Rich Image Capture with Plenoptic Cameras

Todor Georgiev
Adobe Systems
tgeorgie@adobe.com

Andrew Lumsdaine
Indiana University
lums@cs.indiana.edu

Abstract

The plenoptic function was originally defined as a record of both the 3D structure of the lightfield and of its dependence on parameters such as wavelength, polarization, etc. Still, most work on these ideas has emphasized the 3D aspect of lightfield capture and manipulation, with less attention paid to other parameters. In this paper, we leverage the high resolution and flexible sampling trade-offs of the focused plenoptic camera to perform high-resolution capture of the rich "non 3D" structure of the plenoptic function. Two different techniques are presented and analyzed, using extended dynamic range photography as a particular example. The first technique simultaneously captures multiple exposures with a microlens array that has an interleaved set of different filters. The second technique places multiple filters at the main lens aperture. Experimental results validate our approach, producing 1.3Mpixel HDR images with a single capture.

1. Introduction

The plenoptic function was introduced in [2] as a means of representing all of the information necessary to describe the radiance in a scene. It describes the distribution of light rays in 3D space, as well as other properties of those rays, such as wavelength, polarization, and others. The lightfield, introduced to the computer graphics community in [12], established a theoretical framework for analyzing the plenoptic function, specifically the spatial distribution of the light rays. Cameras for capturing the plenoptic function have similarly emphasized capture of the light ray distribution (the 4D lightfield). In this paper, we revisit the richer definition of the plenoptic function and present techniques that can be used to capture information about the scene radiance in addition to the lightfield.

Plenoptic cameras capture the 4D lightfield by multiplexing it onto a conventional 2D sensor. The 4D lightfield represents positional information of the rays in two dimensions and directional information about the rays in two more dimensions. Capturing a richer plenoptic function involves fully using the multiplexing capabilities of the camera to perform extended sampling so that more information is captured with the same limited sensor capabilities.

How sensor pixels are allocated for capturing additional data about the plenoptic function has direct bearing on the achievable resolution of rendered images. Ng's hand-held plenoptic camera [18], used a microlens array in front of the camera sensor and produced one positional sample per microlens. This scheme resulted in a resolution of the final image that was equal to the number of microlenses, 300×300.

More recently, the focused plenoptic camera [14] introduced a different sampling approach that sampled more sparsely in the directional coordinates, resulting in final rendered images comparable to those of regular, non-plenoptic, cameras. The approach that we present in this paper is based on the focused plenoptic camera. In addition to higher spatial resolution, it also enables flexible trade-offs to be made in how the plenoptic function is sampled. Both of these features are essential for efficiently capturing additional information about the plenoptic function.

There are two complementary techniques that we explore for rich image capture with the focused plenoptic camera. The first approach multiplexes the captured plenoptic function at the microlens array, by interleaving microlenses having different properties so as to distinguish the particular property of the plenoptic function that we wish to multiplex. The second approach performs the same type of modulation at the main camera lens, yet still captures the multiplexed data with a microlens array focused on the image.

In this paper, we illustrate our approach to rich image capture with the example of high dynamic range (HDR) imaging. The contributions of this paper are: (1) A camera design capable of capturing extended dynamic range or other modalities of the plenoptic function at high resolution; (2) A technique for rich image capture using a filter array at the microlenses; (3) A technique for rich image capture using a filter array at the main camera lens; and (4) Analysis of the resolution and practical value of the above techniques.

2. Related Work

Capturing richer information about a scene has been studied in the context of both traditional and plenoptic photography. Capturing more information means capturing

more pixels, with some means of allocating pixels for the richer capture. Pixels are typically allocated in two different ways: spatially, which captures the richer information in a single exposure, or temporally, which captures the richer information across multiple exposures.

A familiar instance of spatial allocation is the Bayer filter, used to capture color digital photographs by interleaving red, green, and blue pixels across the CCD [5]. The "assorted pixels" approach extends this idea to other rich image data (such as high dynamic range, polarization, multiple spectra, etc.) [16]. The "split aperture" approach also allocates pixels spatially, but the allocation is done coarsely, by splitting the incoming image, filtering them separately, and then capturing the filtered scenes separately [4].

