SURFACE PLENOPTIC FUNCTION: A TOOL FOR THE SAMPLING ANALYSIS OF IMAGE-BASED RENDERING

Cha Zhang and Tsuhan Chen

Dept. of Electrical and Computer Engineering, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213, USA
{czhang, tsuhan}@andrew.cmu.edu

ABSTRACT

In this paper, we introduce surface plenoptic function (SPF) as a tool for the sampling analysis of image-based rendering (IBR). SPF is a plenoptic function defined on the object surface. It can be mapped onto the normal IBR representation via a coordinate transform. By assuming some properties of the SPF, we can provide insightful analysis on the sampling of the IBR representation, for both uniform and non-uniform methods.

1. INTRODUCTION

Image-based rendering (IBR) has attracted a lot of attentions recently. Various representations of IBR have been proposed, such as plenoptic modeling [1], lightfield [2], lumigraph [3], concentric mosaic [4], unstructured lumigraph [5] etc. IBR has many advantages over the traditional model-based rendering. It requires very little or even no geometrical information about the scene to realistically render it. The rendering speed of IBR is linearly proportional to the image size, but not the scene complexity. Nevertheless, often a huge number of images have to be taken in order to avoid aliasing or ghosting during the rendering. It is therefore extremely important to know how many images are enough to capture a scene. When the IBR is captured uniformly, as in lightfield and concentric mosaics, we could solve the problem by analyzing the Fourier spectrum of them and applying the traditional sampling theory. If we are allowed to sample the scene non-uniformly, we need to decide along the capturing path where to capture more and where to capture less, such that the optimal rendering results can be achieved.

Work has been done on the spectral analysis of uniformly sampled IBR. In [6], Chai et al. proposed a method to study the frequency spectrum of lightfield. Assuming Lambertian surface and no occlusion, they derived that the spectral support of a lightfield signal is bounded by the minimum and maximum depths of objects in the scene only, no matter how complicated the scene is. Knowing the bound of the support, the maximum sampling density can be achieved by compacting copies of the spectral support in the frequency domain. Marchand-Maillet and Vetterli [7] extended the work for scenes with functional surfaces. Lin and Shum [8] performed sampling analysis on both lightfield and concentric mosaic with scale-space theory under constant-depth assumption. The bounds are derived from the aspect of geometry and based on the goal that no "spurious detail" should be generated during the rendering (referred as the causality requirement). All the above work was successful in improving our knowledge about IBR sampling. However they are not easily extendable to analyzing more complex scenes, such as scenes with non-Lambertian surface or occlusions, not to say the non-uniform sampling of IBR.

In this paper, we propose a new method to parameterize the plenoptic function for IBR sampling analysis. We call the new representation *surface plenoptic function* (SPF). There exists a unique mapping between the SPF and the IBR representation. We show that when this mapping is well defined, we are able to do the spectral analysis for generic scenes including occluded or non-Lambertian ones. We can also provide some insightful views to the problem of non-uniform sampling.

The paper is organized as follows. Section 2 introduces the definition of SPF and its relationship to general IBR representations. We use SPF to analyze the Fourier spectrum of IBR in Section 3. Section 4 presents the usage of SPF to non-uniform IBR sampling. Conclusions are given in Section 5.

2. THE SURFACE PLENOPTIC FUNCTION

Any lightfield has its source. Light rays can be either emitted from some light source (e.g., the Sun), or reflected from some object surface. Let the entire surface of all the light sources and objects be **S**. We can always trace a light ray in the free space back to a point on **S**. In the 3D world where radiance does not change along a line unless blocked, the 7D plenoptic function can be reparameterized to 6D including time (1D), wavelength (1D), point on the surface **S** (2D), and azimuth and elevation angles (2D) the light ray is emitted. We name this 6D function as the *surface plenoptic function* (SPF). Notice that the reparameterization does not lose any information given that radiance does not change along its path.

The reason we introduce SPF, is that under the same condition, it will have the same dimensionality as commonly used IBR representations. For example, when time and wavelength is ignored, SPF is a 4D function, the same as lightfield. When cameras and the viewer are constrained on a plane, SPF reduces to 3D (eliminating the elevation angle), as in concentric mosaic. The property of same dimensionality provides us the possibility to easily map between SPF and IBR representations. Such a mapping depends on both the scene geometry and the camera trajectory, but not the surface property such as BRDF. If we have some knowledge about the scene, in other words, if we know some property about the SPF, related property can be derived for the IBR representation. More importantly, the mapping does not require the scene to have no occlusions or Lambertian surface, which is very attractive for IBR sampling analysis.

