Full Resolution Lightfield Rendering

Andrew Lumsdaine
Indiana University
lums@cs.indiana.edu

Todor Georgiev
Adobe Systems
tgeorgie@adobe.com

Figure 1: Example of lightfield, normally rendered image, and full-resolution rendered image.

Abstract

Lightfield photography enables many new possibilities for digital imaging because it captures both spatial and angular information, i.e., the full four-dimensional radiance, of a scene. Extremely high resolution is required in order to capture four-dimensional data with a two-dimensional sensor. However, images rendered from the lightfield as projections of the four-dimensional radiance onto two spatial dimensions are at significantly lower resolutions. To meet the resolution and image size expectations of modern digital photography, this paper presents a new technique for rendering high resolution images from the lightfield. We call our approach *full resolution* because it makes full use of both positional and angular information available in captured radiance data. We present a description of our approach and an analysis of the limits and tradeoffs involved. We demonstrate the effectiveness of our method experimentally by rendering images from a 542 megapixel lightfield, using the traditional approach and using our new approach. In our experiments, the traditional rendering methods produce a 0.146 megapixel image, while with the full resolution approach we are able to produce a 106 megapixel final image.

CR Categories: I.3.3 [Computing Methodologies]: Image Processing and Computer Vision—Digitization and Image Capture

Keywords: fully-resolved high-resolution lightfield rendering

1 Introduction

The lightfield is the radiance density function describing the flow of energy along all rays in three-dimensional (3D) space. Since the description of a ray's position and orientation requires four parameters (e.g., two-dimensional positional information and two-dimensional angular information), the radiance is a four-dimensional (4D) function. Sometimes this is called the *plenoptic function*.

Image sensor technology, on the other hand, is only two-dimensional and lightfield imagery must therefore be captured and represented in flat (two dimensional) form. A variety of techniques have been developed to transform and capture the 4D radiance in a manner compatible with 2D sensor technology [Gortler et al. 1996; Levoy and Hanrahan 1996a; Ng et al. 2005a]. We will call this *flat* or *lightfield* representation of the 4D radiance.

To accommodate the extra degrees of dimensionality, extremely high sensor resolution is required to capture flat radiance. Even so, images are rendered from a flat at a much lower resolution than that of the sensor, i.e., at the resolution of the radiance's positional coordinates. The rendered image may thus have a resolution that is orders of magnitude lower than the raw flat lightfield imagery itself.

For example, with the radiance camera described in Section 7 of this paper, the flat is represented in 2D with a $24,862 \times 21,818$ pixel array. The 4D radiance that is represented is $408 \times 358 \times 61 \times 61$. With existing rendering techniques, images are rendered from this radiance at 408×358, i.e., 0.146 megapixel. Not only is this a disappointingly modest resolution (any cell phone today will have

better resolution), any particular rendered view basically only uses one out of every 3,720 pixels from the flat imagery.

The enormous disparity between the resolution of the flat and the rendered images is extraordinarily wasteful for photographers who are ultimately interested in taking photographs rather than capturing flat representations of the radiance. As a baseline, we would like to be able to render images at a resolution equivalent to that of modern cameras, e.g., on the order of 10 megapixels. Ideally, we would like to render images at a resolution approaching that of the high resolution sensor itself, e.g., on the order of 100 megapixels. With such a capability, radiance photography would be practical almost immediately.

In this paper we present a new radiance camera design and technique for rendering high-resolution images from flat lightfield imagery obtained with that camera. Our approach exploits the fact that at every plane of depth the radiance contains a considerable amount of positional information about the scene, encoded in the angular information at that plane. Accordingly, we call our approach *full resolution* because it makes full use of both angular and positional information that is available in the four-dimensional radiance. In contrast to super-resolution techniques, which create high-resolution images from sub-pixel shifted low-resolution images, our approach renders high-resolution images directly from the radiance data. Moreover, our approach is still amenable to standard radiance processing techniques such as Fourier slice refocusing.

The plan of this paper is as follows. After briefly reviewing image and camera models in the context of radiance capture, we develop an algorithm for full resolution rendering of images directly from flats. We analyze the tradeoffs and limitations of our approach. Experimental results show that our method can produce full-resolution images that approach the resolution that would have been captured directly with a high-resolution camera.

Contributions This paper makes the following contributions.

- We present an analysis of plenoptic camera structure that provides new insight on the interactions between the lens systems.
- Based on this analysis, we develop a new approach to lightfield rendering that fully exploits the available information encoded in the four-dimensional radiance to create final images at a dramatically higher resolution than traditional techniques. We demonstrate a $729\times$ increase in resolution of images rendered from flat lightfield imagery.

2 Related Work

In much of the original work on lightfield rendering (cf. [Gortler et al. 1996; Levoy and Hanrahan 1996b]) and in work thereafter (e.g., [Isaksen et al. 2000; Ng et al. 2005b]), the assumption has been that images are rendered at the spatial resolution of the radiance.

Spatial/Angular Tradeoffs A detailed analysis of light transport in different media, including cameras, is presented in [Durand et al. 2005]. Discussions of the spatial and angular representational issues are also discussed in (matrix) optics texts such as [Gerrard and Burch 1994]. A discussion of the issues involved in balancing the tradeoffs between spatial and angular resolution was discussed in [Georgiev et al. 2006]. In that paper, it was proposed that lower angular resolution could be overcome via interpolation (morphing) techniques so that more sensor real-estate could be devoted to positional information. Nonetheless, the rendering technique proposed still assumed rendering at the spatial resolution of the captured lightfield imagery.

