work is a collaboration with Prof. Dr. Olga Sorkine (ETH Zurich). The work was published in the special issue of the journal Computer Graphics Forum [16], presented at the Eurographics Symposium on Rendering 2011 in Prague, Czech Republic.

5.1.5. Real Time Rough Refractions

Participant: Adrien Bousseau.

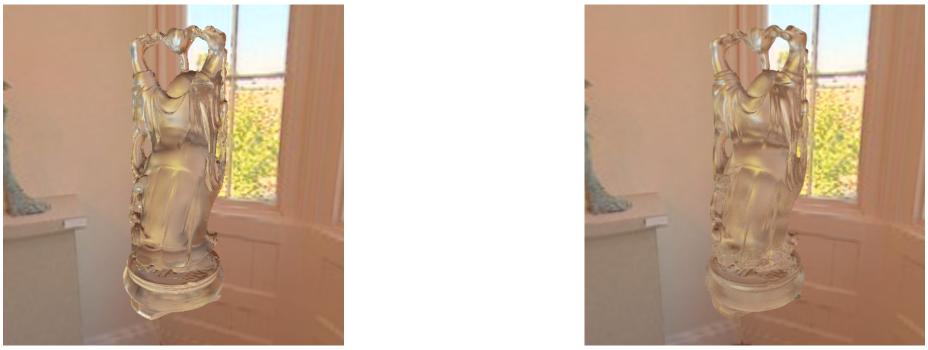

(a) Ground truth (b) Our method

Figure 7. Compared to an expensive ray-traced reference (a), our method produces plausible results in real-time (b).

We have proposed an algorithm to render objects made of *transparent materials with rough surfaces* in real-time, under environment lighting. Rough surfaces such as frosted and misted glass cause wide scattering as light enters and exits objects, which significantly complicates the rendering of such materials. We introduced two contributions to approximate the successive scattering events at interfaces, due to rough refraction: First, an approximation of the Bidirectional Transmittance Distribution Function (BTDF), using spherical Gaussians, suitable for real-time estimation of environment lighting using pre-convolution; second, a combination of cone tracing and macro-geometry filtering to efficiently integrate the scattered rays at the exiting interface of the object. Our method produces plausible results in real-time, as demonstrated by comparison against stochastic ray-tracing (Figure 7).

This work is a collaboration with Charles De Rousiers, Kartic Subr, Nicolas Holzschuch (ARTIS / INRIA Rhône-Alpes) and Ravi Ramamoorthi (UC Berkeley) as part of our CRISP associate team with UC Berkeley. The work was presented at the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games 2011 and won the best paper award.

5.1.6. Single view intrinsic images

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We introduced a new algorithm for decomposing photographs of outdoor scenes into *intrinsic images*, i.e. independent layers for illumination and reflectance (material color).

Extracting intrinsic images from photographs is a hard problem, typically solved with image-driven methods using numerous user indications. Recent methods in computer vision allow easy acquisition of medium-quality geometric information about a scene using multiple photographs from different views. We developed a new algorithm which allows us to exploit this noisy and somewhat unreliable information to automate and improve image-driven propagation algorithms to deduce intrinsic images. In particular, we develop a new approach

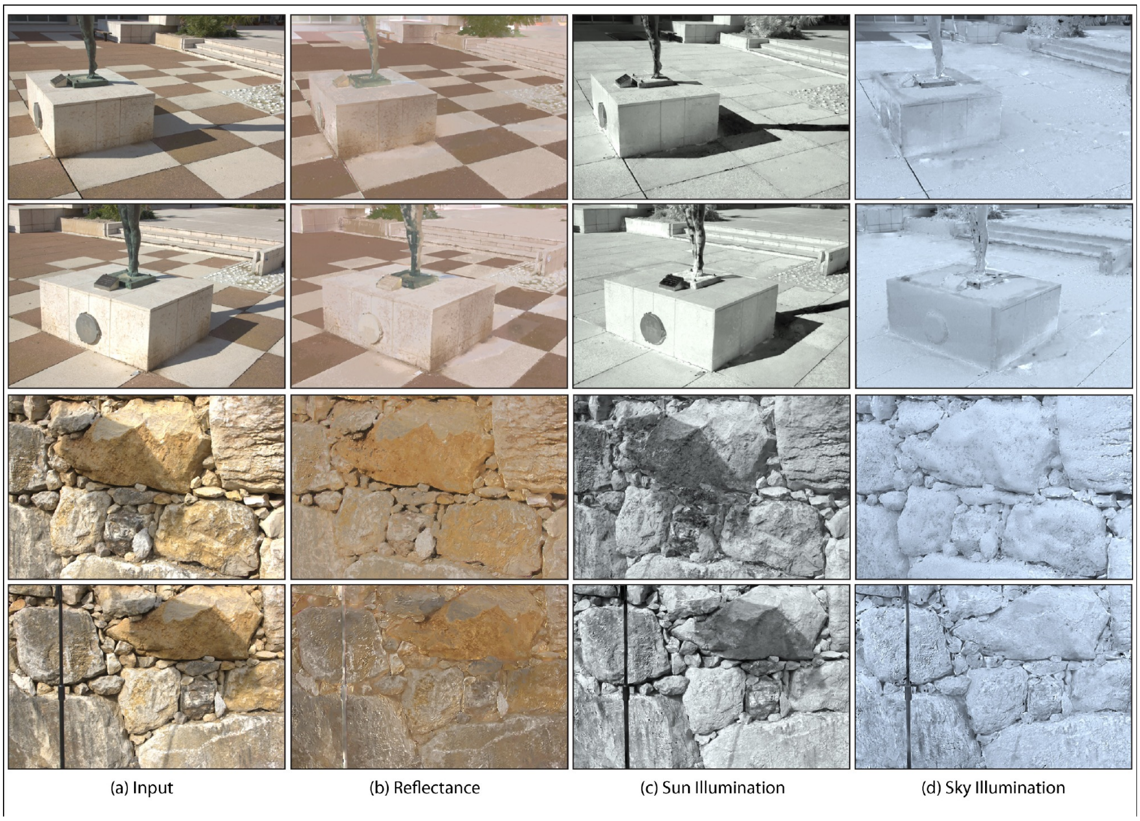

Figure 8. Starting from multiple views of the scene (a), our method automatically decomposes photographs into three intrinsic components — the reflectance (b), the illumination due to sun (c) and the illumination due to sky (d). Each intrinsic component can then be manipulated independently for advanced image editing applications.