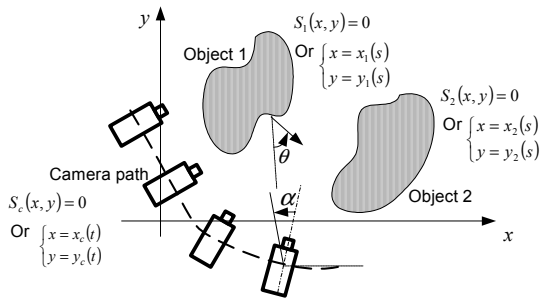

Figure 1 2D SPF and general IBR capturing.

Without loss of generality, we use the 2D world as example throughout this paper for conciseness. The conclusions drawn in this paper are easy to extend to the 3D world. In the 2D world, surface of objects/light sources is described with curves. Ignoring time and wavelength, the SPF is 2D: one dimension for describing a point on a curve, the other for illustrating the direction of the light ray emitted/reflected. An example scene is shown in Figure 1. The surface can be represented by either $S_i(x,y)=0$ or $\{x=x_i(s), y=y_i(s)\}$, where s can be the arc length, i be the index for different objects. For a certain object i, define its SPF as:

$$l_i(s,\theta) \text{ on the curve } \{x=x_i(s), y=y_i(s)\} \qquad (1)$$

where $0 \le \theta < 2\pi$ is the direction of the light ray; $l_i(s,\theta)$ is the radiance of the light ray that can be traced back to the surface point determined by s with direction θ.

In order to capture the plenoptic function or surface plenoptic function, existing IBR approaches align cameras on a path/surface and take images for the scene. For example, cameras are placed on a plane in lightfield, and on a circle in concentric mosaic. In the 2D world, 2D lightfield has cameras on a line, while 2D concentric mosaic has cameras on a circle. In general, the cameras can be put along an arbitrary curve, as is shown in Figure 1. Let the camera path be $S_c(x,y)=0$ or $\{x=x_c(t), y=y_c(t)\}$, where t is the arc length. The image pixels can be indexed by the angle between the captured light ray and the optical axis, as is represented by α in Figure 1. Due to the correspondence between the light rays emitted/reflected from the scene surface and those captured, a coordinate transform between (s,θ) and (t,α) can be well defined. To analyze the sampling problem for the IBR representation, we may first assume some properties on the SPF, and then transform the SPF to the IBR representation, hoping that some related properties can be derived. In the next two sections, we will show examples where we can apply the above strategy to the uniform and non-uniform sampling of IBR.

3. UNIFORM SAMPLING – SPECTRAL ANALYSIS

For IBR uniform sampling, we need to find its Fourier spectrum so that we may apply the traditional sampling theory. Due to the page limit of the paper, we give an example on analyzing the lightfield representation.

3.1. 2D Lightfield Analysis

As shown in Figure 2, 2D lightfield is parameterized by two parallel lines, indexed by t and v, respectively. The t line is the camera line, while the v line is the focal line. The distance between the two lines is f, which is the focal length of the cameras. The SPF, for ease of analysis, is defined based on a global coordinate of θ.

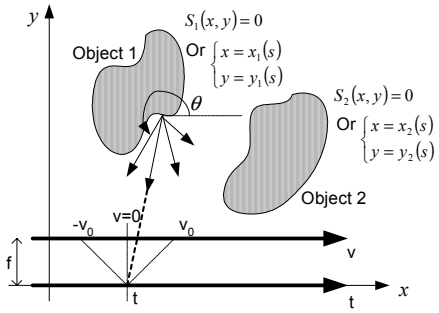

Figure 2 The lightfield parameterization.

It is easy to show that a light ray indexed by pair (t,v) has the following algebraic equation:

$$fx - vy - ft = 0 \qquad (2)$$

Notice that the focal line is indexed locally with respect to where the camera is.