Dappled/Heterodyning In the paper [Veeraraghavan et al. 2007], the authors describe a system for "dappled photography" for capturing radiance in the frequency domain. In this approach, the radiance camera does not use microlenses, but rather a modulating mask. The original high-resolution image is recovered by a simple inversion of the modulation due to the mask. However, the authors do not produce a high-resolution image refocused at different depths.

Super Resolution Re-creation of high-resolution images from sets of low resolution images ("super-resolution") has been an active and fruitful area of research in the image processing community [Borman and Stevenson 1998; Elad and Feuer 1997; Farsiu et al. 2004; Hunt 1995; Park et al. 2003] With traditional super-resolution techniques, high-resolution images are created from multiple low-resolution images that are shifted by sub-pixel amounts with respect to each other. In the lightfield case we do not have collections of low-resolution in this way. Our approach therefore renders high-resolution images directly from the lightfield data.

3 Cameras

Traditional photography renders a three-dimensional scene onto a two-dimensional sensor. With modern sensor technologies, high resolutions (10 megapixels or more) are available even in consumer products. The image captured by a traditional camera essentially integrates the radiance function over its angular portion, resulting in a two-dimensional intensity as a function of position. The angular information of the original radiance is lost.

Techniques for capturing angular information in addition to positional information began with fundamental approach of *integral photography* which was proposed in 1908 by Lippmann [Lippmann 1908]. The large body of work covering more than 100 years of history in this area begins with the first patent filed by Ives [Ives 1903] in 1903, and continues to plenoptic [Adelson and Wang 1992] and hand-held plenoptic [Ng et al. 2005b] cameras today.

3.1 Traditional Camera

In a traditional camera, the main lens maps the 3D world of the scene outside of the camera into a 3D world inside of the camera (see Figure 2). This mapping is governed by the well-known lens

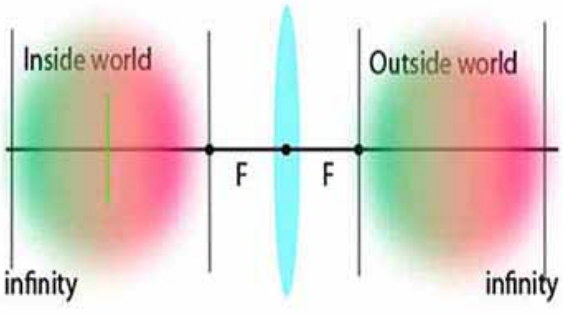

Figure 2: Imaging in a traditional camera. Color is used to represent the order of depths in the outside world, and the corresponding depths inside the camera. One particular film plane is represented as a green line.

equation

$$\frac{1}{A} + \frac{1}{B} = \frac{1}{F}$$

where A and B are respectively the distances from the lens to the object plane and to the image plane. This formula is normally used to describe the effect of a single image mapping between two fixed planes. In reality, however, it describes an infinite number of mappings—it constrains the relationship between, but does not fix, the values of the distances A and B. That is, every plane in the outside scene (which we describe as being at some distance A from the lens) is mapped by the lens to a corresponding plane inside of the camera at distance B. When a sensor (film or a CCD array) is placed at a distance B between F and ∞ inside the camera, it captures an in-focus image of the corresponding plane at A that was mapped from the scene in front of the lens.

3.2 Plenoptic Camera

A radiance camera captures angular as well as positional information about the radiance in a scene. One means of accomplishing this is with the use of an array of microlenses in the camera body, the so-called *plenoptic camera* (see Figure 3).

The traditional optical analysis of such a plenoptic camera considers it as a cascade of a main lens system followed by a microlens system. The basic operation of the cascade system is as follows. Rays focused by the main lens are separated by the microlenses and captured on the sensor. At their point of intersection, rays have the same position but different slopes. This difference in slopes causes the separation of the rays when they pass through a microlens-space system. In more detail, each microlens functions to swap the positional and angular coordinates of the radiance; then this new positional information is captured by the sensor. Because of the swap, it represents the angular information at the microlens. The appropriate formulas can be found for example in [Georgiev and Intwala 2006]. As a result, each microlens image represents the angular information for the radiance at the position of the optical axis of the microlens.

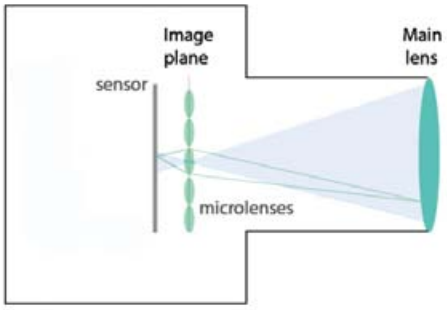

Figure 3: Basic plenoptic camera model. The microlens-space system swaps positional and angular coordinates of the radiance at the microlens. For clarity we have represented only the rays through one of the microlenses.

Images are rendered from the radiance by integrating over the angular coordinates, producing an intensity that is only a function of position. Note, however, the resolution of the intensity function with this approach. Each microlens determines only one pixel in the rendered image. (When you integrate the angular information under one microlens, you only determine one pixel in the rendered image.) If the angular information is finely sampled, then an enormous number of pixels from the flat lightfield imagery are being used to create just one pixel in the rendered image. If the microlens produces, say, a 61×61 array of angular information, we are trading 3,721 pixels in the flat for just one pixel in the rendered image.