to estimate cast shadow regions in the image, by refining an initial estimate available from reconstructed geometric information. We use a voting algorithm in color space which robustly identifies reflectance values for sparse reconstructed 3D points, allowing us to accurately determine visibility to the sun at these points. This information is propagated to the remaining pixels. In a final step we adapt appropriate image propagation algorithms, by replacing manual user indications by data inferred from the geometry-driven shadow and reflectance estimation. This allows us to automatically extract intrinsic images from multiple viewpoints, thus allowing many types of manipulation in images.

As illustrated on Figure 8, our method can extract reflectance at each pixel of the input photographs. The decomposition also yields separate sun illumination and sky illumination components, enabling easier manipulation of shadows and illuminant colors in image editing software.

This work won the best paper award at *"les journées de l'AFIG (Association Francophone d'Informatique Graphique) 2011"*, and will appear in the French journal REFIG [21].

5.1.7. Single view intrinsic images with indirect illumination

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

Building on top of our work published in [21], we extended our intrinsic image decomposition method to handle indirect illumination, in addition to sun and sky illumination. We introduced a new algorithm to reliably identify points in shadow based on a parameterization of the reflectance with respect to sun visibility, and proposed to separate illumination components using cascaded image-based propagation algorithm. We also demonstrated direct applications of our results for the end user in widespread image editing software. This rich intrinsic image decomposition method has been submitted for publication.

5.1.8. Multi-lighting multi-view intrinsic images

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

We compute intrinsic image decompositions using several images of the same scene under different viewpoints and lighting conditions. Such images can be easily gathered for famous landmarks, using photo sharing websites such as Flickr. Our method leverages the heterogeneity of photo collections (multiple viewpoints and lighting conditions) to guide the intrinsic image separation process. With this automatic decomposition, we aim to facilitate many image editing tasks and improve the quality of image-based rendering from photo collections.

This work is in collaboration with Frédo Durand (associate professor at MIT) and Sylvain Paris (researcher at Adobe).

5.1.9. Warping superpixels

Participants: Gaurav Chaurasia, Sylvain Duchêne, George Drettakis.

We are working on developing a novel representation of multi-view scenes in the form of *superpixels*. Such a representation segments the scene into semantically meaningful regions called superpixels which admit variational warps [16], thereby leveraging the power of shape-preserving warps while provinding automatic silhouette and occlusion handling. This entails several contributions namely incorporating geometric priors in superpixel extraction, generating consistent warping constraints and a novel rendering pipeline that assembles warped superpixels into the novel view maintaining spatio-temporal coherence.

This work is a collaboration with Prof. Dr. Olga Sorkine (ETH Zurich).

5.1.10. State-of-the-art Report on Temporal Coherence for Stylized Animations

Participant: Adrien Bousseau.

Non-photorealistic rendering (NPR) algorithms allow the creation of images in a variety of styles, ranging from line drawing and pen-and-ink to oil painting and watercolor. These algorithms provide greater flexibility, control and automation over traditional drawing and painting. The main challenge of computer generated stylized animations is to reproduce the look of traditional drawings and paintings while minimizing distracting flickering and sliding artifacts present in hand-drawn animations. These goals are inherently conflicting and any attempt to address the temporal coherence of stylized animations is a trade-off. This survey is motivated by the growing number of methods proposed in recent years and the need for a comprehensive analysis of the trade-offs they propose. We formalized the problem of temporal coherence in terms of goals and compared existing methods accordingly (Figure 9). The goal of this report is to help uninformed readers to choose the method that best suits their needs, as well as motivate further research to address the limitations of existing methods.

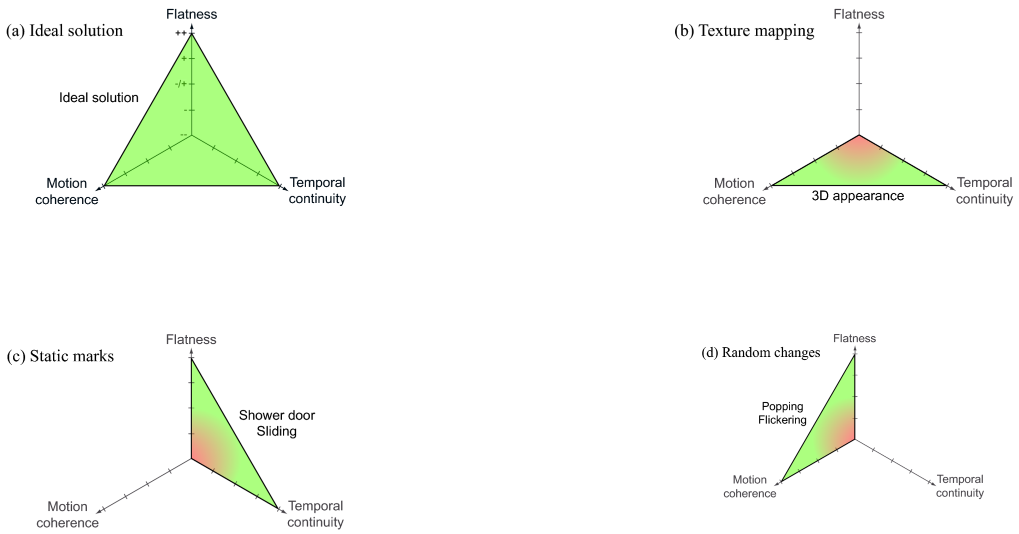

Figure 9. The temporal coherence problem involves three goals represented by the axes of these diagrams. Fully ignoring one of them produces the opposite artifacts.