The temporal allocation technique of "generalized mosaicing" captures a sequence of images and applies different optical settings to each image in the sequence [20]. Generalized mosaicing has been applied to multi-spectral, HDR, and polarization capture [21–23].

Spatial allocation based on plenoptic sampling was proposed in the context of lightfield microscopy [13]. The idea was more fully explored by Horstmeyer et al. [10], which also built on previous work on pupil plane multiplexing [9, 11]. Although the multimodal camera presented in [10] uses pinholes rather than microlenses, the main lens is focused on the pinholes. Since each modal image contains only one pixel from each pinhole, this work would be classified as a traditional plenoptic camera.

The work presented in this paper follows the spatial allocation approach, i.e., we capture a single exposure that contains sets of pixels that have been exposed to different optical settings. Our approach is unique in that it is developed in the context of the focused plenoptic camera [14], which enables higher resolution and flexible trade-offs of captured image parameters.

3. HDR with the Focused Plenoptic Camera

The most popular HDR method allocates pixels temporally. That is, the same scene is photographed multiple times, at different exposures and the resulting images are merged into a single floating-point image (which can support extended dynamic range). In the simplest case, a final image is composed of piecewise linear transfer curves with different slopes. For purposes of display, tone mapping or other HDR compression techniques are often used to produce a low dynamic range output image while preserving contrast and image details [6, 7, 24].

In the case of a moving scene, pixels must be allocated spatially because multiple exposures cannot be taken at different times. A plenoptic camera captures multiple views of the same scene. By appropriately filtering those views, a plenoptic camera can be made to capture multiple exposures of a scene at the same time. With the traditional plenoptic camera, all of the different views of a particular point in the scene are captured by a single microlens. As a result, different views can be separately filtered by filtering at the main lens (as in [10]). With the Lambertian assumption, radiance is constant for all views, so this approach is equivalent to filtering the same scene with different filters. With the focused plenoptic camera, different views are captured by different microlenses. Different views can therefore be separately filtered in a focused plenoptic camera in two ways: by filtering at the microlenses or, as with the traditional plenoptic camera, by filtering at the main lens.

3.1. Filtering at the microlenses

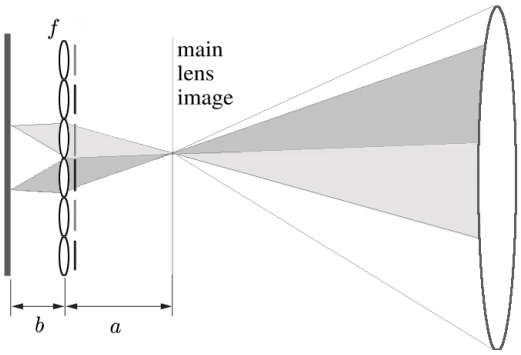

Figure 1: Focused plenoptic camera as a relay system. Interleaved filters on the microlens apertures are shown.

The main lens of the focused plenoptic camera forms an image at its image plane. As shown in [14], an array of microlenses placed behind that plane maps the image from the image plane to the sensor (see Figure 1).

Each microlens creates its own little image of part of the scene, as seen through the main lens aperture—which defines its image shape and size. In other words, each microlens works as a microcamera and its exposure can be determined by placing a Neutral Density (ND) filter in front of it, as shown in Figure 1.

In the focused plenoptic camera the main lens image is formed at distance a in front of the microlenses. We are careful to adjust and fix appropriate distance b to the sensor, so that the lens equation is satisfied

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

with the main lens image placed at distance a. Fine-tuning these distances can make the captured image N times smaller than the main lens image, where

$$N = \frac{a}{b}. \tag{2}$$

Being N times lower, resolution thus depends on the distance from the microlenses to the images created inside the

camera (and, correspondingly, to scene depth). Rendering an entire scene with a fixed resolution will thus produce artifacts if the scene has sufficiently different depths. We require that $N > 2$ so that every point in the scene is seen at least 2 times, i.e., at least once for each filter. At the same time, since we lose a factor of N in the spatial resolution of our final rendered image, we would like to keep N small.