Of course, the availability of this angular information allows us to apply a number of interesting algorithms to the radiance imagery. Nonetheless, the expectation of photographers today is to work with multi-megapixel images. It may be the case that some day in the future, plenoptic cameras with multi-millions of microlenses will be available (with the corresponding multi-gigapixel sensors). Until then, we must use other techniques to generate high-resolution imagery.

3.3 Plenoptic Camera 2.0

In the plenoptic camera the microlenses are placed and adjusted accurately to be exactly at one focal length from the sensor. In more detail, quoting from [Ng et al. 2005a] section 3.1:

"The image under a microlens dictates the directional resolution of the system for that location on the film. To maximize the directional resolution, we want the sharpest microlens images possible. This means that we should focus the microlenses on the principal plane of the main lens. Since the microlenses are vanishingly small compared to the main lens, the main lens is effectively fixed at the microlenses' optical infinity. Thus, to focus the microlenses we cement the photosensor plane at the microlenses' focal depth."

This is the current state of the art.

Our new approach, however, offers some significant advantages. In order to maximize resolution, i.e., to achieve sharpest microlens images, *the microlenses should be focused on the image created by the main lens, not on the main lens.* This makes our new camera different from Ng's plenoptic camera. In the plenoptic camera, microlenses are "cemented" at distance f from the sensor and thus focused at infinity. As we will see in Section 7, our microlenses are placed at distance $4/3f$ in the current experiment. The additional spacing has been created by adding microsheet glass between the film and the microlenses in order to displace them by additional $1/3f = 0.2mm$ from the sensor. In this sense, we are proposing "plenoptic camera 2.0" or perhaps could be called "the 0.2 mm spacing camera" (see Figure 4).

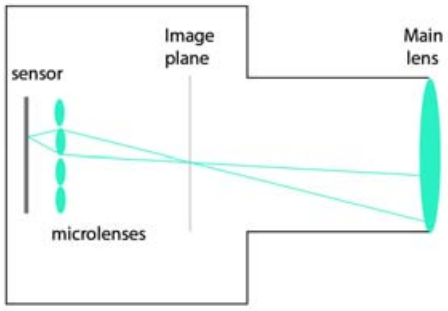

Figure 4: Our proposed radiance camera (plenoptic camera 2.0) with microlens array focused at the image plane.

Analysis in the coming sections will show that focusing on the image rather than on the main lens allows our system to fully exploit positional information available in the captured flat. Based on good

focusing and high resolution of the microlens images, we are able to achieve very high resolution of the rendered image (e.g., a $27\times$ increase in each spatial dimension).

4 Plenoptic Camera Modes of Behavior

The full resolution rendering algorithm is derived by analyzing the optical system of the plenoptic camera. We begin with some observations of captured lightfield imagery and use that to motivate the subsequent analysis.

4.1 General Observations

Figure 5 shows an example crop from a raw image that is acquired with a plenoptic camera. Each microlens in the microlens array creates a microimage; the resulting lightfield imagery is thus an array of microimages. On a large scale the overall image can be perceived whereas the correspondence between the individual microlens images and the large scale scene is less obvious. Interestingly, as we will see, it is this relationship—between what is captured by the microlenses and what is in the overall scene—that we exploit to create high-resolution images.

On a small scale in Figure 5 we can readily notice a number of clearly distinguishable features inside the circles, such as edges. Edges are often repeated from one circle to the next. The same edge (or feature) may be seen in multiple circles, in a slightly different position that shifts from circle to circle. If we manually refocus the main camera lens we can make a given edge move and, in fact, change its multiplicity across a different number of consecutive circles.

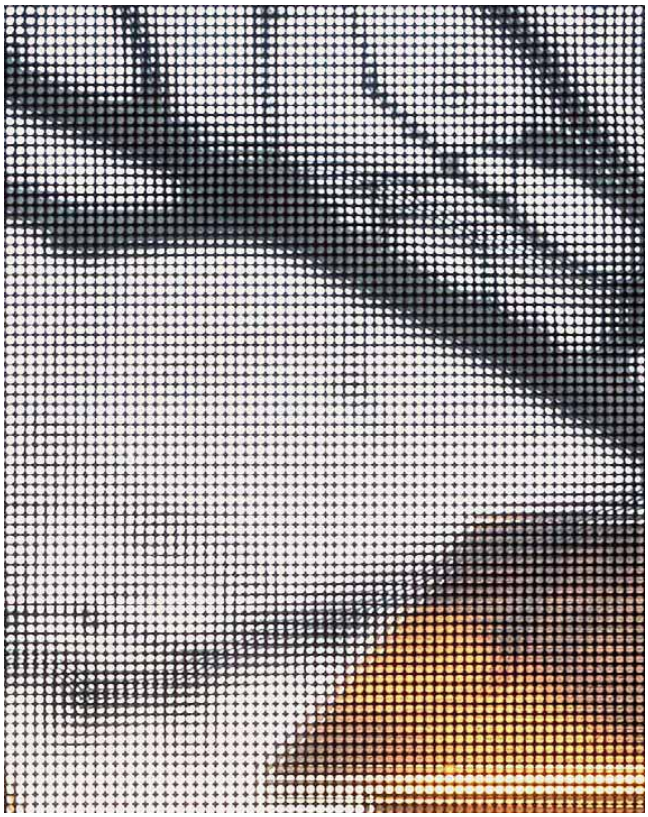

Figure 5: Repeated edges inside multiple circles.