This work is a collaboration with Pierre Bénard and Joëlle Thollot (ARTIS / INRIA Rhône-Alpes) and was published in the journal Computer Graphics Forum [14].

5.1.11. *Improving Gabor Noise*

Participant: Ares Lagae.

We have recently proposed a new procedural noise function, Gabor noise, which offers a combination of properties not found in existing noise functions. In this project, we present three significant improvements to Gabor noise: (1) an isotropic kernel for Gabor noise, which speeds up isotropic Gabor noise with a factor of roughly two, (2) an error analysis of Gabor noise, which relates the kernel truncation radius to the relative error of the noise, and (3) spatially varying Gabor noise, which enables spatial variation of all noise parameters. These improvements make Gabor noise an even more attractive alternative for existing noise functions. Fig. 3 shows a procedural textures generated with our new improved Gabor noise. This work was published in IEEE Transactions on Computer Graphics and Visualization [19]

5.2. Visual Perception and Audio Rendering

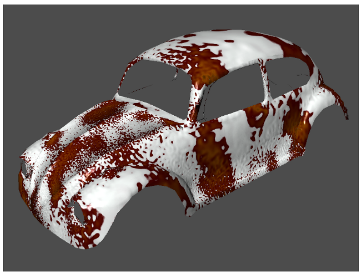

Figure 10. Improving Gabor noise. A procedural texture generated with spatially varying surface Gabor noise. The spatial variation in the size of the rust patterns and the scales is steered by surface curvature.

5.2.1. Perception of Visual Artifacts in Image-Based Rendering of Façades

Participants: Peter Vangorp, Gaurav Chaurasia, Pierre-Yves Laffont, George Drettakis.

(a) (b) (c)

Figure 11. (a) One of the environments used in our perceptual tests, with the input cameras shown. Examples of two of the artifacts we studied, namely (b) parallax distortion and (c) ghosting.

Image-based rendering (IBR) techniques allow users to create interactive 3D visualizations of scenes by taking a few snapshots (Figure 11 (left)). However, despite substantial progress in the field, the main barrier to better quality and more efficient IBR visualizations are several types of common, visually objectionable artifacts. These occur when scene geometry is approximate or viewpoints differ from the original shots, leading to parallax distortions (Figure 11 (mid)), blurring, ghosting (Figure 11(right)) and popping errors that detract from the appearance of the scene. We argue that a better understanding of the causes and perceptual impact of these artifacts is the key to improving IBR methods.

We present a series of psychophysical experiments in which we systematically map out the perception of artifacts in IBR visualizations of façades as a function of the most common causes. We separate artifacts into different classes and measure how they impact visual appearance as a function of the number of images available, the geometry of the scene and the viewpoint. The results reveal a number of counter-intuitive effects in the perception of artifacts.

We summarize our results in terms of the following practical guidelines for improving existing and future IBR techniques:

- When the total number of available images is small, e.g., because of storage limitations, it is preferable to use a sudden transition with its associated popping artifact rather than a gradual blending transition with its associated ghosting artifact.
- Interestingly, the depth range of the façade does not affect the *perceived* parallax distortions, even though it clearly does affect the *objective* parallax distortions. Only the intended output viewing angle should be taken into account when capturing images.
- For Google Street View™-like visualizations, a shorter cross-fading transition would improve the perceived quality.

This work is a collaboration with Roland W. Fleming (Justus-Liebig-University Giessen, Germany). The work was published in the special issue of the journal Computer Graphics Forum [20] and presented at the Eurographics Symposium on Rendering 2011.

5.2.2. Perception of Slanted, Textured Façades

Participants: Peter Vangorp, Adrien Bousseau, Gaurav Chaurasia, George Drettakis.

In large-scale urban visualizations, buildings are often geometrically represented by simple boxes textured with images of the façades. Any depth variations in the façade, such as balconies, are perceived to have distorted angles when the viewer is not at the capture camera position. The *retinal hypothesis* provides the most likely prediction of the magnitude of the perceived distortion. We conduct psychophysical experiments to measure the perceived distortion, thereby validating the retinal hypothesis, and to measure the threshold for detecting any distortion. The result is a prediction of the valid range of viewer motion for a given capture.

This work is a collaboration with Martin S. Banks (UC Berkeley).

5.2.3. Binocular and Dynamic Cues to Glossiness

Participants: Peter Vangorp, George Drettakis.

Recent advances in display technology have made it possible to present high quality stereoscopic imagery with accurate head tracking. This improves not only depth perception but also affects the perception of glossy materials. Previous work has shown that these conditions can increase the perceived gloss by a small amount. We conduct psychophysical experiments to measure this effect quantitatively.

This work is a collaboration with Roland W. Fleming (Justus-Liebig-University Giessen, Germany) and Martin S. Banks (UC Berkeley).

5.2.4. Sound Particles

Participants: Charles Verron, George Drettakis.

This research deals with a sound synthesizer dedicated to particle-based environmental effects, and intended to be used in interactive virtual environments. The synthesis engine is based on five physically-inspired basic elements (sound atoms) that can be parameterized and stochastically distributed in time and space. Physically-inspired controls simultaneously drive graphics particle models (e.g., distribution of particles, average particles velocity etc.) and sound parameters (e.g., distribution of sound atoms, spectral modifications etc.). The simultaneous audio/graphics controls result in an intricate interaction between the two modalities that enhances the naturalness of the scene. The approach is currently illustrated with three environmental phenomena: fire, wind, and rain.

5.2.5. Sound Synthesis for Crowds

Participants: Charles Verron, George Drettakis.

We are currently investigating new methods for synthesis of crowd sounds in virtual environments. Crowd sounds are constituted of many overlapping voices spatialized at different positions in the environment. A novel level of detail for crowd sounds is desirable: the cost of spatializing many individual voice sounds can be replaced by an efficient babble noise synthesis model. Furthermore, high-level control should be proposed to modify the crowd sounds by semantic parameters, related to the crowd emotional state (e.g., calm, angry...). This research should result in a new real-time crowd sound synthesizer with semantic controls for virtual environments.