Consider the phase space diagram of image capture in Figure 2. Following the position-direction convention of [8], we denote the 4D radiance at a given plane perpendicular to the optical axis as $r(x, \nu)$, where x describes the location of a ray in the plane and ν describes its direction. In 2D space, $\nu = \tan\theta$, where θ is the angle of the ray relative to the optical axis. As shown in [14], the image formed behind each microlens samples the plenoptic function at the main lens image plane. Sampling is done in a tilted fashion defined by the optical transfer matrices of the system. We adjust the parameters such that images in the microlenses overlap, i.e., the same spatial coordinate (x) corresponds to two images or more. By placing ND filters on the microlenses we interleave different types of sampling as shown with the bright and dark slanted lines in Figure 2 (middle). In this case we show simple modulation with two filters, but the extension to more filters should be clear.

Figure 3 shows a crop from the array of microimages created by the microlenses. Since different microlens filters are interleaved, we observe respectively bright and dark microimages (each with the shape of the main lens aperture). This is the interleaved pattern in the phase space diagram, Figure 2 (middle). Also, notice that the main lens aperture is made to be square so that microimages are square and tile together with little loss of space between them.

The microimages captured in this way with the thousands of microcameras are reduced in size relative to the focal plane image. As discussed above, the reduction is by a factor of N, typically chosen to be between 2 and 5. In our design we interleave transmission of 100% and $1/16$ in a checkerboard pattern.

3.2. Filtering at the main lens

Different views can be separately filtered by placing filters at the main lens. This solution is easier to implement because physical dimensions are on the order of millimeters and can be manipulated by hand. Considered at one group of microlenses, the light modulation pattern would be similar to that in the microlens filter design (see Figure 4).

The phase space diagram (Figure 2) shows the difference between the two approaches to filtering. When filtering at the microlenses, the pattern of k elements is spread over k microlenses, whereas when filtering at the main lens, the whole pattern is visible in each microimage. (Figure 2 shows simple modulation with $k = 2$ filters.)

In Figure 5 we see a crop from the array of microimages created by the microlenses when using 4 filters on the main lens aperture. We observe the 4 bright and dark areas within each microimage. Compare this with the phase space diagram of Figure 2 (right). The change of pattern appears twice as fast along each axis, compared to a similar image captured with interleaved filters on the microlenses in Figure 3. Again, the main lens aperture is made to be square so that microimages are square and tile together with little loss of space between them.

3.3. Main Lens Filtering: Resolution Analysis

The approach of putting filters at the main lens can also be used with the traditional plenoptic camera [10]. Since the microlenses are focused at the main lens aperture, this approach is good for discriminating between different filters—for each pixel in the image we know exactly by which filter it has been influenced.

In the focused plenoptic camera, the microlenses are not focused on the main camera lens, but rather on the main camera lens image. As a result, structures on the main lens aperture will be blurred in each microimage.

Thus, we have the following trade-offs (which we will quantify below):

1. Focus the microlenses on the main image. Filter boundaries will be blurry, but the image of the scene will be sharp. The ability to distinguish the filtering for particular pixels may be limited.
2. Focus the microlenses on the main lens. Filter boundaries will be sharp, but the image of the scene blurry. Overall image resolution will be limited.

To quantify these trade-offs, we estimate the main lens aperture blur as it appears in a microlens image. Consider the camera system in Figure 6. As in the traditional plenoptic camera [3], microlenses are small and the main lens can be assumed to be at optical infinity. The main lens aperture P is imaged from optical infinity to P', a distance f behind each microlens, while the image of the scene is formed at distance b behind the microlens, where the sensor is located. We are interested in the blur caused by this misfocus of the main lens aperture.

Using the lens equation in Newton's form

$$(a - f)(b - f) = f^2, \tag{3}$$

the distance between the image of the scene and the image of the main aperture is

$$b - f = \frac{f^2}{a - f}. \tag{4}$$

If the microlenses have microapertures of diameter d, the diameter of the circle of confusion for the image of the main aperture would be

