

Efficient Light Field Based Camera Walk

Aviral Pandey

Institute of Technology
Banaras Hindu University
Varanasi, India 221005

aviral_itbhu@yahoo.com

Biswarup Choudhury

Computer Science & Engg. Dept.
IIT Bombay
Mumbai, India 400076

biswarup@cse.iitb.ac.in

Sharat Chandran

Computer Science & Engg. Dept.
IIT Bombay
Mumbai, India 400076

sharat@cse.iitb.ac.in

Abstract

The light field rendering method is an interesting variation on achieving realism. Once authentic imagery has been acquired using a camera gantry, or a handheld camera, detailed novel views can be synthetically generated from various viewpoints.

One common application of this technique is when a user “walks” through a virtual world. In this situation, only a subset of the previously stored light field is required, and considerable computation burden is encountered in processing the input light field to obtain this subset. In this paper, we show that appropriate portions of the light field can be cached at select “nodal points” that depend on the camera walk. Once sparsely and quickly cached, scenes can be rendered from any point on the walk efficiently.

1 Introduction

The traditional approach for “flying” through scenes is by repeated rendering of a three-dimensional geometric model. One well known problem with “geometry-based” modeling is that it is very difficult to achieve photo-realism due to the complex geometry and lighting effects present in nature. A relatively newer approach is Image-Based Rendering (IBR) [?], which uses a confluence of methods from computer graphics and vision. The IBR approach is to generate novel views from virtual camera locations from pre-acquired imagery. Synthetic realism is achieved, so to speak, using real cameras.

Light Field Rendering (LFR) [8] (or Lumigraphs [6]) is an example of IBR. The approach here is to store samples of the plenoptic function[1] which describes the directional radiance distribution for every point in space. The subset of this function in an occlusion-free space outside the scene can be represented in the form of a four-dimensional function. The parameterization scheme is shown in Figure 1. Every viewing ray from the novel camera location C passing through the scene is characterized by a pair of points (s, t) and (u, v) on two planes. By accessing the previously acquired radiance associated with this four tuple, we are able to generate the view from C . In order to view a scene from any point in surrounding space, six light slabs are combined so that the six viewpoint planes cover some box surrounding




Figure 1. Two plane parameterization

the scene.

1.1 Statement of the Problem

The key to LFR lies in re-sampling and combining the pre-acquired imagery. In a typical walk-through situation, a person is expected to walk along a trajectory in three space and “suitably” sample the light field. The problem we pose in this paper is “Given the light field on disk, and a camera walk, how fast can the scene be rendered?”.

For best results in light field based IBR, we expect that the size of the light field data-structure drastically increases with the increase in the resolution of image and the sampling density. As mentioned above, ray-tracing is performed as an intermediate step of the rendering procedure, a computationally intensive operation [11].

1.2 Contributions

By definition, one needs to store the complete light field in volatile memory for interactive rendering of the scene, whereas only a subset of this is needed for the camera walk. Prior methods do not effectively address this issue. In this paper, we show how caching the light field *suitable* for the camera walk, dramatically reduces the computational burden, as seen in Figure 11(a).

- We compute the optimal location of a sparse set of “nodal points.” The lightweight “light field” stored at these nodes is enough to render the scene from any of the infinite points – termed *query* points – on the camera path.
- The method in [11] uses homography to reduce the ray shooting computational burden in producing the image

from one query point; multiple query points are treated afresh since no notion of nodal points was required therein. We use an alternative Taylor series method for reducing the ray shooting queries.

- The correctness of our scheme is shown using a mathematical characterization of the geometry of the light field. Experimental results validate our scheme.

The rest of this paper is organized as follows. In Section 2 and Section 3, we give details of our approach. Sample results are shown in Section 4. We end with some concluding remarks in the last section.

2 Our Approach

As in the original work [8], the field of view of the query camera is expected to be identical to the cameras that generated the light field. Likewise, sheared perspective projection [8] handles the problem of aligning the plane of projection with the light-slab. Coming to the camera walk, in this section we provide the mathematical basis for the location and spacing of nodal points. For brevity, the description in this work restricts the center of projection of the camera to move along a plane parallel to the UV and the ST plane. When this condition is not satisfied, a generalization of Lemma 2.3.2 is required to compute the location of nodal points.

2.1 Fixed Direction

The algorithms in this section tell us where to place nodal points for a specific query point q assuming a fixed direction determined by some point s . This condition is relaxed later.

For motivation, consider a setup similar to the two slab setup where planes (UV and ST) are replaced by lines U and S. The query points lie on line C.