5.3. Interaction and Design for Audiovisual Virtual Environments

5.3.1. Lighting Design for Material Depiction

Participants: Adrien Bousseau, Emmanuelle Chapoulie.

Shading, reflections and refractions are important visual features for understanding the shapes and materials in an image. While well designed lighting configurations can enhance these features and facilitate image perception (Figure 12b), poor lighting design can lead to misinterpretation of image content (Figure 12a).

(a) Our optimized lighting emphasizes materials (b) Poor lighting de-emphasizes materials

Figure 12. Our method (a) automatically optimizes the lighting to enhance material-specific visual features. The lighting reveals the thin and thick parts of the subsurface scattering wax candle, it accentuates the Fresnel reflections along the side of the porcelain vase and it adds strong specular highlights to emphasize the shiny chrome metal of the sculpture. Poorly designed lighting (b) diminishes these characteristic visual features of the materials. The candle appears more like solid plastic, the vase appears to be made of diffuse clay and the sculpture no longer looks like it is made of chrome.

We have presented an automated system for optimizing and synthesizing environment maps that enhance the appearance of materials in a scene. We first identified a set of lighting design principles for material depiction. Each principle specifies the distinctive visual features of a material and describes how environment maps can emphasize those features. We then proposed a general optimization framework to solve for the environment map that best fulfill the design principles. Finally we described two techniques for transforming existing photographic environment maps to better emphasize materials. Our approach generates environment maps that enhance the depiction of a variety of materials including glass, metal, plastic, marble and velvet.

This work is a collaboration with Ravi Ramamoorthi and Maneesh Agrawala (UC Berkeley) as part of our CRISP associate team with UC Berkeley. The work was published in the special issue of the journal Computer Graphics Forum, presented at the Eurographics Symposium on Rendering 2011.

5.3.2. A Multimode Immersive Conceptual Design System for Architectural Modeling and Lighting

Participants: Marcio Cabral, Peter Vangorp, Gaurav Chaurasia, Emmanuelle Chapoulie, Martin Hachet, George Drettakis.

We developed a system which allows simple architectural design in immersive environments. The user is able to define the initial conceptual design of the model and can take into account the effects of daylight. Our system allows the manipulation of simple elements such as windows, doors and rooms while the overall model is automatically adjusted to the manipulation. The system runs on a four-sided stereoscopic, head-tracked immersive display. We also provide simple lighting design capabilities, with an abstract representation of sunlight and its effects when shining through a window. Our system provides three different modes

of interaction: a miniature-model *table mode*, a fullscale *immersive mode* and a combination of table and immersive which we call *mixed mode* (see Figure 13). Our goal is to study direct manipulation for basic 3D modeling in an immersive setting, in the context of conceptual or initial design for architecture.
We performed an initial pilot user test to evaluate the relative merits of each mode for a set of basic tasks such as resizing and moving windows or walls, and a basic light-matching task. The study indicates that users appreciated the immersive nature of the system, and found interaction to be natural and pleasant. In addition, the results indicate that the mean performance times seem quite similar in the different modes, opening up the possibility for their combined usage for effective immersive modeling systems for novice users.

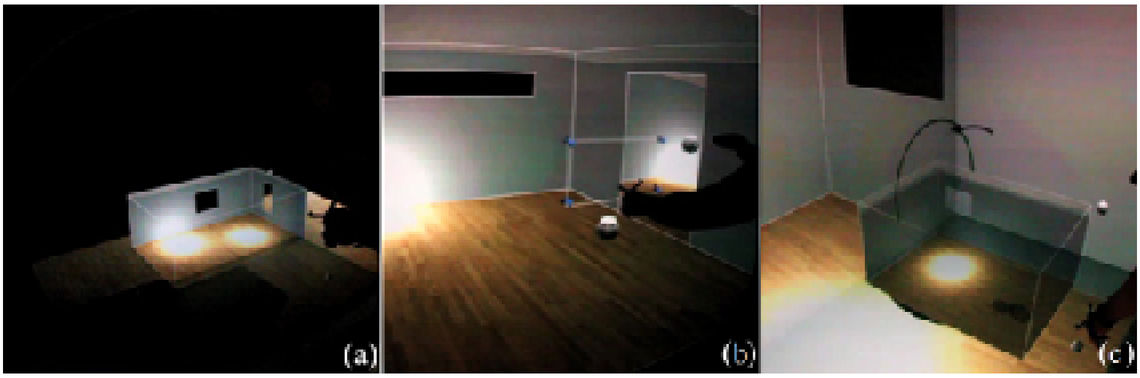

Figure 13. (a) Table mode, (b) Immersive mode, (c) Mixed mode

5.3.3. *Walking in a Cube: Novel Metaphors for Safely Navigating Large Virtual Environments in Restricted Real Workspaces*

Participants: Peter Vangorp, Emmanuelle Chapoulie, George Drettakis.

Immersive spaces such as 4-sided displays with stereo viewing and high-quality tracking provide a very engaging and realistic virtual experience. However, walking is inherently limited by the restricted physical space, both due to the screens (limited translation) and the missing back screen (limited rotation). We propose three novel navigation techniques that have three concurrent goals: keep the user safe from reaching the translational and rotational boundaries; increase the amount of real walking; and finally, provide a more enjoyable and ecological interaction paradigm compared to traditional controller-based approaches. We notably introduce the “Virtual Companion”, which uses a small bird to guide the user through VEs larger than the physical space. We evaluate the three new techniques through a user study with pointing and path following tasks. The study provides insight into the relative strengths of each new technique for the three aforementioned goals.

This work is a collaboration with Gabriel Cirio, Maud Marchal, and Anatole Lécuyer (VR4I / INRIA Rennes) in the context of ARC NIEVE (Section 6.2.1) and has been accepted for publication [17].

5.3.4. *Inferring Normals Over Design Sketches*

Participant: Adrien Bousseau.

We are currently working on a sketch-based tool to infer normals over a 2D drawing. Our tool should allow users to apply realistic and non-photorealistic shading over the drawing, with applications in product design.