$$\Delta = \frac{f^2}{a - f}\frac{d}{f}. \tag{5}$$

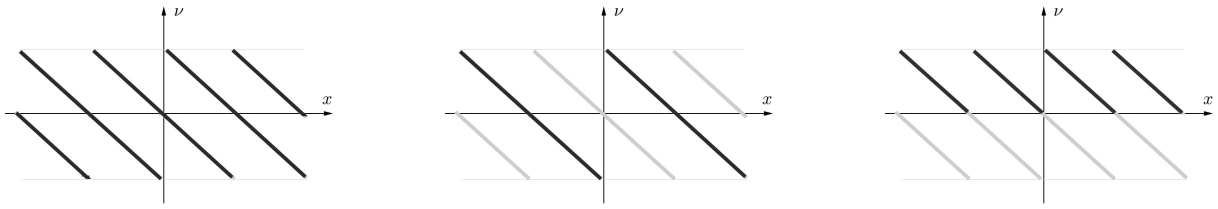

Figure 2: Phase-space diagram of lightfield capture in the focused plenoptic camera. The basic focused plenoptic camera is shown on the left, the camera with interleaved microlens images in the center, and the camera with filters at the main lens on the right. Similarly filtered samples are shown in the same color (black or grey).

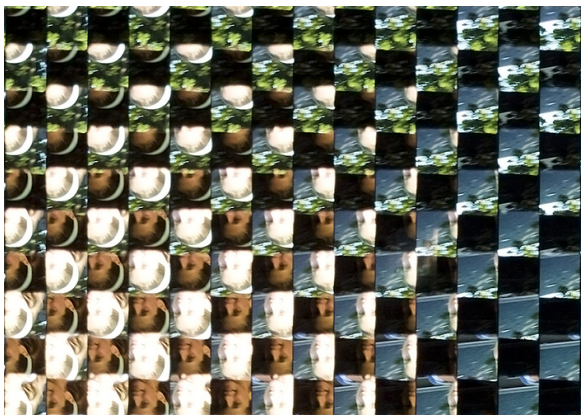

Figure 3: Multiple capture of the image by microlenses with interleaved filters. Notice the alternating brightness.

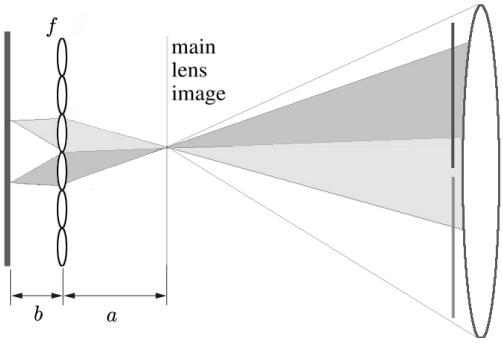

Figure 4: Focused plenoptic camera as a relay system with filters on the main camera lens aperture.

From the lens equation (1) we have

$$\Delta = \frac{bd}{a}. \tag{6}$$

Another useful expression for (6) is in terms of the magnification factor $N = a/b$, i.e.,

$$\Delta = \frac{d}{N}. \tag{7}$$

The upper limit of this blur is $\Delta < \frac{d}{2}$, which corresponds to the minimal applicable $N = 2$. A typical blur that we get

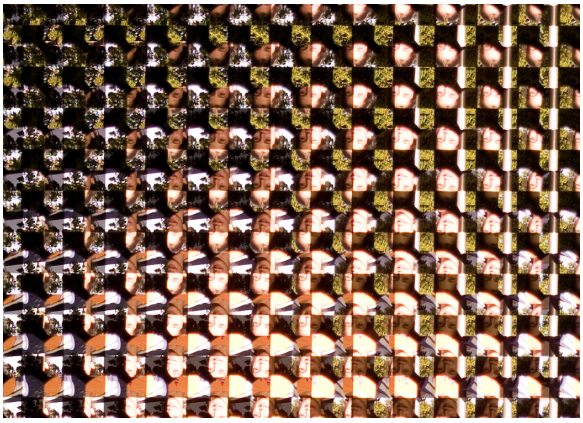

Figure 5: Capture of the image created by the main lens with 4 filters. Notice the alternating brightness effect–each microimage is sharply split into 4 parts.

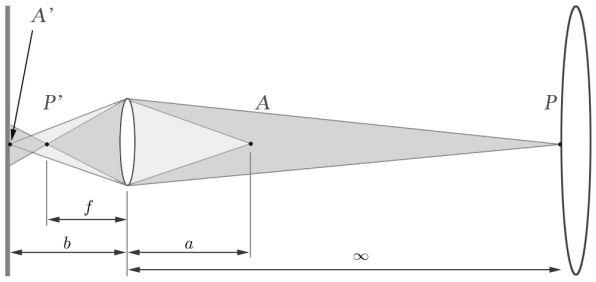

Figure 6: A point P on the main lens aperture (at optical infinity) is imaged to point P' behind the microlens and creates a circle of confusion on the sensor. At the same time, a point A from the main lens image is imaged to A' exactly on the sensor.

with our imaging is at around $N = 5$, which would be about 3 pixels with our equipment (see Section 4).

