

Digital Correction of Lens Aberrations In Light Field Photography

Ren Ng and Pat Hanrahan

Computer Science, Stanford University, Gates 368, Stanford, California 94305, USA

ABSTRACT

Digital light field photography consists of recording the radiance along all rays (the 4D light field) flowing into the image plane inside the camera, and using the computer to control the final convergence of rays in final images. The light field is sampled with integral photography techniques, using a microlens array in front of the sensor inside a conventional digital camera. Previous work has shown that this approach enables refocusing of photographs after the fact. This paper explores computation of photographs with reduced lens aberrations by digitally re-sorting aberrated rays to where they should have terminated. The paper presents a test with a prototype light field camera, and simulated results across a set of 35mm format lenses.

1. INTRODUCTION

Digital image sensor resolution is growing exponentially, and today it is not uncommon to see commodity cameras with ten megapixels (MP) of resolution [1]. Current trends suggest that cameras with over 100 MP are possible. For example, if one were to use the VLSI designs for the smallest pixels present in low-end cameras (1.9 micron pitch) on the large sensor die sizes in high-end commodity cameras (24 mm \times 36 mm), one could print a sensor with close to 250 MP. However, there is a growing consensus that raw sensor resolution is starting to exceed the resolving power of lenses and the output resolution of displays and printers. For example, for the most common photographic application of printing 4" \times 6" prints, more than 2 MP provides little perceptible improvement [2].

Recent work [3–5] on *digital light field photography* has focused on devoting some of this excess sensor resolution to measuring ray directional information. The idea is to capture not just a 2D photograph of the irradiance on the imaging plane, but to sample the full radiance distribution inside the camera body. The four-dimensional radiance field is often called the *light field*. The purpose of capturing the additional two dimensions of data is to enable computation of the final photograph flexibly and at higher-quality than in the conventional case. This is done by resampling the rays of light, sorting them to where they would have terminated if the camera had been configured as desired. For example, *digital refocusing* simulates a virtual camera focused at a different depth in the world to compute a photograph focused at that depth.

Section 2 reviews this previous work on digital light field photography, including the optical principles, processing concepts, and performance of digital refocusing. The remainder of the paper discusses applications to reducing the effects of lens aberrations.

In the 1850s, Maxwell proved that no optical system can achieve ideal imaging at all focal depths, because such a system would necessarily violate the basic mechanisms of reflection and refraction [6]. Lens designers must make compromises in which parts of the focal and image space will be best corrected for aberrations. Correction has traditionally been viewed as an optical design problem. The usual approach has been to combine lens elements of different shapes and glass types, balancing the aberrations of each element to improve the image quality of the combined system. Today the process of correcting aberrations by combining glass elements has been carried to remarkable extremes. For example, commodity 35 mm zoom lenses generally contain no fewer than ten different glass elements, and some have as many as twenty-three [7].

This paper introduces a new, pure-software approach to compensating for lens aberrations after the photograph is taken. This approach complements the classical optical techniques. The central concept is simple: since a light field camera records the light traveling along all rays flowing into the photograph, we can use the computer to re-sort aberrated rays of light to where they should ideally have converged. Digital correction of this kind improves the quality of final images by reducing residual aberrations present in any given optical recipe. To simplify the analysis, this paper assumes monochromatic light, neglecting chromatic aberrations.

2. REVIEW OF DIGITAL LIGHT FIELD PHOTOGRAPHY

2.1. Recording a Photograph's Light Field

To sample the light field contributing to an ordinary photograph, we modify a conventional camera (photographic main lens and photosensor) with a microlens array inserted in front of the sensor. The light field camera model is illustrated in Figure 1A. We have built a prototype camera that uses a medium format camera body with a 4096×4096 sensor. A 296×296 microlens array is mounted above the sensor. The width of each microlens is 125 microns. The main lens and microlens each have an aperture of $f/4$. More details about the prototype camera are available in Ng [5]. This kind of a camera is sometimes called a plenoptic camera [8], and it is based on integral photography techniques [9, 10].

Figure 1-B1 illustrates the Cartesian *ray-space diagram* corresponding to this camera configuration. Each point (x, u) on the diagram represents the ray passing through x on the imaging plane inside the camera, and u on the exit pupil, as shown on Figure 1A. In general we consider rays (x, y, u, v) inside the camera that pass through (x, y) on the image plane and intercepting the exit pupil at (u, v) . For conceptual simplicity, we draw u as an intersection point on the aperture of the lens on the ray-trace diagrams, ignoring pupil effects.