The relationship between lightfield $l_c(t,v)$ and the SPF $l_i(s,\theta)$ is as follows. For the same light ray emitted/reflected from a surface point, it must be captured at the corresponding angle. That is:

$$\tan(\theta) = f/v \text{ or } \theta = 3\pi/2 - \tan^{-1}(v/f) \qquad (3)$$

where $-v_0 \le v \le v_0$ and $2\tan^{-1}(v_0/f)$ tells the field of view (FOV). If the cameras have a limited FOV, which often happens in real-life systems, Equation (3) can be linearized as:

$$\theta = 3\pi/2 - v/f \qquad (4)$$

Another constraint is that the light ray (t,v) can be traced back to a cross point on the object surface, whose arc length s can be obtained through solving:

$$\begin{cases} x=x_i(s), y=y_i(s) \\ fx - vy - ft = 0 \end{cases} \qquad (5)$$

When multiple objects exist in the scene or some objects occlude themselves, Equation (5) may have multiple answers. We have to figure out which cross point is the closest to the cameras. The closest point will occlude all the others. This may make scenes with occlusions hard to analyze. However, for simple occluded scenes this is still doable [9]. We next show a simple example for analyzing a scene at constant depth. More complex examples can be found in [9].

3.2. An example

The simplest scene we can have for lightfield is one at a constant depth. Let the SPF of the scene be $l(s,\theta)$, whose Fourier transform be $L(\Omega_s,\Omega_\theta)$. The object surface can be described by:

$$\{x=x_0(s)=s, y=y_0(s)=d_0\} \qquad (6)$$

We can solve Equation (5) without concerning about occlusion:

$$fs - vd_0 - ft = 0 \Rightarrow s = vd_0/f + t \qquad (7)$$

The lightfield spectrum can be easily derived as:

$$L_c(\Omega_t,\Omega_v) = fL(\Omega_t, d_0\Omega_t - f\Omega_v)e^{j\frac{3\pi}{2}(d_0\Omega_t - f\Omega_v)} \qquad (8)$$

We can see that the spectrum of the lightfield at constant depth is a rotated version of the SPF spectrum, with some constant factor in magnitude and some shift in phase. The rotation angle is determined by the scene depth d_0 and the focal length f.

We may give different assumptions about the SPF. For example, if the object surface is Lambertian, which means the light rays from the same surface point have identical radiance, we may let $L(\Omega_s,\Omega_\theta)=L_s(\Omega_s)\delta(\Omega_\theta)$. Therefore,

$$L_c(\Omega_t,\Omega_v)=fL_s(\Omega_t)\delta(d_0\Omega_t-f\Omega_v)e^{j\frac{3\pi}{2}(d_0\Omega_t-f\Omega_v)} \quad (9)$$

This is a tilted line in the (Ω_t,Ω_v) space, which is the same conclusion as that in [6].

When the object surface is non-Lambertian, the spectrum of the SPF expands along Ω_θ. If the scene surface is not very specular, thus light rays from the same surface point change slowly about the direction θ, we may assume that the SPF is band-limited along Ω_θ. Therefore, we may let $L(\Omega_s,\Omega_\theta)=L(\Omega_s,\Omega_\theta)I(\Omega_\theta)$, where

$$I(\Omega_\theta)=\begin{cases}1, & \text{if } |\Omega_\theta|<B_\theta,\\ 0, & \text{otherwise.}\end{cases} \quad (10)$$

and B_θ defines the bandwidth. Consequently:

$$L_c(\Omega_t,\Omega_v)=fL(\Omega_t,d_0\Omega_t-f\Omega_v)e^{j\frac{3\pi}{2}(d_0\Omega_t-f\Omega_v)}I(d_0\Omega_t-f\Omega_v) \quad (11)$$

The spectrum is also tilted, but this time it has a finite width $2B_\theta/\sqrt{d_0{}^2+f^2}$ perpendicular to the tilted spectrum (or $2B_\theta/d_0$ horizontally) because of the indicator function.

The above analysis is illustrated in Figure 3. A scene at constant depth has two sinusoids (different frequency) pasted on it as texture, as shown in Figure 3 (a). Figure 3 (b) is the epipolar image (EPI) when the scene is Lambertian. Figure 3 (c) is its Fourier transform. The spectrum has several peaks because the texture on the scene object is pure sinusoids. It basically lies on a tilted line with some small horizontal and vertical windowing artifact that is due to the truncation of the range of s and θ. We ignored the windowing artifacts in our analysis for simplicity. Figure 3 (d) shows the EPI for a non-Lambertian case at the same depth. It can be seen that because of the non-Lambertian property, its Fourier transform in Figure 3 (e) is expanded towards the normal direction of the tilted line.

4. NON-UNIFORM SAMPLING – SAMPLING DENSITY FUNCTION

It is not straightforward how to extend the above method to the analysis of non-uniform sampling. In this section, we define sampling density function as another tool to solve the problem together with SPF.