Repetition of features across microlenses is an indication that that part of the scene is out of focus. When an object from the large scale scene is in focus, the same feature appears only once in the array of microimages.

In interpreting the microimages, it is important to note that, as with the basic camera described above, the operation of the basic plenoptic camera is far richer than a simple mapping of the radiance function at some plane in front of the main lens onto the sensor. That is, there are an infinite number of mappings from the scene in front of the lens onto the image sensor. For one particular distance this corresponds to a mapping of the radiance function. What the correspondence is for parts of the scene at other distances—as well as how they manifest themselves at the sensor—is less obvious. This will be the topic of the remaining part of this section.

Next we will consider two limiting cases which can be recognized in the behavior of the the plenoptic camera: *Telescopic* and *Binocular*. Neither of those cases is exact for a true plenoptic camera, but their fingerprints can be seen in every plenoptic image. As we show later in this paper, they are both achievable exactly, and very useful.

4.2 Plenoptic Camera: Telescopic Case

We may consider a plenoptic camera as an array of (Keplerian) telescopes with a common objective lens. (For the moment we will ignore the issue of microlenses not being exactly focused for that purpose.) Each individual telescope in the array has a micro camera (an eyepiece lens and the eye) inside the big camera: Just like any other camera, this micro camera is focused onto one single plane and maps the image from it onto the retina, inverted and reduced in size. A camera can be focused only for planes at distances ranging from f to infinity according to $1/a + 1/b = 1/f$. Here, a, b, and f have the same meaning as for the big camera, except on a smaller scale. We see that since a and b must be positive, we can not possibly focus closer than f. In the true plenoptic camera the image plane is fixed at the microlenses. In [Georgiev and Intwala 2006] we have proposed that it would be more natural to consider the image plane fixed at fistance f in front of the microlenses. In both cases micro images are out of focus.

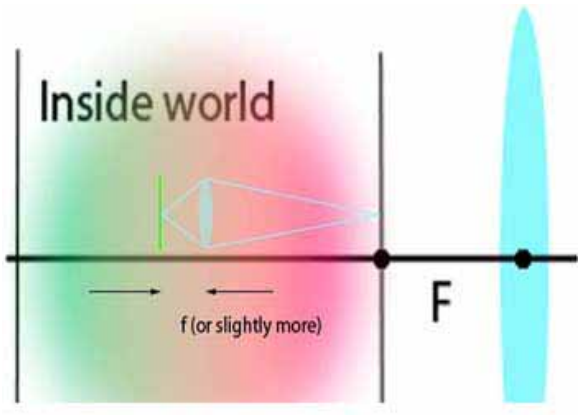

Figure 6: Details of "telescopic" imaging of the focal plane in a pleoptic camera. Note that the image is inverted.

As we follow the movement of an edge from circle to circle, we can readily observe characteristic behavior of telescopic imaging in the flat lightfield. See Figure 7, which is a crop from the roof area in Figure 5. As we move in any given direction, the edge moves relative to the circle centers in the same direction. Once detected in a given area, this behavior is consistent (valid in all directions in that area). Careful observation shows that images in the little

circles are indeed inverted patches from the high resolution image, as if observed through a telescope.

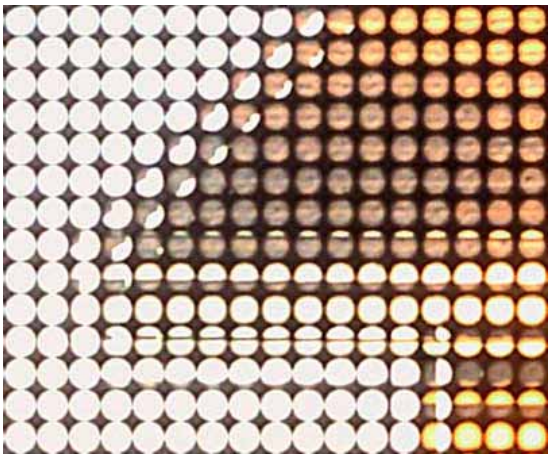

Figure 7: "Telescopic" behavior shown in close up of the roof edge in Figure 5. We observe how the edge is repeated 2 times as we move away from the roof. The further from the roof a circle is, the further the edge appears inside that circle.

4.3 Plenoptic Camera: Binocular Case

We may also consider a plenoptic camera as an "incompletely focused" camera, i.e., a camera focused behind the film plane (as in a Galilean telescope/binoculars). If we place an appropriate positive lens in front of the film, the image would be focused on the film. For a Galilean telescope this is the lens of the eye that focuses the image onto the retina. For a plenoptic camera this role is played by the microlenses with focal length f. They need to be placed at distance smaller than f from the film. Note also that while the telescopic operation inverts the inside image, the binocular operation does not invert it.

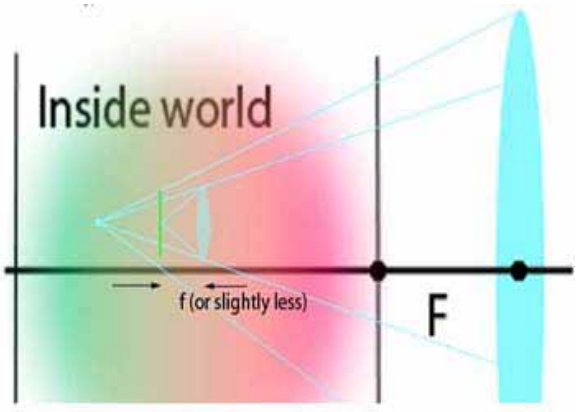

Figure 8: Details of "binocular" imaging in lightfield camera. Note that the image is not inverted.

