1
Introduction

One of the primary goals in Computer Graphics (CG) is photorealistic rendering. CG tries to solve a well-defined problem: given the geometry, material, lighting and shading information for a virtual scene, create an image that looks as close as possible to one that a camera would capture of a real version of the described scene. However, despite all the advancements in more classical areas of CG, it is still hard to compete with images of real scenes.

This limitation on quality is inherent to geometry-based rendering processes. Even when high quality rendering techniques, such as ray tracing, are used to mimic the real world physics involved in object's illumination interaction, still the synthesized image's quality is limited to the model description and the approximations of physical models describing light reflection and transmission. For instance, realistic rendering of hair and skin remains a challenge to CG's community until nowadays.

Image-based rendering (IBR) is a powerful alternative to traditional geometry-based techniques for image synthesis. The main idea is to use images rather than geometry as the main primitives for rendering novel views. Computer vision (CV) algorithms are used to extract a model from existing images and videos, a process called image-based modeling. Model and images

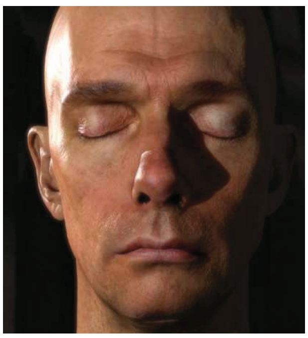

1.1(a): Skin rendering example (image courtesy of Craig Donner and Henrik Wann Jensen [12].

1.1(b): Hair rendering example [25].

Figure 1.1: Skin and hair rendering are examples of challenges to purely geometry-based approaches.

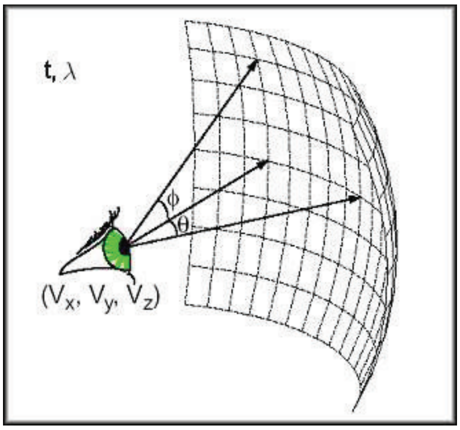

Figure 1.2: Plenoptic function.

then work as input for rendering methods that can take advantage of real-world samples of the scene's radiance and lighting properties, potentially leveraging synthesized images higher visual accuracy.

Ideally, the obtained model describing a scene would be equivalent to the $plenopticfunction$ [3]:

$$P_7(\theta, \phi, \lambda, t, V_x, V_y, V_z) \tag{1-1}$$

To measure this 7D function, one can imagine placing a pinhole camera's center at every 3D location (V_x, V_y, V_z) at every possible angle (θ, ϕ), for every wavelength λ, at every time instant t. Indeed, image-based modeling and rendering can be defined as a means of sampling the plenoptic function, representing it in a compact and useful manner, compressing all this data and rendering novel views from it. In other words, IBR can be viewed as a set of techniques to reconstruct a continuous representation of the plenoptic function using observed samples as input.

The generation of novel views from acquired images is motivated by several applications in computer games, sports broadcast, TV advertising, cinema and entertainment industry. In case of ambiguity in a soccer game, for instance, many input views may be used to synthesize a new view at a different angle to help referees inspect for events such as fouls or offsides.

The goal of this work is to develop a method of rendering synthetic novel views of a scene captured by real cameras, capable of generating visually accurate images at interactive rendering rates. Depth images, i.e. color images along with their dense depth maps, are used as the sole input of our algorithm. Results demonstrate the efficiency and quality of the proposed system.

In short, our method has the following characteristics:

– Real-time performance for virtual view synthesis.

– Visually-accurate synthesized views: rendered images have quality comparable to the input photos, with few visible artifacts.

- Transitions between views are smooth: visible changes when changing from one viewpoint to another are not easily noticeable.

The main contribution of this work is an IBR method running entirely on the GPU. That not only guarantees good performance, but also leaves the CPU free to perform other tasks like input video decoding, for instance. An additional contribution is that our method depends solely on depth images as input, without any pre-processing stage of the input.

This document is organized as follows. Chapter 2 presents a review of related research in IBR, delineating reasons for our design choices. Chapters 3 and 4 present the basics on depth image representation and how images can be composited for virtual synthesis. Chapter 5 describes how the proposed method modifies existing techniques to better suit rendering at the GPU. Results and performance numbers are depicted in Chapter 6, where publicly available datasets are used to test our algorithm. Finally, in Chapter 7 we conclude and present future work directions to further improve and extend our architecture.