Denote Δl to be the constant distance $d[G_i, G_{i+1}]$ between two consecutive grid points on the UV plane, i.e., the distance between the acquired camera locations.

2.1.1 Fixed Direction Algorithm

Given q , our algorithm computes N_1 and N_2 as follows. Draw the ray from q to s , for a given s , to obtain q' on U. Mark points N_1' and N_2' on U at a distance $d = \frac{\Delta l}{2}$ apart on either side of q' . This determines the points N_1 and N_2 as shown in Figure 2. The radiance (in the direction of s) is presumably cached at points N_1 and N_2 . We need to make use of this cache.

Denote $\text{assoc}(p)$, where p is a point on C, to be the closest grid vertex G (on U) to the ray \overline{ps} . Suppose $\text{assoc}(q)$ is G_i . We set $L[q] = L[N_1]$ if $\text{assoc}(N_1) = G_i$, otherwise $L[q] = L[N_2]$, where $L[q]$ represents radiance at q .

In the two-dimensional case, given q , our algorithm computes four nodal points N_1, N_2, N_3 and N_4 . Draw the ray from q to s for a given s to obtain q' on UV. Now, mark four




Figure 2. N_1 and N_2 , the nodal points for q are marked such that $d[q'N_1'] = d[q'N_2'] = \frac{\Delta l}{2}$

points $(q'.u \pm \frac{\Delta l}{2}, q'.v \pm \frac{\Delta l}{2}, z_{uv})$, where $q'.u$ and $q'.v$ represent the component of q' along u and v respectively. These four points correspond to four nodal points on the camera COP (center of projection) plane.

2.1.2 Comments

Notice that if the distance d is more than $\frac{\Delta l}{2}$, as in Figure 3, we can have an incorrect value of $L[q]$. When d is as specified in the algorithm, it is easy to observe that either $\text{assoc}(N_1) = G_1$ or $\text{assoc}(N_2) = G_1$; it cannot be the case that $\text{assoc}(N_1) = G_0$ and $\text{assoc}(N_2) = G_2$. A choice less than $\frac{\Delta l}{2}$ might be suitable to maintain correctness, but will increase the number of nodal points, and hence decrease our efficiency.




Figure 3. $\text{assoc}(N_1)$ is G_0 and $\text{assoc}(N_2)$ is G_2 .

2.2 Changing the view direction

The results in this section assert that the nodal points may be chosen by arbitrarily picking any direction and applying the algorithm given in section 2.1.1. That is, the selection of nodal points is independent of the direction. The one dimensional case is given to illustrate the idea.

Lemma 2.2.1 *The set of nodal points for a given point s on S serves as the set of nodal points for all s .*

Proof: Omitted in this version. Figure 4 illustrates the idea. \square




Figure 4. The choice of nodal points is independent of the direction. $N_1'q'$ is equal to $N_1''q''$

Next, we consider the corresponding lemma for the two dimensional case.

Lemma 2.2.2 Given a query point, nodal points may be decided using any point (s, t) provided the camera planes are parallel.

Proof: As in Figure 5, let N_1, N_2, N_3, N_4 be the nodal




Figure 5. Choice of nodal points is independent of direction (2 dimensions case)

points for query point q as determined by the algorithm in section 2.1.1. Let $S_a = (s_1, t_1, z_s)$ be the intersection point of the query ray from q on the ST plane, and let $X_a = (x_1, y_1, z_u)$ be the intersection point of that ray on the UV plane. Similarly, define S_b and X_b . The proof (omitted

here) uses the relationships (equations 1 and 2) below, to prove equation 3.

$$\overline{S_b N_1} = \overline{S_a N_1} + \overline{(S_a - S_b)} \quad (1)$$

$$\overline{S_b q} = \overline{S_a q} + \overline{(S_a - S_b)} \quad (2)$$

$$\overline{S_b N_1} = k(\overline{S_b X_b} + (-\frac{\Delta l}{2}, \frac{\Delta l}{2}, 0)) \quad (3)$$

□

2.3 The Power of Nodal Points

Once nodal points are selected, there are a range of query points for which these nodal points are valid, as stated below.

Lemma 2.3.1 The nodal points N_1, N_2 of a query point q_1 are sufficient for determining the radiance of any query point in the interval $[N_1, N_2]$.

Proof:

Consider any point q_2 , between N_1 and N_2 , and presume that the nodal points as determined by our Algorithm in Section 2.1.1 are N_3 and N_4 . The lemma asserts that, for q_2 , the radiance values stored at N_1 and N_2 are sufficient.