This work is a collaboration with Alla Sheffer (University of British Columbia), Cloud Shao and Karan Singh (University of Toronto).

5.3.5. Using natural gestures into virtual reality immersive space

Participants: Emmanuelle Chapoulie, George Drettakis.

We are studying the use of gestures which are as natural as possible in a context of virtual reality environments. We define a scenario which is a sequence of tasks (hiding, finding, pushing, pulling, grabbing, picking up, putting down objects) that the users will perform with hands and with wands, in order to evaluate the usability of our interaction approach. Each task will be used to evaluate a specific criterion.

6. Contracts and Grants with Industry

6.1. Grants with Industry

6.1.1. Industrial Contracts and Donations

6.1.1.1. Technology Transfer with 3Delight

Participants: Ares Lagae, George Drettakis.

The Gabor noise technology transfer with 3Delight http://www.3delight.com/, developer of high performance rendering software, has been finalized.

6.1.1.2. Autodesk

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

Autodesk has offered a significant cash donation to REVES in support of our work on intrinsic images. We are in the process of establishing a working collaboration with the local office in Sophia-Antipolis on this topic and negociating a technology transfer agreement for the RID system on single-lighting condition intrinsic images. Autodesk has also donated several licenses of Maya, 3DS Max and SketchBookPro.

6.1.1.3. Adobe

Participants: Pierre-Yves Laffont, Adrien Bousseau, George Drettakis.

In the context of our collaboration with Adobe (project with S. Paris and F. Durand from MIT), we have received a cash donation in support of our research.

6.1.1.4. NVIDIA

We have received several graphics cards as part of the professor partnership program.

6.2. National Initiatives

6.2.1. ARC NIEVE: Navigation and Interfaces in Emotional Virtual Environments

Participants: Peter Vangorp, Adrien David, George Drettakis, Gaurav Chaurasia, Emmanuelle Chapoulie.

The goal of this joint research project is to develop and evaluate improved interfaces for navigation in immersive virtual environments (VEs) such as the 4-wall stereoscopic ISpace system in the Immersive Space Gouraud-Phong.

This project is in collaboration with G. Cirio, M.Marchal, A. Lécuyer (VR4I, Rennes) and I. Viaud-Delmon (IRCAM).

There is evidence of significant overlap in brain structures related to spatial memory and orientation and those related to emotion. We will examine the influence of high-quality 3D visual and auditory stimuli on the emotions evoked by the virtual environment. Our study will focus on the phobia of dogs as a way to modulate emotion in audiovisual VEs (Figure 14).

Navigation in VEs involves the use of different views, i.e., egocentric ("first person") and allocentric ("bird's eye") views during navigation tasks. We will study appropriate visual representations for each view (for example, the level of realism ranging from abstract map-like rendering for top-down views to photorealistic rendering for first-person views), and appropriate transitions between the different views.

We will develop an appropriate methodology to evaluate such navigation interfaces in stressful environments, based on the insights gained by the emotion modulation study in phobic settings. This novel methodology can be seen as a "stress-test" for navigation interfaces: if the navigation interfaces developed are successful even in stressful setups, they will definitely be successful under "normal conditions".

ARC NIEVE has resulted in several publications this year: [22], [20], [16], [17].

Figure 14. A person immersed in a virtual environment where the behaviors of several dogs will evoke different levels of anxiety.

This is a joint research project with Isabelle Viaud-Delmon (IRCAM, CNRS), Anatole Lécuyer and Maud Marchal (VR4I / INRIA Rennes), and Jean-Christophe Lombardo (DREAM / INRIA Sophia Antipolis). Interact3D (Section 6.2.3) is associated with this ARC.

6.2.2. ANR ALTA

Participants: Emmanuelle Chapoulie, Adrien David, George Drettakis.

The ANR ALTA project started in October 2011, and focuses on the development of novel algorithms for realistic and efficient global illumination. The project is coordinated by the Grenoble INRIA group ARTIS (N.Holzschuch), and the Bordeaux INRIA group IPARLA (X. Granier) is also a partner.

Our participation will mainly be in the study of error bounds for these algorithms and the development of interactive global illumination solutions that can be used in Virtual Reality solutions, for example in the context of the immersive space.

6.2.3. ADT Interact3D

Participants: Adrien David, George Drettakis.

This ADT involves half time software development for ARC NIEVE, and the other half general support to the new Immersive Space Gouraud-Phong in Sophia-Antipolis (supervised by Jean-Christophe Lombardo of the DREAM service). The main contribution was the complete rewrite of our VR application environment with the development of the Imerse software. This platform will allow first experiments, and the development of a generic Virtual Reality framework addressing neuroscience/psychology applications. This generic platform is based on osgVR which aims at a high-quality context abstraction to be usable in several domains, as well as distributed rendering capacities. These improvements, deployable for a variety of applications to come, are tightly coupled with the current ARC NIEVE, thus contributing to its implementation. Future prospects for the ADT Interact 3D include developing novel multimodal interaction techniques for example for gesture-based interaction etc.

6.2.4. National French Bilateral Collaboration

We have ongoing collaborations with J. Thollot (INPG/INRIA Grenoble) [14], , M. Hachet (IPARLA, Bordeaux) [22], and B. Galerne (ENST/ENS Cachan) and Sylvain Lefebvre (ALICE/Nancy) [19]