As with telescopic imaging, we can readily observe characteristic behavior of binocular imaging in the plenoptic camera. See Figure 9, which is a crop from the top left corner in Figure 5. If we move in any given direction, the edge moves relative to the circle centers in the opposite direction. Once detected in a given area, this behavior is consistent (valid in all directions in that area). It is due to the depth in the image at that location. Careful observation shows that images in the little circles are in fact patches from the corresponding area in the high resolution image, only reduced in size. The more times the feature is repeated in the circles, the smaller it appears and thus a bigger area is imaged inside each individual circle.

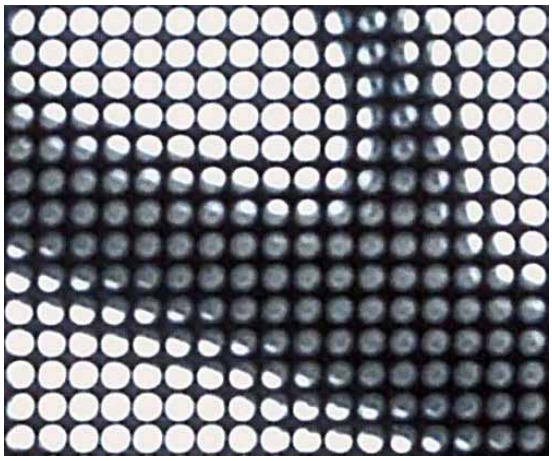

Figure 9: "Binocular" behavior shown in close up of Figure 5. Note how edges are repeated about 2 or 3 times as we move away from the branch. The further from the branch we are, the closer to the branch the edge appears inside the circle.

4.4 Images

To summarize, our approximately focused plenoptic camera can be considered as an array of micro cameras looking at an image plane in front of them or behind them. Each micro camera images only a small part of that plane. The shift between those little images is obvious from the geometry (see Section 5). If at least one micro camera could image all of this plane, it would capture the high resolution image that we want. However, the little images are limited in size by the main lens aperture.

The magnification of these microcamera images, and the shift between them, is defined by the distance to the image plane. It can be at positive or negative distance from the microlenses, corresponding to the telescopic (positive) and binocular (negative) cases. By slightly adjusting the plane of the microlenses (so they are exactly in focus), we can make use of the telescopic or binocular focusing to patch together a full-resolution image from the flat. We describe this process in the following sections.

5 Analysis

Often, microlenses are not focused exactly on the plane we want to image, causing the individual microlens images to be blurry. This limits the amount of resolution that can be achieved. One way to improve such results would be deconvolution. Another way would be to stop down the microlens apertures.

In Figure 10 we consider the case of "plenoptic" camera using pinhole array instead of microlens array. In ray optics, pinhole images produce no defocus blur, and in this way are perfect, in theory. In the real world pinholes are replaced with finite but small apertures and microlenses.

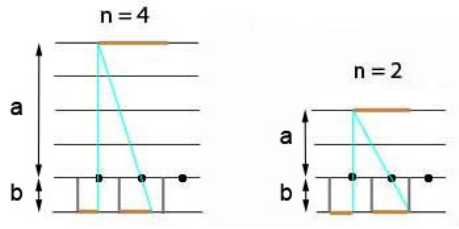

Figure 10: An array of pinholes (or microlenses) maps the areal image in front of them to the sensor. The distance $a = nf$ to the areal image defines the magnification factor $M = n\text{-}1$.

From the lens equation

$$\frac{1}{a} + \frac{1}{b} = \frac{1}{f}$$

we see that if the distance to the object is $a = nf$, the distance to the image would be

$$\begin{aligned} b &= \frac{nf}{n-1} \\ n &= \frac{b}{b-f} \end{aligned}$$

We define the geometric magnification factor as $M = a/b$, which by substition gives us

$$M = n - 1.$$

Figure 10 shows the ray geometry in the telescopic cases for $n = 4$ and $n = 2$. Note that the distance b from the microlenses to the sensor is always greater than f (this is not represented to scale in the figure). Looking at the geometry in Figure 10, the images are M times smaller, inverted, and repeated M times.

6 Algorithm

Section 4 describes two distinct behaviors (telescopic and binocular), and our algorithm executes a different action based on which behavior was observed in the microimages.

Telescopic: If we observe edges (or features) moving relative to the circle centers in the same direction as the direction in which we move, invert all circle images in that area relative to their individual centers.

Binocular: If we observe edges moving relative to the circle centers in a direction opposite to the direction we move, do nothing.

The small circles are, effectively, puzzle pieces of the big image, and we reproduce the big image by bringing those circles sufficiently close together.

The big image could also have been reproduced had we enlarged the pieces so that features from any given piece match those of adjacent pieces. Assembling the resized pieces reproduces exactly the high resolution image.

In either of these approaches the individual pieces overlap. Our algorithm avoids this overlapping by dropping all pixels outside the square of side m.

Prior work did not address the issue of reassembling pixels in this way because the plenoptic camera algorithm [Ng 2005] produces one pixel per microlens for the output image. Our remarkable gain in resolution is equal to the number of pixels m in the original patches.

