Surface Light Field Rendering for Virtual Product Design

Jan MESETH, Gero MÜLLER, Reinhard KLEIN
Universität Bonn, Institut für Informatik II, Römerstr. 164, 53117 Bonn, Germany

Abstract. Interior and material designers currently refuse to employ Virtual Prototyping software due to potential errors in computer generated images. In this paper we propose real-time rendering of a precomputed global illumination solution including accurate materials, which is stored as an outgoing Surface Light Field (SLF), as a method to improve upon this problem. In addition, suitable SLF compression and computation techniques are presented and the quality of rendered images is demonstrated.

1. Introduction

Virtual prototypes are a de-facto standard in manufacturing industries like the automotive industry. They are used, e.g., for performing virtual crash tests or for making decisions concerning ergonomic constraints. In the area of interior and material design the use of Virtual Reality is still uncommon since the designers do not trust the computer generated images despite the great advances in computer graphics in the last decades. And indeed, errors can be and usually are introduced during the whole modelling and rendering process. Sources of errors range from inaccurate geometric modelling and parameterization, which results in unrealistic application of material properties, over inadequate modelling of the reflectance and material properties, to physically invalid rendering algorithms (especially if interactive rendering in a Virtual Reality (VR) environment is required).

In this paper we address some of these problems in the context of interactive material preview. The first problem is the lack of realistic material representations that can reproduce all the subtle effects that define the unique look and feel of a material. For this purpose we advocate usage of measured material properties, namely BTFs [4]. The second problem is the physical invalidity of interactive rendering algorithms that typically rely on artificial illumination representations like point lights and neglect indirect illumination (see Figure 1 for a comparison of the achievable realism). State-of-the-art approaches for interactive computation and rendering of global illumination solutions [8] are currently far from practical application. Existing approaches for rendering precomputed global illumination solutions in Virtual Reality are typically limited to Radiosity solutions which assume completely diffuse materials. As a consequence, the material's

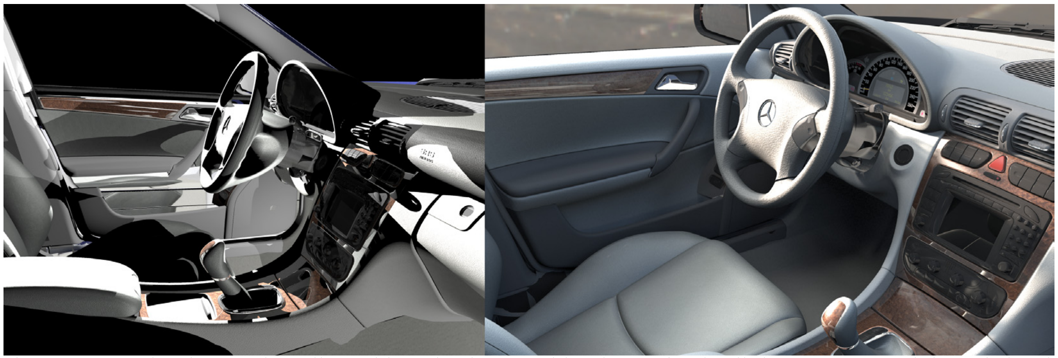

Figure 1: Despite the good quality achieved using combinations of point lights, BTF materials and simple shadows (left), the achievable realism is clearly superior using area lights and global illumination (right).

unique look and feel as visible in Figure 1 cannot be reproduced using such approaches.

Therefore, inspired by the approach of Chen et al. [1], we propose real-time surface light field (SLF) rendering as a suitable method for VR applications. In contrast to Radiosity solutions, SLFs, which encode the outgoing radiance of surface points for arbitrary view directions given static lighting, support complex material properties like BTFs which leads to drastically improved realism in rendered images. Compared to real-time raytracing approaches, SLF rendering requires a single, commodity PC only.

2. Surface Light Field Compression and Rendering

Precomputing a global illumination solution and storing it in a SLF is an offline task which takes a large amount of time (hours to days). The result is a set of images as shown in Figure 2: one for each sampled local view direction.

While precomputation of the SLF is an offline task without tight timing constraints, visualization of the SLF data has to be done in real-time to be usable for Virtual Prototyping. This requirement has two implications: first, the large amount of memory required by the SLF data (storing the SLF data shown in Figure 2 as losslessly compressed 16 bit OpenEXR images requires about 220 MB) has to be reduced by applying suitable compression methods. Second, the data has to be transferred into a data format which can efficiently be used for rendering exploiting the features of modern graphics processing units (GPUs).

A large number of compression algorithms have been proposed for image-based rendering data (see e.g. [9][1][4]) all of which have specific advantages and disadvantages. For SLF rendering, Light Field Mapping [1], which uses principal component analysis (PCA) to compress SLF data, has been the state of the art. Nevertheless, as shown in the related area of BTF rendering [4], PCA compression is not optimal w.r.t. the trade-off between compression ratio, reconstruction quality and run-time efficiency. Instead, a combination of vector clustering and PCA called local PCA [4] was found to yield a better trade-off.