Without loss of generality, assume q_2 to be nearer to N_2 than N_1 . We observe that either $d[N_1', \text{assoc}(N_3)] < \frac{\Delta l}{2}$ or $d[N_2', \text{assoc}(N_3)] < \frac{\Delta l}{2}$.




Figure 6. $\text{assoc}(N_3)$ is closer to N_1 than N_2 . Notice that $\text{assoc}(N_3) = \text{assoc}(N_1)$.

- **Case 1:** $d[N_1', \text{assoc}(N_3)] < \frac{\Delta l}{2}$ (shown in Figure 6) $\Rightarrow \text{assoc}(N_3) = \text{assoc}(N_1)$. So, $\text{assoc}(N_2) = \text{assoc}(N_4)$ i.e, $L[N_1] = L[N_3]$, $L[N_2] = L[N_4]$. Thus, the radiance of q_2 can be obtained from N_1 or N_2 .
- **Case 2:** $d[N_2', \text{assoc}(N_3)] < \frac{\Delta l}{2}$ (Shown in Figure 7) $\Rightarrow \text{assoc}(N_3) = \text{assoc}(N_2)$. Further, $d[q_2', \text{assoc}(N_1)] > \frac{\Delta l}{2}$ and $d[q_2', \text{assoc}(N_4)] > \frac{\Delta l}{2}$.

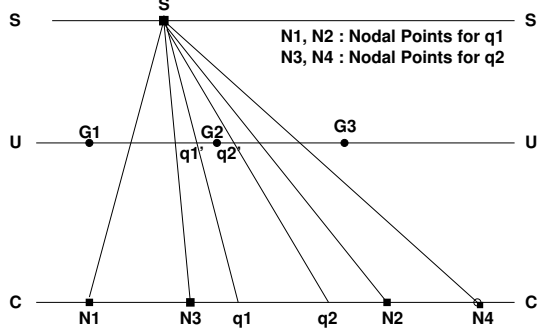


Figure 7. $\text{assoc}(N_3)$ is closer to N_2 than N_1 . Notice that $\text{assoc}(N_3) = \text{assoc}(N_2)$.

So, $L[N_3] = L[N_2]$. Thus, the radiance of q_2 can be obtained from N_2 . \square

Notice that, unlike in Figure 2, the nodal points N_1 and N_2 are not equidistant from q_2 .

Next, we consider the corresponding lemma for the two dimensional case:

Lemma 2.3.2 *The nodal points N_1, N_2, N_3 and N_4 of a query point q_1 on the camera plane are sufficient for determining radiance at any query point in the rectangular region bounded by these nodal points.*

Proof: Omitted in this version. \square

A generalization of this lemma for the case when the camera motion is *not* restricted to the plane has not been provided here. In this situation, the relevant nodal points form a truncated pyramid instead of a rectangle.

2.4 Ray Intersection

Finding the camera ray intersections with planes is costly [11] and should be avoided. In this section we show how to avoid ray intersection using the Taylor's theorem. Specifically, consider $X_1 = [X_c, Y_c, Z_c]$, which is center of projection for a virtual camera, and $I = (I_x, I_y, C_1)$, which is a point on the ST plane. Then a ray from X_1 to I , intersects UV plane at $g(X_1) = [g_x, g_y, C_2]$. Moving the COP to the location $X_2 = X_1 + \Delta X$, the change in x co-ordinate of the point of intersection with the UV plane is given by:

$$\Delta g_x = \frac{C_2 - C_1}{Z_c - C_1} \Delta X_c + \frac{(I_x - X_c)(C_1 - C_2)}{(Z_c - C_1)^2} \Delta Z_c \quad (4)$$

A similar equation is derived for Δg_y . The error associated with approximation is given by E_{g_x} ,

$$E_{g_x} = g_x(X_1 + \Delta X) - \bar{g}_x(X_1 + \Delta X) \quad (5)$$




Figure 8. Nodal Points surrounding a Camera Path

where, $\bar{g}_x(X_1 + \Delta X)$ is the Taylor's estimate. From the first order analysis

$$E_{g_x} = \frac{2(I_x - X_c)(C_1 - C_2)}{(C_1 - Z_c)^2} \Delta Z_c \quad (6)$$

A similar equation is derived for E_{g_y} . If the camera motion is on any arbitrary path in a plane parallel to the ST plane ($\Delta Z_c = 0$), then the error $E_{g_x} = 0, E_{g_y} = 0$. In addition, the computational complexity involved in calculating the new UV intersection point decreases substantially, as (4) reduces to

$$\Delta g_x = \frac{C_2 - C_1}{Z_c - C_1} \Delta X_c \quad (7)$$

which is independent of the direction of the ray. This implies that a regular camera motion results in a regular shift of intersection points.