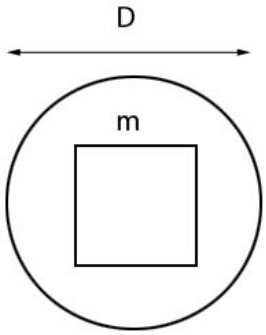

Figure 11: A lens circle of diameter D and a patch of size m.

That is, we produce $m \times m$ pixels instead of one. See Figure 11.

Above we have shown that the magnification $M = n - 1$. Now we see that also $M = D/m$. It therefore follows that

$$n = 1 + \frac{D}{m}.$$

The distance (measured in number of focal lengths) to the image plane in front of the microlens is related to D and m.

It is important to note that lenses produce acceptable images even when they are not exactly in focus. Additionally, out of focus images can be deconvolved, or simply sharpened. That's why the above analysis is actually applicable for a wide range of locations of the image plane. Even if not optimal, such a result is often a useful tradeoff. That's the working mode of the plenoptic camera, which produces high quality results [Ng 2005].

The optics of the microlens as a camera is the main factor determining the quality of each micro image. Blurry images from optical devices can be deconvolved and the sharp image recovered to some extent. In order to do this we need to know the effective kernel of the optical system. While there are clear limitations in this related to bit depth and noise, in many cases we may hope to increase resolution all the way up to m times the resolution of the plenoptic camera. In this paper we demonstrate $27\times$ increase of resolution in one plane, and 10 times increase of resolution in another plane without any deconvolution.

7 Experimental Results

7.1 Experimental Setup

Camera For this experiment we used a large format film camera with a 135mm objective lens. The central part of our camera is a microlens array. See Figure 12. We chose a film camera in order to avoid the resolution constraint of digital sensors. In conjunction with a high resolution scanner large format film cameras are capable of 1 gigapixel resolution.

The microlens array consists of 146 thousand microlenses of diameter 0.25 mm and focal length 0.7 mm. The microlens array is custom made by Leister Technologies, LLC. We crafted a special mechanism inside a 4 X 5 inch film holder. The mechanism holds the microlens array so that the flat side of the glass base is pressed against the film. We conducted experiments both with and without inserting microsheet glass between the array and the film.

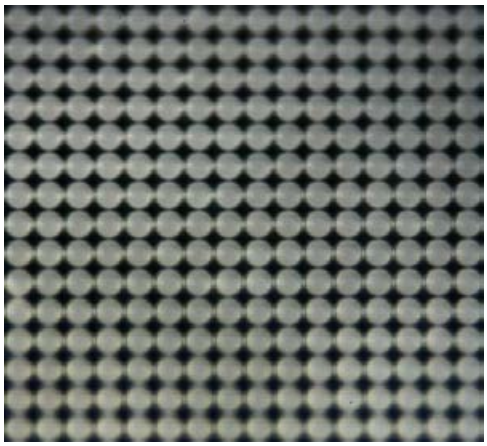

Figure 12: A zoom into our microlens array showing individual lenses and (black) chromium mask between them.

The experiments where the microsheet glass was inserted provided spacing in a rigorously controlled manner.

In both cases our microlenses' focal length is $f = .700$ mm; The spacings in the two experimental conditions differ as follows:

- $b = 0.71$ mm so that $n = 71$ and $M = 70$ which is made possible directly by the thickness of the glass; and
- $b = 0.94$ mm based on microsheet glass between microlens array and film. As a result $n = 3.9$ (almost 4) and $M = 3$, approximately.

Computation The software used for realizing our processing algorithm was written using the Python programming language and executed with Python version 2.5.1. The image I/O, FFT, and interpolation routines were resepectively provided by the Python Imaging Library (version 1.1.6) [pil], Numerical Python (version 1.0.3.1) [Oliphant 2006], and SciPy (version 0.6.0) [Jones et al. 2001–]. All packages were compiled in 64-bit mode using the Intel icc compiler (version 9.1).

The computational results were obtained using a computer system with dual quad-core Intel L5320 Xeon processors running at 1.86 Ghz. The machine contained 16GB of main memory. The operating system used was Red Hat Enterprise Linux with the 2.6.18 kernel.

The time required to render an image with our algorithm is proportional to the number of microlenses times the number of pixels sampled under each microlens. In other words, the time required to render an image with our algorithm is directly proportional to the size of the output image. Even though no particular attempts were made to optimize the performance of our implementation, we were able to render 100 megapixel images in about two minutes, much of which time was actually spent in disk I/O.

7.2 High-Resolution Rendering Results

Figures 13 through 16 show experimental results from applying the full resolution rendering algorithm. In particular, we show the operation of rendering in botrh the telescopic case and the binocular case.

The original image was digitized with the camera, film, and scanning process described above. After digitization, the image measures 24,862 × 21,818 pixels. A small crop from the lightfield image was shown in Figure 5. A larger crop from the flat lightfield is shown in Figure 13.

An image rendered from the lightfield in the traditional way is shown in Figure 14. Also shown in the figure (upper right hand) is a crop of the curb area rendered at full resolution. On the upper left is shown zoom in of the same area cropped directly from the traditionally rendered image. Note that each pixel appears as a 27 × 27 square, and the enormous increase in resolution.

In Figure 15 we show a full resolution rendering of the experimental lightfield, rendered assuming the telescopic case. For this rendering, the scaling-down factor M was taken to be approximately 2.4, so that the full resolution rendered image measured 11016 × 9666, i.e., over 100 megapixels. In this paper we only show a 2,250 × 1,950 region. The image is well-focused at full resolution in the region of the house but not well-focused on the tree branches.