Applying local PCA to the data set of Figure 1 (using a high-quality setting of 128 clusters and 4 PCA components per cluster) reduces the storage requirements to about 9 MB – a compression ratio of 1:24. This process takes about two hours on a single, commodity PC with 2.4 GHz and 1 GB of main memory. Distributing the work onto several CPUs reduces the amount of required time approximately linearly with the number of employed CPUs since the overall time is dominated by assigning the per-pixel SLF data to appropriate clusters which requires no synchronization between the individual CPUs. In our example the amount of time reduces to about 35 minutes when using four CPUs instead of a single one.

Efficiently rendering the compressed data on current GPUs can be achieved using vertex and fragment shaders, which reconstruct the outgoing luminance for a specific surface position and view direction from the SLF‘s local PCA representation. Resampling the view directions to correspond to the directions of an 8×8 pixels parabolic map [3] allows for efficient interpolation of the sampled view directions using graphics hardware. Interpolation in the spatial domain is achieved by manually computing mip-map levels and hardware-supported anti-aliasing.

Results of our SLF rendering algorithm for a test scene consisting of a box with colored and partially mirroring walls and an area light are shown in Figure 3. The seat cushion is covered

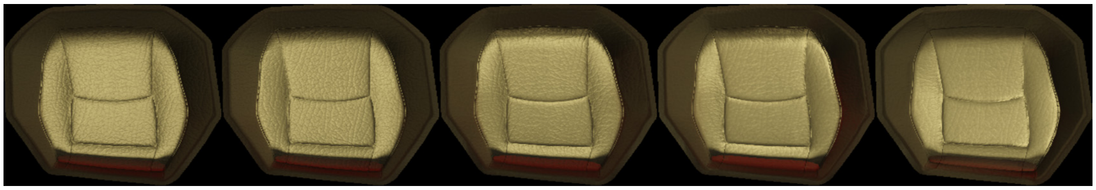

Figure 2: Five slices of the seat cushion SLF

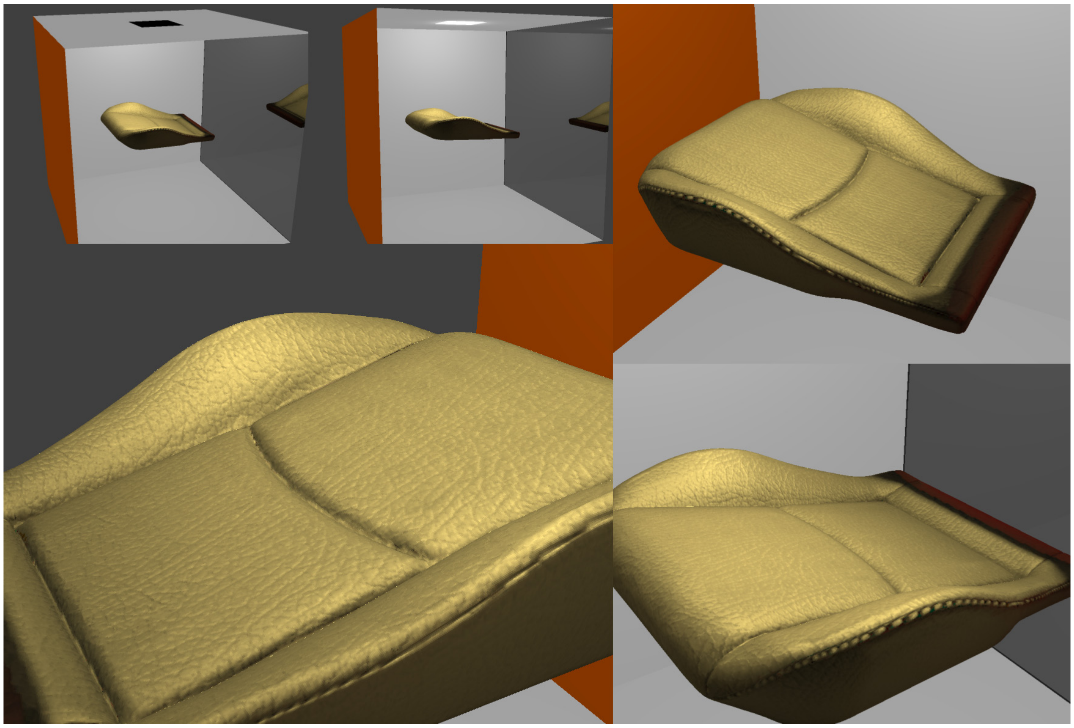

Figure 3: Results for a test scene. Small images: overview of the scene with colored and partially mirroring walls, and the seat cushion placed in the center. Remaining images: different views of the seat cushion. Some parts of the object appear very dark due to missing tone mapping functionality.

with a measured BTF of a light-colored synthetic leather. The images shown are computed as 16 bit high-dynamic range images [2] at about 30 fps on a NVidia Geforce 6800 Ultra graphics card with 256 MB memory. Due to missing appropriate tone-mapping functionality, which is required to map the 16-bit values to 8 bit values employed for display on monitors and in standard image formats, some parts of the seat cushion like the downwards-pointing sides appear overly dark. Adding such functionality remains for future work.