3 Nodal Point Caching

We now have the mathematical apparatus to select the nodal points, given a camera walk. The algorithm is straightforward. Starting from the initial position on the camera path curve, we mark nodal points at a distance $\Delta x = \Delta l \times R$ where R is the ratio of the distance between the camera plane and the ST plane and the distance between the UV and ST planes. For simplicity, the nodal points are selected parallel to the u and v directions as shown in Figure 8.

The light field is cached at these nodal points. The precise computation of the light field from the nodal points can take advantage of the methods suggested in [11] or Section 2.4, instead of the original method [8]. Once the radiance at nodal points is known, Lemma 2.3.2 assures us that for any query point, we can fetch the radiance from neighboring nodal points. We denote the time taken for this operation as k_2 (Section 4.2).

An alternate way to pick nodal points is region-based, as shown in Figure 9. Any query on the camera walk in the rectangular region defined by the convex hull of the nodal points can be answered efficiently.

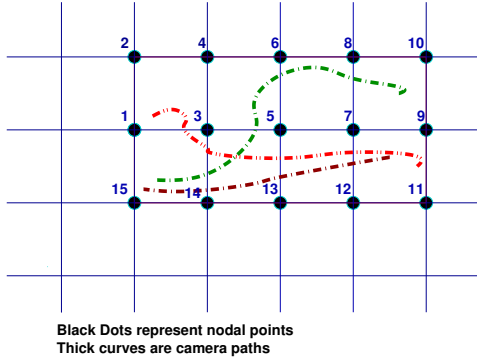


Figure 9. Nodal Points Covering Domain of Camera Motion

3.1 Quadrilinear Versus Nearest Neighbor Approximation

Once nodal points are known, nearest neighbor approximation or quadrilinear interpolation can be used. In generating views using the nearest neighbor approximation, four nodal points will suffice for all information that is needed for intermediate query points. For quadrilinear interpolation 16 nodal points are needed to provide information (radiance) for a query point.

4 Sample Results

In this section, we first provide evidence that the results obtained by the use of the method in Section 3 matches those obtained by the implementation given in [8]. Later we show that our method requires less resources.

4.1 Buddha and Dragon

Figure 10(a) shows the result obtained using the nearest neighbor approximation as suggested in [8]. Figure 10(b) shows what is obtained using the method from Section 3. The two images are identical as returned by `diff` in Unix. The virtual camera viewpoint was at $(0, 0, 3)$ and the nodal points were situated at $(\pm 0.09375, 0, \pm 0.09375, 3.00000)$. The origin was at the centre of the ST plane. The input images were those obtained using 32×32 cameras.

Identical behavior is observed when we render the dragon (Figure 10(c)); the light field for this was acquired using 8×8 cameras. Here the virtual camera was located at $(0.03, 0.02, 2.00)$ and the nodal points at $(\pm 0.5, \pm 0.5, 2)$.

4.2 Computational Advantage

We now proceed to show the computational advantage when a camera walk is introduced. As discussed earlier, advantages arise due to nodal light field caching, and avoiding ray intersection calculations. Let n be the number of query points, and p the number of nodal points. Denote k_1 to be

the time taken for ray intersection computations in the original method [8] for one query; if we use homography (from [11]), then this value is negligible.

When the input light field is densely sampled, and is at a high resolution, it may or may not be possible to place the light field in memory. We penalize access to the light field (for both methods) by the factor t in the following equation. The expected gain in our method is

$$\frac{n(t + k_1)}{(p(t + k_1) + nk_2)} \quad (8)$$

If the light field does not fit in memory, t represents disk access time. Therefore $t \gg k_1$, and the gain is approximately $\frac{n}{p}$.


4.2.1 Resource Usage

For the purpose of comparison, and to hand an advantage to the original light field implementation, we have chosen not to use the optimization in Section 2.4 in the experiments. Nevertheless, the results are worth noting. Our time results are based on an Intel Pentium IV 2.4GHz Linux based computer with 1 GB memory.

1. To simulate low RAM situations, we used only 32 MB of the 1 GB available and rendered Buddha on various camera paths located at different distances from the original camera gantry.

We note that there is considerable gain as seen in figure 11(a), where the real time taken by the two approaches is plotted with respect to the z co-ordinate of COP. The average value of p for $n = 100$ is approximately 19, so the time gain is approximately 5.26 times, which is what theoretically equation 8 promises.