In Figure 16 we show a full resolution rendering of the experimental lightfield, rendered assuming the binocular case. Note that in contrast to the image in Figure 15, this image is well-focused at full resolution in the region of the tree branches but not well-focused on the house.

8 Conclusion

In this paper we have presented an analysis of lightfield camera structure that provides new insight on the interactions between the main lens system and the microlens array system. By focusing the microlenses on the image produced by the main lens, our camera is able to fully capture the positional information of the lightfield. We have also developed an algorithm to render full resolution images from the lightfield. This algorithm produces images at a dramatically higher resolution than traditional lightfield rendering techniques.

With the capability to produce full resolution rendering, we can now render images at a resolution expected in modern photography (e.g., 10 megapixel and beyond) without waiting for significant advances in sensor or camera technologies. Lightfield photography is suddenly much more practical.

References

ADELSON, T., AND WANG, J. 1992. Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 99–106.

BORMAN, S., AND STEVENSON, R. 1998. Super-resolution from image sequences-a review. *Proceedings of the 1998 Midwest Symposium on Circuits and ...* (Jan).

DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. 2005. A frequency analysis of light transport. *ACM Trans. Graph.*, 1115–1126.

ELAD, M., AND FEUER, A. 1997. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured .... *Image Processing*.

FARSIU, S., ROBINSON, D., ELAD, M., AND MILANFAR, P. 2004. Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology*.

GEORGIEV, T., AND INTWALA, C. 2006. Light-field camera design for integral view photography. *Adobe Tech Report*.

GEORGIEV, T., ZHENG, K., CURLESS, B., SALESIN, D., AND ET AL. 2006. Spatio-angular resolution tradeoff in integral photography. *Proc. Eurographics Symposium on Rendering*.