Note that this blur is only at the boundaries between filtered regions in each microimage and so only a small percentage of pixels become unavailable for use in reconstructing the final image. Also, in terms of such blur we are typically 5 times better than the same camera using an array of pinholes, which has blur of no less than the pinhole diame-

ter $\Delta = d$.

A similar calculation shows that if we focus the microlenses on the main lens aperture, the blur of the image of the scene will be of approximately the same amount:

$$\Delta = \frac{d}{N+1}. \quad (8)$$

However, now the blur is across the entire microimage.

This analysis shows that focusing on the image has significant advantages over focusing on the main lens. When focusing on the image we lose a few pixels at the filter edges, but the image itself is sharp. The loss in pixel count in our case is about 10%. Even this can be compensated by computing the influence of defocus blur of the filters on individual pixel values at the filter image boundaries, and restoring their "unmixed" values. On the other hand, if we focus the microlenses on the main lens aperture, the filter edge boundaries are sharp, but the image itself is blurred. For example, a 3 pixel blur would be equivalent to $3 \times 3 = 9$ times lower resolution in terms of pixel count. The difference in resolution is fundamental to these two methods, with significantly better resolution provided by focusing on the image.

Thus, using the focused plenoptic camera for rich image capture makes efficient use of sensor pixels. For example, without any optimization for pixel usage, our current camera produces 1.3 Mpixel images (1,300 $\times$ 1,000 pixels). Our images are slightly blurry (we have reduced image size for this paper), but this is due to the particular optics used, and not a fundamental constraint of our imaging model. To compare to related efforts, the images reported in [18] and [10] were 300 $\times$ 300 and 177 $\times$ 177, respectively. With more careful camera design we should be able to approach a factor of k reduction in the total sensor resolution when capturing k modes. In our case of 39Mpixel sensor, four separate modes could be captured at 9Mpixels each.

Note that the focused plenoptic rendering could be applied to the traditional plenoptic camera if the main lens image is appropriately spaced from the microlenses. However because the microlenses are defocused in the plenoptic camera design, this image would necessarily be blurred. Additional resolution could be achieved by deblurring (e.g., with deconvolution techniques). However, there remains a fundamental difference between the traditional and the focused plenoptic camera. In the focused plenoptic camera the sensor plane is conjugate to the object plane, while in the traditional plenoptic camera it is not.

In the case of specular highlights (or other non-Lambertian scenes) we remark that there is a difference between main lens filtering and microlens filtering. From Figure 2, middle, we see that for large N the configuration with filters at the microlenses will sample densely in the angular dimension, meaning it will work well for non-Lambertian/specular scenes as well as for other plenoptic rendering tasks such as refocusing. The configuration with filters at the main lens always has certain views tied to certain filters and so cannot take advantage of large numbers of views in the same way (Figure 2, right).

4. Experimental Setup

Our camera is medium format with an 80-mm lens and a 39-megapixel digital back from Phase One. The lens is mounted on the camera with a 13-mm extension tube, which provides the needed spacing a. The microlens array is custom made by Leister Axetris. The microlenses have focal length of 1.5 mm so that they can be placed directly on the cover glass of the sensor. We are able to provide additional spacing of up to 0.5 mm to enable fine tuning of the microlens focus. The pitch of the microlenses is 500 μm with precision 1 μm. The sensor pixels are 6.8 μm. The value of $b \approx 1.6$ mm was estimated with precision 0.1 mm from known sensor parameters and independently from the microlens images at different F/numbers.

Our microlens apertures are circular, with diameters 100 μm, and are formed by the black chromium mask deposited on the microlenses. While the small apertures extend depth of field and make microlenses diffraction limited, they also introduce high F/number and associated with it diffraction blur and longer exposure times. Our camera would only benefit from apertures of larger size. The small apertures that we have used are not required, especially with good optical quality microlenses.