6.3. European Initiatives

6.3.1. FP7 Projet

6.3.1.1. VERVE

Title: VERVE:
Type: COOPERATION (ICT)
Defi: Services to promote E-inclusion using socially realistic virtual environments
Instrument: Integrated Project (IP)
Duration: October 2011 - September 2014
Coordinator: Trinity College - Dublin (Ireland)
Others partners: DFKI (Germany), CNRS, IRCAM, U. of Zaragoza (Spain), Testaluna (IT), KAINOS (UK)
See also: http://www.verveconsortium.eu/
Abstract: Social exclusion has many causes, but major factors are the fear and apathy that often accompany a disability. The European e-Inclusion policy stresses the importance of ICT in improving the quality of life in potentially disadvantaged groups, including older people and persons with disabilities. In this project, we will develop ICT tools to support the treatment of people who are at risk of social exclusion due to fear and/or apathy associated with a disability. These tools will be in the form of personalised VR scenarios and serious games specifically designed for therapeutic targets and made broadly available via a novel integration of interactive 3D environments directly into Web browsers. We will perform cutting edge research into rendering and simulating personalised and populated VR environments, 3D web graphics, and serious games. These technical efforts will be underpinned by our clinical/laboratory and industry partners, who will be fully involved throughout in the requirements, design and evaluation of VERVE, and liaison with the stakeholders (i.e., participants, carers/family, and health professionals). They will implement the VERVE interventions in three use-cases, each targeting a different group of participants: Fear of falling, Apathy related to cognitive decline and behavioural disturbances, and other emotional disturbances linked to anxiety. While developing clinical assessment methods and interventions for the first two patient groups is our primary focus, our results will be applicable to a much wider range of potentially disadvantaged individuals.

7. Partnerships and Cooperations

7.1. Bilateral Collaborations

7.1.1. France-USA

Participants: Gaurav Chaurasia, Pierre-Yves Laffont, Adrien Bousseau, Carles Bosch, George Drettakis, Peter Vangorp.

We have an ongoing collaboration (C. Bosch) with Yale University (Holly Rushmeier and Julie Dorsey), on weathering, resulting in the publication [12]. We continue this collaboration of fracture-related weathering.

We have an ongoing collaboration with Adobe Research (Sylvain Paris) and Fredo Durand (MIT) on intrinsic images for mutliple lighting conditions. In this context Pierre-Yves Laffont visited MIT this summer July 6 - September 2.

We also collaborate with M. Banks, R. Ramamoorthi and M. Agrawala from the University of California, Berkeley in the context of our CRISP associate team, resulting in the publication [23].

7.1.2. France-Switzerland

Participants: Gaurav Chaurasia, George Drettakis.

We collaborate with O. Sorkine at ETH Zurich on imaeg-based rendering, which resulted in the publication [16].

7.1.3. France-Germany

Participants: Peter Vangorp, George Drettakis.

We collaborate with the Justus-Liebig-University Giessen, Germany on perception techniques for rendering, notably with R. Fleming. This resulted in the following publication [20]. P. Vangorp is now a postdoctoral researcher there and this collaboration continues.

7.1.4. France-Spain

Participants: Pierre-Yves Laffont, Carles Bosch, George Drettakis.

We collaborate with C. Bosch who is now at the University of Girona (Spain), on weathering and normal mapping.

7.1.5. France-Canada

Participant: Adrien Bousseau.

We collaborate with K. Singh (University of Toronto) and Alla Schaeffer (U. British Columbia, Vancouver), on sketching techniques for materials.

7.1.6. France-Belgium

Participant: George Drettakis.

We have continued the collaboration with A. Lagae and P. Dutré and the Catholic University of Leuven, resulting in the publications [19], [18].

7.2. Visiting Researchers

We hosted several researchers this year:

- Toshiya Hachisuka (Univ. of California San Diego), in February
- Alla Sheffer (Univ. of British Columbia), in March
- Ares Lagae (KU Leuven, FWO Belgium), in March, June, and November
- Martin Banks (Univ. Of Berkeley), in April and October
- Bernd Froehlich (Bauhaus-Universitaet Weimar), in May
- Olga Sorkine (New York Univ.), in May
- Eugene Fiume (Univ. of Toronto), in June
- Christian Lessig (Univ. of Toronto), in June
- James O'Shea (Univ. of Berkeley), in July
- J.P. Lewis (Weta Digital, Wellington, NZ) in October
- Maria Roussou (Makebelieve, Greece), in November
- Carles Bosch (Univ. of Girona), in November
- Insu Yu (UCL London), in November
- Cloud Shao (Univ. of Toronto) in Nov/Dec.

7.3. International Initiatives

7.3.1. INRIA Associate Teams

7.3.1.1. CRISP

Title: Human Perception

INRIA principal investigator: George Drettakis

International Partner:

Institution: University of California Berkeley (United States)

Laboratory: Electrical Engineering and Computer Science

Duration: 2011 - 2013

See also: http://www-sop.inria.fr/reves/crisp/

The goal of the CRISP associate team between REVES and University of California (UC) Berkeley is to investigate novel ways to create, render and interact with images based on the study of human Perception. This novel and emerging area has been the focus of ongoing collaborations between researchers from the REVES research group at INRIA (Adrien Bousseau, George Drettakis) and researchers in Computer Science and Vision Science at UC Berkeley (Maneesh Agrawala, Ravi Ramamoorthi, Martin S. Banks (Human Vision Science)). All of the researchers involved in CRISP share a common interest in creating and manipulating effective synthetic imagery. To achieve this goal we will focus on understanding how people perceive complex material, lighting and shape, on developing new rendering algorithms based on this understanding, and on building interactive tools that enable users to efficiently specify the kind of image they wish to create. More specifically, we will explore the following research directions :Perception: Images are generated from the interaction of lighting, material, and geometry. We will evaluate how people perceive material, lighting, and geometry in realistic images such as photographs, and non realistic images such as drawings and paintings. This knowledge of human perception is essential for developing efficient rendering algorithms and interaction tools that focus on the most important perceptual features of an image. We have started several projects on the perception of materials in realistic and non realistic images, with promising results. Rendering: We will develop rendering algorithms that generate images that are plausible with respect to the user's intent and allocate resources on the visual effects that best contribute to perception. Current projects on rendering include work on enhancing material variations in realistic and non realistic rendering.Interaction: We will facilitate the creation of material, lighting, and geometric effects in synthetic images by developing novel user interfaces for novice and professional users. We are currently working on interfaces to draw object appearance and to relight photographs.Our contributions have the potential to benefit different applications of image creation such as illustration (archeology, architecture, education); entertainment (video games, movies) and design (sketching, photograph editing). This research naturally falls in INRIA's strategic objective of interacting with real and virtual worlds.