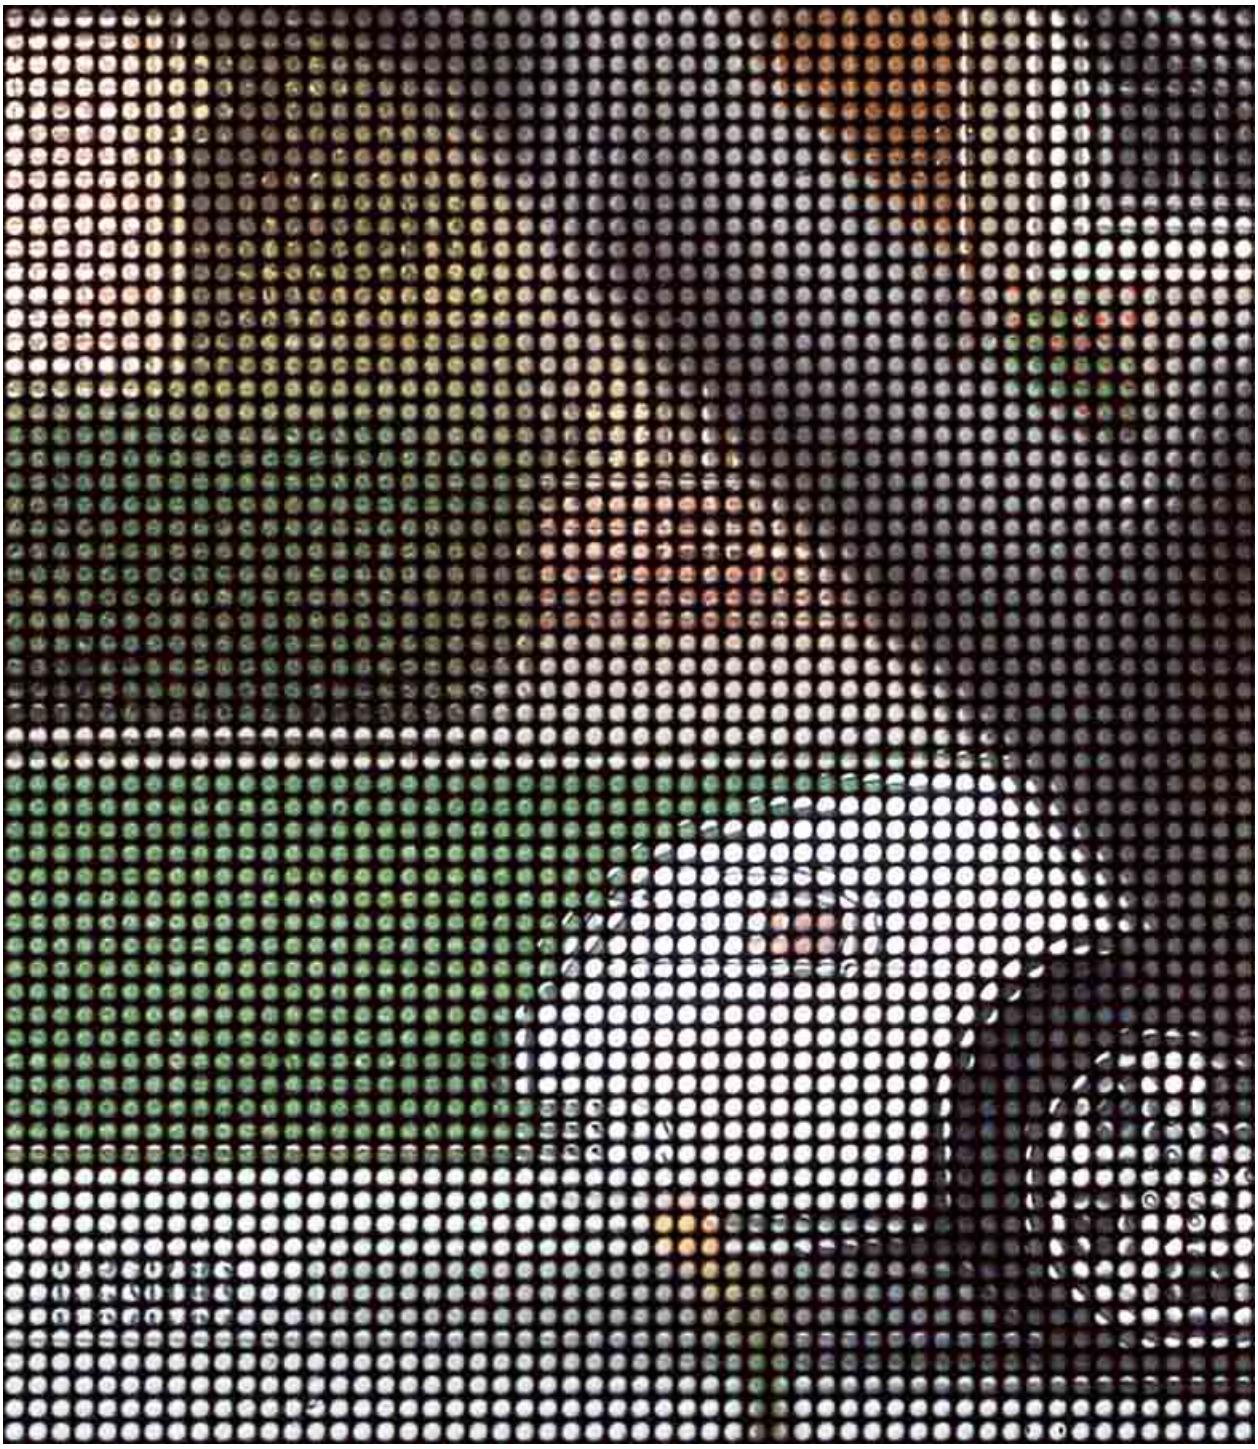

Figure 13: Crop of our lightfield. The full image is 24,862 × 21,818 pixels, of which 3,784 × 3,291 are shown here. This region of the image is marked by the red box in Figure 14.

Figure 14: The entire lightfield rendered with the traditional method, resulting in a 408×358 pixel image. Above are shown two small crops that represent a $27\times$ magnification of the same curb area. The left one is generated with traditional lightfield rendering; the right one is generated with full resolution rendering. A comparison demonstrates the improvement that can be achieved with the proposed method. The red box marks the region shown in Figure 13. The green box marks the region that is shown in Figures 15 and 16.

Figure 15: A crop from a full resolution rendering of the experimental lightfield. Here, the entire image is rendered assuming the telescopic case. We take the scaling down factor M to be approximately 2.4, resulting in a 11016 $\times$ 9666 full resolution image (100 megapixel). A 2,250 $\times$ 1,950 region of the image is shown here. Note that in this case the image is well-focused at full resolution in the region of the house but not well-focused on the tree branches. This region of the image is marked by the green box in Figure 14.

Figure 16: A crop from a full resolution rendering of the experimental lightfield. The entire image is rendered assuming the binocular case. The same 2,250 × 1,950 region as in Figure 15 is shown here. Note that in this case the image is well-focused at full resolution in the region of the tree branches but not well-focused on the house. In other words, only blocks representing the branches match each-other correctly. This region of the image is marked by the green box in Figure 14.

GERRARD, A., AND BURCH, J. M. 1994. Introduction to matrix methods in optics.

GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. *ACM Trans. Graph.*, 43–54.

HUNT, B. 1995. Super-resolution of images: algorithms, principles, performance. *International Journal of Imaging Systems and Technology*.

ISAKSEN, A., MCMILLAN, L., AND GORTLER, S. J. 2000. Dynamically reparameterized light fields. *ACM Trans. Graph.*, 297–306.

IVES, F. 1903. Patent us 725,567.

JONES, E., OLIPHANT, T., PETERSON, P., ET AL., 2001–. SciPy: Open source scientific tools for Python.

LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. *Proceedings of the 23rd annual conference on Computer Graphics and Interactive Techniques*.

LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. *ACM Trans. Graph.*, 31–42.

LIPPMANN, G. 1908. Epreuves reversibles donnant la sensation du relief. *Journal of Physics 7*, 4, 821–825.

NG, R., LEVOY, M., BREDIF, M., DUVAL, G., HOROWITZ, M., ET AL. 2005. Light field photography with a hand-held plenoptic camera. *Computer Science Technical Report CSTR*.

NG, R., LEVOY, M., BRDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light field photography with a hand-held plenoptic camera. *Tech. Rep.*.

NG, R. 2005. Fourier slice photography. *Proceedings of ACM SIGGRAPH 2005*.

OLIPHANT, T. E. 2006. *Guide to NumPy*. Provo, UT, Mar.

PARK, S., PARK, M., AND KANG, M. 2003. Super-resolution image reconstruction: a technical overview. *Signal Processing Magazine*.

Python imaging library handbook. http://www.pythonware.com/library/pil/handbook/index.htm.

VEERARAGHAVAN, A., MOHAN, A., AGRAWAL, A., RASKAR, R., AND TUMBLIN, J. 2007. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph. 26*, 3, 69.