2. When the memory is sufficiently large to accommodate the large light field, Figure 11(b) shows that the time taken by our method is comparable to the original method. This indicates that going to the nodal cache is not very expensive. However, the total memory that we used was even less than 2% of the memory requirement of method in [8]. This is due to the fact that the method in [8] uses $u \times v \times r$ units of memory for a $u \times v$ camera gantry with an image resolution of r .
3. To quantify disk access, we rendered Buddha on various camera paths located at different distances from the original camera gantry. Starting at $(-1.5, -1.5, z_{Cord})$ and going to $(1.5, 1.5, z_{Cord})$, the virtual camera was made to follow different zig-zag paths at different values of z co-ordinate denoted z_{Cord} . The query points were chosen randomly along these paths. In the experiments on the Buddha image, the origin of the co-ordinate system was located at the




(a) Buddha (from [8]).

(b) Buddha using our method.


(c) Dragon (from [8]).

(d) Dragon using our method.


Figure 10. Rendered images of Buddha and Dragon using our method and the traditional method are identical.



(a) The time taken by the proposed method is considerably lesser than the original method.



(b) Time taken by our approach is comparable. However, our approach uses only a very tiny fraction of the system memory.



(c) Number of disk accesses. Our approach performs significantly better.

Figure 11. All results are for 100 query points.

centre of the ST plane. Figure 11(c) shows the relative gain in terms of disk accesses. The graph shows the number of accesses to the disk storage required by various techniques when the nearest neighbor approximation is used, at $z = 3, 4, 6, 8, 12$, where z is the distance of the COP from the ST plane.

As a point to note, when the value of z increases, the number of nodal points required for the same camera path decreases and so we get a quantitative difference in the number of disk accesses.

5 Final remarks

In virtual reality and in gaming applications, the light field is useful because no information about the geometry or surface properties is needed. However, there are some disadvantages.

In this paper, we have looked at the problem of reducing the computational burden in dealing with the rich and densely sampled light field when a user walks through a virtual world. We have achieved this by recognizing that in-

stead of considering the complete light field, it is enough to consider a sparse set of nodal points. The number of nodal points, and the distance between them have been characterized to ensure that the rendering of the scene is identical to what may have been done without the cache. The proofs of these characterizations have been shown for a restricted case of arbitrary, but planar motion, for the sake of brevity.

Our description does not explicitly deal with decompression issues (indeed, in the first stage [8] of rendering, the entire light field is decompressed as it is read into memory from disk.) However, there should not be any conceptual blockade in applying the general caching strategy and the mathematical elements even in this case.

Acknowledgments

We are grateful to N. N. Kalyan who contributed to Section 2.4, and other members of ViGIL, IIT Bombay. The images of Buddha and the dragon were taken from <http://www-graphics.stanford.edu/software/lightpack/lifs.html>

References

- [1] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. *M. Landy and J. A. Movshon, (eds) Computational Models of Visual Processing*, 1991.
- [2] D. Aliaga, T. Funkhouser, D. Yanovsky, and I. Carlbom. Sea of images. In *Proceedings of IEEE Visualization*, pages 331–338, 2001.
- [3] C. Buehler, M. Bosse, L. McMillan, S. J. Gortler, and M. F. Cohen. Unstructured lumigraph rendering. In E. Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 425–432. ACM Press / ACM SIGGRAPH, 2001.
- [4] D. Burschka, G. D. Hager, Z. Dodds, M. Jgersand, D. Cobzas, and K. Yerex. Recent methods for image-based modeling and rendering. In *Proceedings of the IEEE Virtual Reality 2003*, page 299. IEEE Computer Society, 2003.
- [5] P. Debevec and S. Gortler. Image-based modeling and rendering, 1998.
- [6] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. *Computer Graphics*, 30(Annual Conference Series):43–54, 1996.
- [7] A. Isaksen, L. McMillan, and S. J. Gortler. Dynamically reparameterized light fields. In K. Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, pages 297–306. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
- [8] M. Levoy and P. Hanrahan. Light field rendering. *Computer Graphics*, 30(Annual Conference Series):31–42, 1996.
- [9] M. S. Presented. Tutorial on light field rendering.
- [10] H. Schirmacher, W. Heidrich, and H.-P. Seidel. Adaptive acquisition of lumigraphs from synthetic scenes. In P. Brunet and R. Scopigno, editors, *Computer Graphics Forum (Eurographics '99)*, volume 18(3), pages 151–160. The Eurographics Association and Blackwell Publishers, 1999.
- [11] P. Sharma, A. Parashar, S. Banerjee, and P. Kalra. An uncalibrated lightfield acquisition system. In *Third Indian Conference on Computer Vision, Graphics and Image Processing ICVGIP*, pages 25–30, 2002.