8. Dissemination

8.1. Animation of the scientific community

8.1.1. Program Committees

G. Drettakis, A. Bousseau and A. Lagae served on the Eurographics Symposium on Rendering program committee. A. Bousseau served on the NPAR program committee.

8.1.2. Invited Talks

C. Verron presented his research at the annual meeting of SFA in the context of his award. Carles Bosch gave two invited talks in the context of his ERCIM fellowship: Pisa (Oct. 2010), ETH Zurich (March 2011).

8.1.3. Thesis Committees

G. Drettakis was a rapporteur for Ph.D. committee of Christian Eisenacher (Erlangen, Germany) and a member of the Ph.D. committee of G. Cirio (INRIA Rennes).

8.1.4. Community service

G. Drettakis chairs the Eurographics (EG) Awards Committee and the EG Working group on Rendering and is part of the EG conference steering committee. He is an associate editor-in-chief for IEEE Transactions on Computer Graphics and Visualization. G. Drettakis is a member of the “comité d’animation scientifique” for the Interaction, Cognition and Perception theme of INRIA since October 2009.

8.1.5. Conference Presentations and Attendance

At SIGGRAPH 2011 in Vancouver, Ares Lagae presented the paper [18]. Carles Bosch presented the paper [12] (August). P. Vangorp, G. Chaurasia and A. Bousseau presented the papers [20] respectively at the Eurographics Symposium on Rendering in Prague (June). P-Y. Laffont presented his work [21] the AFIG conference in November in Bidart and attended the summer school “Toward petaflop numerical simulation on parallel hybrid architectures”, June 6-10 2011 in Sophia-Antipolis. G. Drettakis and A. Bousseau attended EUROGRAPHICS 2011. E. Chapoulie and G. Chaurasia ENS/INRIA Visual Recognition and Machine Learning Summer School in Paris, France, from July 25th, 2011 to July 29th, 2011. P. Vangorp attended APGV in Toulouse (September).

8.2. Teaching

Licence: E. Chapoulie teaches Analogical electronics 32h, UNSA EPU (L1). P-Y. Laffont teaches Object Oriented Programming 2 (L3), UNSA EPU (64h).

Masters: G. Drettakis organizes and teaches Computer Graphics at the ECP (Paris) (9h), A. Bousseau teaches 3h at the same course. G. Drettakis organizes and teaches the M1 Module in Computer Graphics in the PENSUNS Masters (6h), A. Bousseau teaches 6h in the course. G. Drettakis teaches 6h in the MAPI M1 Module (Jeux Video), A. Bousseau teaches 3h in this program.

Ph.D.s in progress:

Emmanuelle Chapoulie is continuing her Ph.D. on "Lighting Algorithms and Virtual Reality" (Supervisor: G.Drettakis, started Oct.2010).

Gaurav Chaurasia is continuing his Ph.D. on "Texture Synthesis Rendering" (Supervisor: G.Drettakis, started Sept. 2010)

Sylvain Duchene started in Nov. his Ph.D on "Lightweight Capture, Display and Manipulation of Urban Environments" (Supervisor: G.Drettakis),

Pierre-Yves Laffont is continuing his Ph.D. on Image- and Video-Based Relighting for Virtual Environments (Supervisor: G.Drettakis (co-supervisor A. Bousseau), started Oct.2009).

8.3. Demonstrations and Press

8.3.1. Demonstrations

Participants: Adrien David, George Drettakis.

We performed many demonstrations this year, mostly in the Immersive Space but also of our various research results. Specifically, we performed demos to the companies Computmaps, EON, Stonetrip and Technicom. We also presented demos to students of the UNSA EPU, the Franhofer Institute, the Prefet de la Region and a delegation from Taiwan.

8.3.2. Press

Participant: George Drettakis.

During the visit of Tony DeRose from Pixar, the INRIA Press agency interviewed G. Drettakis and Dr. DeRose, resulting in a press article posted on the INRIA site (as well as some coverage in the local press): http://www.inria.fr/en/centre/sophia/news/tony-derose-pixar.

9. Bibliography

Major publications by the team in recent years

[1] N. BONNEEL, G. DRETTAKIS, N. TSINGOS, I. VIAUD-DELMON, D. JAMES. *Fast Modal Sounds with Scalable Frequency-Domain Synthesis*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", August 2008, vol. 27, n^{o} 3, http://www-sop.inria.fr/reves/Basilic/2008/BDTVJ08.

[2] C. DACHSBACHER, M. STAMMINGER, G. DRETTAKIS, F. DURAND. *Implicit Visibility and Antiradiance for Interactive Global Illumination*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", August 2007, vol. 26, n^{o} 3, http://www-sop.inria.fr/reves/Basilic/2007/DSDD07.

[3] 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/Basilic/2002/DCSD02/Vis2002Deussen.pdf.

[4] F. DUGUET, G. DRETTAKIS. *Robust Epsilon Visibility*, in "Proceedings of ACM SIGGRAPH 2002", J. HUGHES (editor), ACM Press / ACM SIGGRAPH, July 2002, http://www-sop.inria.fr/reves/Basilic/2002/DD02/RobustEpsilonVisibility.pdf.

[5] F. DUGUET, G. DRETTAKIS. *Flexible Point-Based Rendering on Mobile Devices*, in "IEEE Computer Graphics and Applications", July-August 2004, vol. 24, n^{o} 4.

[6] A. LAGAE, S. LEFEBVRE, G. DRETTAKIS, P. DUTRÉ. *Procedural Noise using Sparse Gabor Convolution*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", August 2009, vol. 28, n^{o} 3, http://www-sop.inria.fr/reves/Basilic/2009/LLDD09.

[7] A. RECHE, I. MARTIN, G. DRETTAKIS. *Volumetric Reconstruction and Interactive Rendering of Trees from Photographs*, in "ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)", July 2004, vol. 23, n^{o} 3, http://www-sop.inria.fr/reves/Basilic/2004/RMD04/FinalPaper_0167.pdf.