The relatively large pitch of the microlenses is chosen in order to match the F-number of the main camera lens, which can be as low as $F/3$. This large pitch is needed because our microlenses are at large distance $(1.6mm)$ from the sensor, defined by the cover glass. We opened the main camera lens and introduced filters and a square aperture at the original aperture location. We have experimented with ND filters, polarizers, and a large set of color filters. Due to space limitation, the printed version of this paper shows only ND filters result, but in the electronic version we show some of the other results.

We also use an array of square filters of side $0.5mm$, deposited as chromium masks on a thin $0.21mm$ silica wafer. This forms our array of interleaved ND filters. The ND microfilters form a checkerboard pattern: 100% and $1/16$ transmission, to achieve dynamic range expansion of 4 stops. No anti-reflectant coating was applied in order to test the artifacts produced without it.

5. Experimental Results

To make rich image data captured by the focused plenoptic camera available for subsequent processing, we render separate images for each independently captured mode,

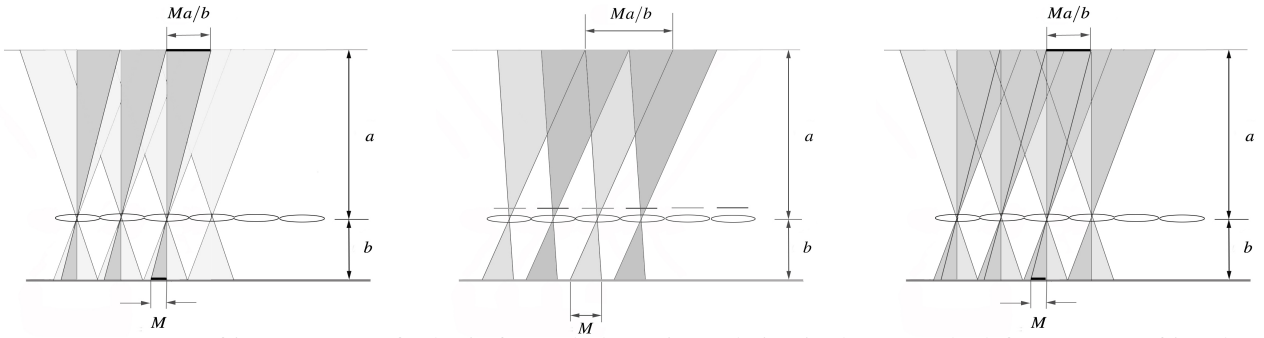

Figure 7: Geometry of image capture for basic focused plenoptic rendering is shown on the left, geometry of interleaved microlens filters is shown in the center, and geometry of imaging with filters at the main camera lens (at infinity) is shown on the right.

where the separate images are each rendered using the focused plenoptic rendering algorithm as described in [14]. In this rendering process, the captured radiance is represented as a 4D array. The algorithm selects contiguous samples of size M from each microimage, spaced at equal distances, and patches them together to generate a view. This process is shown in phase space in Figure 2 and geometrically in Figure 7. We remark that the basic rendering approach of [14] often results in artifacts between extracted samples when the samples do not exactly match at their boundaries (due, e.g., to parallax, incorrect scaling, or optical variations between microcameras). Approaches to eliminating these artifacts are the subject of current work and will be reported in a forthcoming paper.

5.1. Interleaved filters on the microlenses

To render separate images when we use interleaved filters on the microlenses, we modify the focused plenoptic rendering algorithm to accommodate the allocation of pixels to differently filtered images. In the case of interleaved filters, we have interleaved plenoptic images. We render the final image by applying the basic focused plenoptic algorithm using only images of one mode with modified "pitch" size (distance between microlens images). It is multiplied by a constant factor, in our case – doubled. Consider the phase space diagram in Figure 2 and the geometric diagram in Figure 7. By sampling at twice the normal pitch, the focused plenoptic rendering algorithm will create a final image from all of the black samples or all of the grey samples (depending on choice of initial sample).

In Figure 8 (left and middle) we patch together a final image using respectively the unfiltered aperture microlenses and the filtered aperture microlenses (Figure 3). The neutral density filter in this case reduces light by a factor of 16. The figure also shows the captured images merged into a final HDR image, using Photoshop. (For the merged image, we simply use a standard solution because our main point is image capture, not tone mapping and HDR processing.)