Since the major goal of this work is to achieve high-quality, real-time visualization of global illumination solutions for Virtual Reality applications, we decided to integrate our rendering algorithm with the open source scene graph OpenSG [6]. The SLF material is realized by a ChunkMaterial which is assigned a number of textures (as TextureChunks) and GLSL vertex and fragment programs (as SHLChunks).

3. Fast Surface Light Field Generation

Computing an accurate outgoing SLF is a time-consuming task and can take from several hours to days on a single PC. For applications like material preview in rapid prototyping such a delay is hardly acceptable. Therefore, we propose a hybrid technique which enables fast generation of an outgoing SLF for a given BTF by sacrificing self-interreflections. The idea is to precompute an accurate incoming SLF for the object of interest. The outgoing SLF of the object covered with specific BTF materials is then generated by convolving the incoming SLF with the material at each sample point on the object. Self-interreflections are neglected this way because the incoming radiance is set to zero where the ray hits another part of the object surface. Therefore, we must assume not too specular materials such that the visual effect of self-interreflections can be ignored safely.

Computing the outgoing SLF $L_o(\mathbf{x},\omega_o)$ for sample points $\mathbf{x}$ on the surface given the incoming SLF $L_i(\mathbf{x},\omega_i)$ and material properties $f_x^*(\omega_i,\omega_o)$ comprises evaluating the reflection integral

$L_o(\mathbf{x},\omega_o)=\int_{\Omega_i} f_x^*(\omega_i,\omega_o)L_i(\mathbf{x},\omega_i)d\omega_i$, which can be interpreted as a convolution of material properties and incoming SLF. Evaluating this integral in a naive way corresponding to the projection of the factors $L_i(\mathbf{x},\omega_i)$ and $f_x^*(\omega_i,\omega_o)$ onto a directional basis of the incoming hemisphere Ω_i is hardly a speed up. E.g., generating an outgoing SLF of a spatial resolution of 1024^2 texels and a hemicube of 768 outgoing directions corresponds to evaluating the above integral about 800 million times. If the incoming SLF is also sampled with 768 directions more than 1.2 trillion arithmetic operations have to be executed per color channel.

Projecting onto more appropriate bases (Spherical Harmonics [7] or Wavelets [5]) improves the situation significantly by linear or non-linear approximation of the coefficient vector using only 25-100 non-zero coefficients per factor leading to a speed up by a factor of 10-20.

4. Conclusions

In this paper we have presented an approach for rendering global illumination solutions of scenes with complex material properties in real-time using surface light fields, which allows the user to move in a static scene at real-time frame rates. We presented a compression and real-time rendering scheme for SLF data and its integration into the scene graph OpenSG, and proposed an improved SLF generation method.

Obviously our approach has both advantages and drawbacks. Compared to other existing methods for visualization of global illumination solutions our approach achieves high-quality results in real-time on a single PC while supporting very general material representations, which makes it applicable to a wide variety of design studies. Drawbacks are the significant memory requirements, which depend mainly on the spatial resolution of the SLF data. Therefore, this approach is currently limited to rather small objects or scenes with smoothly varying illumination. Developing even more efficient compression algorithms will remove this limitation, which we plan as future work. Compared to real-time raytracing and precomputed radiance transfer approaches, SLF rendering is less flexible since both the geometry and the lighting need be static. Yet, as shown in the previous section, materials can be replaced very efficiently, which is usually much more important for design studies than e.g. changing the geometry. As a result, SLF rendering represents a reasonable, easy-to-implement solution suitable for currently existing, commodity hardware configurations.

The work presented within this paper was done as part of the EU project RealReflect (IST-2001-34744). We thank our project partners from INRIA Grenoble and DaimlerChrysler AG for providing the incoming SLFs and the models underlying figures 1-3 respectively.

References

[1] W.-C. Chen *et al.*, Light Field Mapping: Efficient Representation and Hardware Rendering of Surface Light Fields, Proc. of SIGGRAPH 2002, pp. 447-456, 2002.

[2] P. Debevec *et al.*, High-Dynamic Range Imaging, SIGGRAPH 2004 Course #13, 2004.

[3] W. Heidrich and H.-P. Seidel, View-Independent Environment Maps, Proc. of Eurographics/SIGGRAPH Workshop on Graphics Hardware '98, pp. 39-45, 1998.

[4] G. Müller *et al.*, Acquisition, Synthesis and Rendering of Bidirectional Texture Functions, *Computer Graphics forum*, **24(1)**, 2005, 83-109

[5] R. Ng *et al.*, All-Frequency Shadows using Non-Linear Wavelet Lighting Approximation, *ACM Transactions on Graphics*, **22(3)**, pp.376-381, 2003

[6] www.opensg.org

[7] P.-P. Sloan *et al..*, Precomputed Radi-ance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments, Proc. of SIGGRAPH 2002, pp. 527-536, 2002.

[8] I. Wald *et al.*, Realtime Ray Tracing and its use for Interactive Global Illumination, Eurographics State of the Art Reports, 2003

[9] D. Wood *et al.*, Surface Light Fields for 3D Photography, Proc. of SIGGRAPH 2000, pp. 287-296, 2000.