[8] M. STAMMINGER, G. DRETTAKIS. *Perspective Shadow Maps*, in "Proceedings of ACM SIGGRAPH 2002", J. HUGHES (editor), ACM Press/ ACM SIGGRAPH, July 2002, http://www-sop.inria.fr/reves/Basilic/2003/SDD03/.

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

[10] B. WALTER, G. DRETTAKIS, D. GREENBERG. *Enhancing and Optimizing the Render Cache*, in "Proceedings of the Eurographics Workshop on Rendering", P. DEBEVEC, S. GIBSON (editors), ACM Press, June 2002, http://www-sop.inria.fr/reves/Basilic/2002/WDG02.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] M. CABRAL. *Reshape and Relighting for Interactive Content Creation and Manipulation*, University of Nice Sophia-Antipolis, May 2011, http://www-sop.inria.fr/reves/Basilic/2011/Cab11.

Articles in International Peer-Reviewed Journal

[12] C. BOSCH, P.-Y. LAFFONT, H. RUSHMEIER, J. DORSEY, G. DRETTAKIS. *Image-Guided Weathering: A New Approach Applied to Flow Phenomena*, in "ACM Transactions on Graphics", 2011, p. 20:1–20:13, (Also presented at SIGGRAPH 2011), http://hal.inria.fr/inria-00606830/en.

[13] A. BOUSSEAU, E. CHAPOULIE, R. RAMAMOORTHI, M. AGRAWALA. *Optimizing Environment Maps for Material Depiction*, in "Computer Graphics Forum", July 2011, vol. 30, n^o 4, http://hal.inria.fr/inria-00606838/en.

[14] P. BÉNARD, A. BOUSSEAU, J. THOLLOT. *State-of-the-Art Report on Temporal Coherence for Stylized Animations*, in "Computer Graphics Forum", October 2011 [*DOI :* 10.1111/J.1467-8659.2011.02075.X], http://hal.inria.fr/inria-00636210/en.

[15] M. CABRAL, N. BONNEEL, S. LEFEBVRE, G. DRETTAKIS. *Relighting Photographs of Tree Canopies*, in "IEEE Transactions on Visualization and Computer Graphics", 2011 [*DOI :* 10.1109/TVCG.2010.236], http://hal.inria.fr/inria-00607042/en.

[16] G. CHAURASIA, O. SORKINE, G. DRETTAKIS. *Silhouette-Aware Warping for Image-Based Rendering*, in "Computer Graphics Forum", 2011, vol. 30, n^o 4, http://hal.inria.fr/inria-00607039/en.

[17] G. CIRIO, P. VANGORP, E. CHAPOULIE, M. MARCHAL, A. LÉCUYER, G. DRETTAKIS. *Walking in a Cube: Novel Metaphors for Safely Navigating Large Virtual Environments in Restricted Real Workspaces*, in "IEEE Transactions on Visualization and Computer Graphics (Proceedings of IEEE Virtual Reality)", 2012, (conditionally accepted).

[18] A. LAGAE, G. DRETTAKIS. *Filtering Solid Gabor Noise*, in "ACM Transactions on Graphics", August 2011, vol. 30, n^o 3, http://hal.inria.fr/inria-00606907/en.

[19] A. LAGAE, S. LEFEBVRE, P. DUTRÉ. *Improving Gabor Noise*, in "IEEE Transactions on Visualization and Computer Graphics", August 2011, vol. 17, n^o 8, p. 1096 - 1107 [*DOI :* 10.1109/TVCG.2010.238], http://hal.inria.fr/inria-00601695/en.

[20] P. VANGORP, G. CHAURASIA, P.-Y. LAFFONT, R. FLEMING, G. DRETTAKIS. *Perception of Visual Artifacts in Image-Based Rendering of Façades*, in "Computer Graphics Forum", 2011, vol. 30, n^o 4, http://hal.inria.fr/inria-00606832/en.

Articles in National Peer-Reviewed Journal

[21] P.-Y. LAFFONT, A. BOUSSEAU, G. DRETTAKIS. *Images intrinsèques de scènes en extérieur à partir de multiples vues*, in "REFIG (Revue Electronique Francophone d'Informatique Graphique)", 2011, vol. 5, n^o 2, p. 57-67, In French. Presented at *Journées de l'AFIG 2011*, **best paper award**..

International Conferences with Proceedings

[22] M. CABRAL, P. VANGORP, G. CHAURASIA, E. CHAPOULIE, M. HACHET, G. DRETTAKIS. *A Multimode Immersive Conceptual Design System for Architectural Modeling and Lighting*, in "3D User Interfaces (3DUI), 2011 IEEE Symposium on", Singapour, Singapore, JOSEPH J. JR. LAVIOLA, M. HACHET, M. BILLINGHURST (editors), IEEE, 2011, p. 15-18 [*DOI :* 10.1109/3DUI.2011.5759211], http://hal.inria.fr/inria-00606833/en.

[23] C. DE ROUSIERS, A. BOUSSEAU, K. SUBR, N. HOLZSCHUCH, R. RAMAMOORTHI. *Real-Time Rough Refraction*, in "ACM Siggraph Symposium on Interactive 3D Graphics and Games (I3D)", San Francisco, United States, 2011 [*DOI :* 10.1145/1944745.1944764], http://hal.inria.fr/inria-00545955/en.

References in notes

[24] 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/Basilic/2002/DCSD02/Vis2002Deussen.pdf.

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

[26] 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/Basilic/2002/DD02/RobustEpsilonVisibility.pdf.

[27] F. DURAND, G. DRETTAKIS, C. PUECH. *The 3D Visibility Complex, a new approach to the problems of accurate visibility*, in "Rendering Techniques'96 (7th Eurographics Workshop on Rendering)", Springer Verlag, June 1996, p. 245–257.

[28] 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.

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

[30] T. FUNKHOUSER, I. CARLBOM, G. ELKO, G. PINGALI, M. SONDHI, J. WEST. *A Beam Tracing Approach to Acoustic Modeling for Interactive Virtual Environments*, in "ACM Computer Graphics (SIGGRAPH'98 Proceedings)", July 1998.

[31] 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.

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

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

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

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