5.2. Images with filters on the main lens

As with the previous case, we also have interleaved plenoptic images, but the interleaving is within microlens images. We therefore can render the final image simply by applying the basic focused plenoptic rendering to specific portions of each microlens image. Consider the phase space diagram in Figure 2 and the geometric diagram in Figure 7. By choosing samples from the left side of each microimage (positive v), we render a final image from the black samples. With samples from the right side of each microimage (negative v), we render a final image from the grey samples.

In Figure 9 (left and middle) we have patched together a final image using respectively the unfiltered part of the main lens and the filtered part of the main lens (see Figure 5). The neutral density filter in this case reduces light by a factor of 16. The figure also shows the captured images merged into a final HDR image, using Photoshop. (For the merged image, we simply use a standard solution because our main point is image capture, not tone mapping and HDR processing.)

6. Analysis

Capturing dynamic rather than static scenes is a unique power of our approach. It can only be partially matched with methods performing multiple sampling of each pixel, like [19], or methods with variable pixel size like [17], and related. Our method is optical, and can be used on top of any other method.

Comparison with variable size pixels [17] shows great similarities. Both methods are related to loss in resolution proportional to the number of different exposures. While the variable pixels has resemblance to the Bayer array approach and requires a blurring (antialiasing) filter, our method has perfect sharpness but it produces reduced size image by a factor equivalent to the blur. However, our approach has slightly better quality, similar to the way a Foveon sensor is better than a Bayer array sensor based on sampling each color at exactly the same location [1]. Under the Lambertian assumption, our "assorted pixels" are effec-

Figure 8: Two bracketed images from a single exposure using microlens filters, and final tone-mapped HDR image (right).

Figure 9: Two bracketed images from a single exposure using main lens filter, and final tone-mapped HDR image (right).

tively on top of each other, i.e., they differ only in angular and not in spatial coordinates. In the case where there is visible parallax, our approach may not necessarily exhibit these advantages.

A unique feature of our approach is that it avoids blooming. If a pixel is very brightly illuminated so that neighboring pixel values in the sensor are damaged by blooming, we can always pick the representation of those pixels from another microlens image captured with a darker ND filter (albeit with loss of dynamic range). For comparison, in approaches like [17], if a pixel with a darker filter is a direct neighbor to a bright one, such that there is blooming, the damage is irreparable.

Considering our microlens filter design, Figure 10, the use of two separate wafers increases inter-reflectance and (micro) lens glare. The filter (mask) wafer reflects light back towards the microlens wafer, which in turn reflects some of it back. Several inter-reflections cause lens glare as can be seen in our final images in Figure 8.

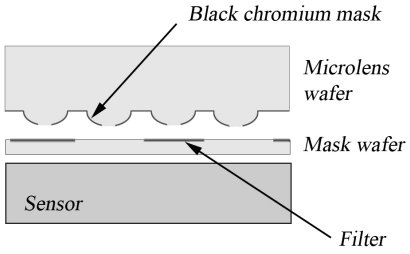

Figure 10: Geometry of our microlens array with apertures and individual ND micro filters on a separate mask wafer.

This effect was intentional—we did not request anti-reflectant coating for the mask and microlens wafers so that we could evaluate the importance of the anti-reflectant coating on our microlenses and microlens filters. We conclude that antireflectant coating is very important.

A comparison with the main lens aperture modulation approach shows the difference. Consider Figures 8 and 9. The effect of microlens glare is especially strong in dark microimages that are close to bright ones. Since we used photographic quality filters on the main lens, the images in Figure 9 do not have as strong a problem with glare and are generally of better sharpness and quality.

Another issue with our approach is that since the microimages sample at a tilt in phase space, there is some small parallax between the images produced from different modes. For applications where the modes are combined to a final image (such as HDR), the parallax may need to be accounted for. However, in many multi-modal applications, the need to properly align image features for the combining process is an established part of the process. Parallax depends on main aperture diameter—it is small with point and shoot cameras and with any camera at large distances.

Note the small white squares, mainly concentrated in the bottom left area of the image. These are due to defective microlenses and are not significant for evaluating the method since better manufacturing can eliminate them completely.