4.1. Stochastic Sampling and the Sampling Density Function

Stochastic sampling is a Monte Carlo technique [11] that takes samples at irregular locations rather than at regular spaced locations. A simple implementation of stochastic sampling is to select the sampling locations based on a certain distribution. In this paper, we refer the probability density function of such a distribution as *sampling density function* (SDF).

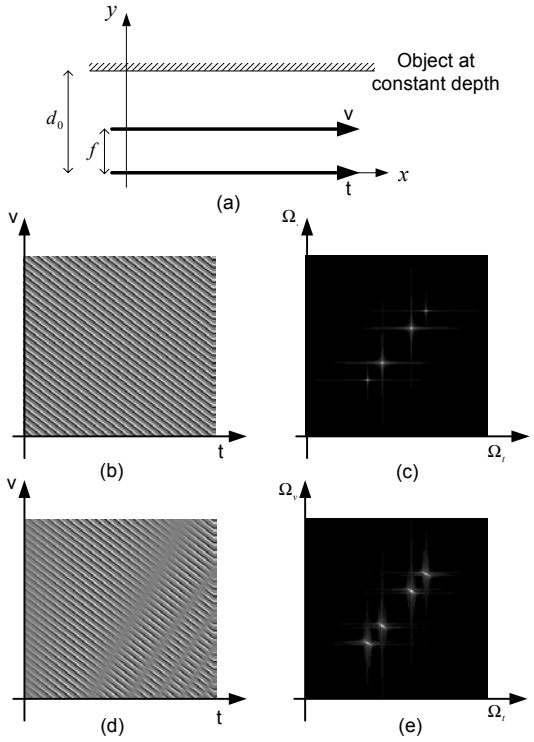

Figure 3 Lambertian and non-Lambertian scenes at constant depth.

Although stochastic sampling itself is sometimes referred as non-uniform sampling in the literature, we define the uniformity of sampling in a different way. A signal is *uniformly sampled* if and only if its SDF is uniform. This is intuitively true because when the SDF is uniform, all the locations have equal probability to be chosen as a sample.

In the following analysis on non-uniform sampling of IBR, we will show that the SDF of the IBR representation tend to be non-uniform, even if the optimal SDF of the surface plenoptic function is uniform.

4.2. Method and Example

The SDFs of the SPF and the IBR representation is related by the mapping between them. Therefore, one can transform one of the SDF to the other. Let the SDF of the SPF be $S_{spf}(s,\theta)$, and that of the IBR representation be $S_{ibr}(t,\alpha)$. If the mapping between SPF and IBR representation has an explicit form:

$$(t,\alpha)=M(s,\theta) \quad (12)$$

under the condition that $M(s,\theta)$ is monotonically increasing or decreasing, we have:

$$S_{ibr}(t,\alpha)=\frac{S_{spf}(s,\theta)}{|J|} \quad (13)$$

where $J=\begin{vmatrix}\partial t/\partial s & \partial t/\partial\theta\\ \partial\alpha/\partial s & \partial\alpha/\partial\theta\end{vmatrix}$ is the Jacobi Determinant of the transform.

Unfortunately, the mapping between SPF and IBR representation rarely has an explicit form, not to say monotonically in-

creasing or decreasing. We solve the problem by Monte Carlo method. Assume that the SDF of the SPF is known. We simply perform stochastic sampling on the light rays from SPF, and record their corresponding coordinates in the IBR representation. The SDF of the IBR representation is achieved by counting the frequencies of the light rays falling into different bins.

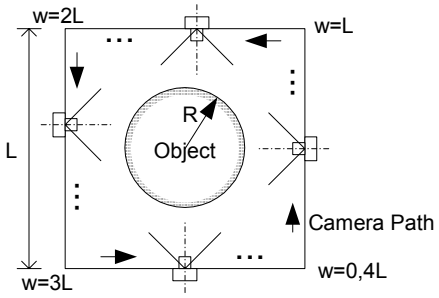

Figure 4: An example IBR scene.

An example 2D scene is shown in Figure 4. The object is a perfect circle. We assume that the optimal SDF for its SPF is uniform. The object is captured with a standard inward-looking lightfield. Cameras are placed on a concentric box with their optical axes perpendicular to the camera path. In the following experiment, we set object radius R = 1, capturing box side length L = 4, field of view of the cameras FOV = 90°. The camera path is indexed with w, whose value is from 0 to 4L, divided into 1000 bins. The FOV is divided into 200 bins. One billion (10^9) light rays are randomly drawn in our Monte Carlo experiment.