7. Extended Color and Polarization Capture

The use of filters to capture additional dimensions or types of data in plenoptic cameras is a general idea, not restricted only to HDR image capture. Placing filters at the main lens aperture has the advantage that the type of filtering being done can be easily (and economically) changed. See the supplemental electronic material provided with this paper for images that demonstrate extended color capture and polarization capture, respectively using color filters and polarization filters at the main lens aperture.

8. Conclusion and Future Work

In this paper we have demonstrated rich image capture with the focused plenoptic camera, using the specific examples of HDR capture. Our design is based on the focused plenoptic camera and used interleaved microlens aperture filters or filters at the main lens aperture, allowing high spatial resolution as well as increased dynamic range capture. Future work includes extending this same approach to other modes, such as color and polarization, and to developing artifact-free rendering algorithms.

References

[1] http://www.foveon.com/article.php?a=74.

[2] E. Adelson and J. Bergen. The plenoptic function and the elements of early vision. In *Computational models of visual processing*. MIT Press, 1991.

[3] E. Adelson and J. Wang. Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 99–106, 1992.

[4] M. Aggarwal and N. Ahuja. Split aperture imaging for high dynamic range. *International Journal of Computer Vision*, Jan 2004.

[5] B. E. Bayer. Color imaging array. US Patent 3971065, 1976.

[6] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. *ACM Transactions on Graphics, SIGGRAPH 1997 Conference Proceedings, San Diego, CA*, pages 369–378, 1997.

[7] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics, SIGGRAPH 2002 Conference Proceedings, San Antonio, TX*, pages 257–266, 2002.

[8] F. Durand, N. Holzschuch, C. Soler, E. Chan, et al. A frequency analysis of light transport. *International Conference on Computer Graphics and Interactive Techniques*, Jan 2005.

[9] R. Horstmeyer, R. Athale, and G. Euliss. Modified light field architecture for reconfigurable multimode imaging. In *Adaptive Coded Aperture Imaging, Non-Imaging, and Unconventional Imaging Sensor Systems*. SPIE, 2009.

[10] R. Horstmeyer, G. Euliss, R. Athale, and M. Levoy. Flexible multimodal camera using a light field architecture. In *Proceedings ICCP 2009*, 2009.

[11] R. Horstmeyer, G. W. Euliss, R. A. Athale, R. L. Morrison, R. A. Stack, and J. Ford. Pupil plane multiplexing for multi-domain imaging sensors. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7096, Aug. 2008.

[12] M. Levoy and P. Hanrahan. Light field rendering. *ACM Transactions on Graphics*, pages 31–42, 1996.

[13] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz. Light field microscopy. *SIGGRAPH '06: SIGGRAPH 2006 Papers*, Jul 2006.

[14] A. Lumsdaine and T. Georgiev. The focused plenoptic camera. In *IEEE International Conference on Computational Photography (ICCP)*, 2009.

[15] M. McGuire, W. Matusik, H. Pfister, B. Chen, J. F. Hughes, and S. K. Nayar. Optical splitting trees for high-precision monocular imaging. *IEEE Comput. Graph. Appl.*, 27(2):32–42, 2007.

[16] S. G. Narasimhan and S. K. Nayar. Enhancing resolution along multiple imaging dimensions using assorted pixels. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):518–530, 2005.

[17] S. K. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. *Proceedings IEEE CVPR*, 2000.

[18] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a hand-held plenoptic camera. Technical Report CSTR 2005-02, Stanford Univ, Apr. 2005.

[19] J. H. Park et al. An ultra wide dynamic range cmos image sensor with a linear response. *Proceedings SPIE-IS&T Electronic Imaging*, SPIE Vol. 6068, 2006.

[20] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing. In *ICCV*, pages 17–25, 2001.

[21] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing: Wide field of view multispectral imaging. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(10):1334–1348, 2002.

[22] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing: High dynamic range in a wide field of view. *International Journal of Computer Vision*, 53(3):245–267, 2003.

[23] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing: Polarization panorama. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):631–636, 2005.

[24] J. Tumblin and G. Turk. A boundary hierarchy for detail-preserving contrast reduction. *ACM Transactions on Graphics, SIGGRAPH 1999 Conference Proceedings, Los Angeles, CA*, pages 83–90, 1999.