Figure 5 (a) shows the SDF of the IBR representation of the setup in Figure 4. The horizontal axis is the camera path w, the vertical axis is the pixel index or bin index of the field of view. A brighter pixel means denser samples. It can be seen that the SDF is non-uniform. Notice that in practice when we take an image, the sampling pattern along the pixel index axis is fixed (most likely uniform). Therefore, we are only allowed to sample the camera path w non-uniformly. The marginal distribution of the SDF along the camera path is shown in Figure 5 (b). The vertical axis is the number of light rays falling into the corresponding w bin. It tells us that we should capture more images around the center of each side of the capturing box.

3.4. Discussions

Transforming the SDF of the SPF to that of the IBR representation is generally doable with Monte Carlo method, provided that the former SDF and the geometry of the scene are known. Given the SDF of the IBR representation, it is possible to plan the view positions before capturing, which will yield optimal rendering quality.

In practice, the scene geometry is usually unknown. An IBR viewer often assumes some rough geometry of the scene, e.g., a constant-depth plane. In this case, the above method is still applicable, but the source SPF has to be defined on the rendering surface instead of the true object surface.

The difficulty lies at finding the SDF of the SPF. A uniform SDF of the SPF is unlikely to be optimal in practice, given that the scene surface can be specular, and the rendering geometry can be arbitrary chosen. In [10], we explore algorithms that can approach similar results (near-optimal non-uniform sampling) but are more practical.

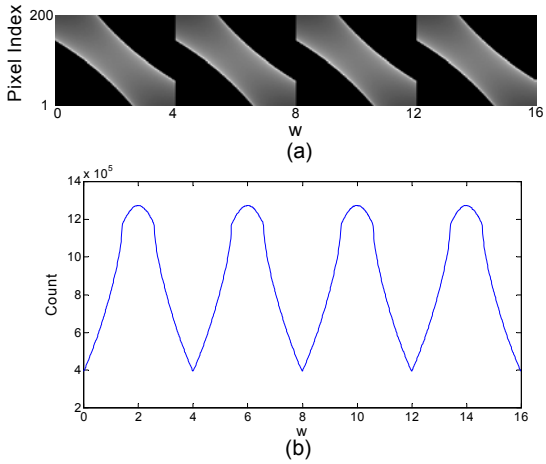

Figure 5: (a) SDF of the IBR representation in Figure 4. (b) Marginal distribution along camera path w.

5. CONCLUSIONS

We introduced surface plenoptic function in this paper as a tool to analyze the sampling problem of image-based rendering. We showed its effectiveness for both uniform and non-uniform sampling analysis. Although we may need to assume certain properties for the SPF, we believe SPF is still a useful tool for IBR sampling analysis.

REFERENCES

[1] L. McMillan and G. Bishop, "Plenoptic modeling: an image-based rendering system", *Computer Graphics (SIGGRAPH'95)*, pp. 39-46, Aug. 1995.

[2] M. Levoy and P. Hanrahan, "Light field rendering", *Computer Graphics (SIGGRAPH'96)*, pp. 31, Aug. 1996.

[3] S. J. Gortler, R. Grzeszczuk, R. Szeliski and M. F. Cohen, "The Lumigraph", *Computer Graphics (SIGGRAPH'96)*, pp. 43-54, Aug. 1996.

[4] H.Y. Shum and L.-W. He. "Rendering with concentric mosaics", *Computer Graphics (SIGGRAPH'99)*, pp.299-306, Aug. 1999.

[5] C. Buehler, M. Bosse, L. McMillan, S. Gortler and M. Cohen, " Unstructured Lumigraph Rendering", *SIGGRAPH 2001*, pp. 425-432.

[6] J.X. Chai, X. Tong, S.C. Chan and H. Y. Shum, "Plenoptic sampling", *Computer Graphics (SIGGRAPH'00)*, pp.307-318, July 2000.

[7] D. Marchand-Maillet and M. Vetterli, "Sampling Theory for Image-Based Rendering", Master thesis, EPFL, Apr. 2001.

[8] Z. C. Lin and H. Y. Shum, "On the Number of Samples Needed in Light Field Rendering with Constant-Depth Assumption", *Proc. CVPR 2000*.

[9] C. Zhang and T. Chen, "Surface Plenoptic Function: Spectral Analysis for Image-Based Rendering", submitted to *CVPR 2003*.

[10] C. Zhang and T. Chen, "Active Image-Based Rendering", submitted to *CVPR 2003*.

[11] J. H. Halton, "A Retrospective and Prospective Survey of the Monte Carlo Method", *SIAM Rev.* pp. 1-63, Vol. 12, No. 1, Jan. 1970.