The ray-space diagram depicts the light field sampling pattern that is captured by the sensor, as an overlay of grid of boxes on the diagram. Each grid box represents the set of light rays integrated by one photosensor pixel. We computed the grid by ray-tracing through the example double-gauss main lens. Note that the grid has a vernier offset within each vertical strip, because there is a non-integral number of pixels across each image that appears under a microlens [11]. Each column of the ray-space diagram represents the set of rays that strike one microlens. The boxes within a column are the photosensors underneath the same microlens – summing these photosensor values computes the value of a pixel in a conventional photograph that has the resolution of the microlens array.

2.2. Computing Final Photographs

Diagrams B1–B3 on Figure 1 illustrate how we process the recorded photosensor values to compute final photographs. The shaded strips on the ray-space of Figure 1-B1 illustrates the cones of rays that appears on Figure 1. A strip of constant x represents a cone of rays from the optical focal plane; a slanted strip represents a cone of




Figure 1. Overview of recording a light field and processing it to compute one pixel in an output photograph refocused closer in the world.

rays emanating from a point on a virtual focal plane. By summing along strips, we can form images focused at different depths.

The central concept of resampling is that the light traveling along each of these rays can be found by tracing the ray through the main lens optics, through the microlens array, and down to the sensor surface as illustrated macroscopically on Figure 1A and in close-up on Figure 1-B2. The radiance along each ray is estimated from sensor pixel values near the point where it intersects the sensor. Summing these radiance values computes the value of the output pixel corresponding to the desired cone of rays on Figure 1-A. Figure 1-B3 illustrates the ray-space view of resampling, in which we estimate the desired slanted strip with weighted sums of the overlapping grid cells.

Theory predicts [3] and experiments with a prototype camera corroborate [4] that refocusing power scales linearly with directional resolution. For example, if each image that forms under a microlens spans $N \times N$ pixels, then digital refocusing can compute output images where the blur is reduced by a factor of N compared to the blur that would have appeared in a conventional photograph. A different way of saying this is that one can extend the depth of focus by a factor of N .

Taking into account the aberrations of real lenses requires an extra level of detail. There are *two* sets of ray-spaces inside the camera: the *ideal* ray-space and the *aberrated* ray-space. Ideal rays are what we wish we had recorded with the light field camera, and aberrated rays are what we actually recorded. Let us differentiate between these two spaces by denoting an ideal ray as (x, y, u, v) and an aberrated ray as (x', y', u', v') .

The two ray-spaces are connected by the common space of rays in the world. An aberrated camera ray maps to a world ray via geometric refraction through the glass elements of the main lens. In contrast, an ideal camera ray maps to a world ray via tracing through an idealized approximation of the lens' optical properties that is free of aberrations. In this paper, the ideal lens is a Gaussian thick lens approximation of the real lens.

These two mappings into the world space define a mapping, C , directly from the aberrated space to the ideal space (i.e. $C(x', y', u', v') = (x, y, u, v)$). We call this map the *ray correction function*, and its inverse the *ray distortion function*. These are the fundamental mappings that must be calculated in computing digitally corrected images. Conceptually, C results from composing the mapping from aberrated rays to world rays with the inverse of the mapping from ideal rays to world rays. A procedure to compute the forward map is to take the input aberrated camera ray, trace it out into the world through the real optics of the lens, and then compute its Gaussian conjugate back into the camera.

The light field sampling grid that is recorded by the light field camera is rectilinear in the aberrated ray-space. Plotting this sampling pattern in the ideal ray-space causes it to become curved. Figure 2A shows the rays refracting through a double-convex lens, and Figure 2B shows the distortion of the sampling grid caused by the lens' aberrations. The vertical lines on the grid become S-shaped curves that are essentially ray-intercept curves [12] at different points across the field.

3. DIGITAL CORRECTION ALGORITHMS

The visualization of the correction function on the ideal ray-space provides a different way to visualize the action of traditional optical correction techniques. Two classical techniques are stopping down the lens, and adding lens elements to balance aberrations.

The logic of stopping down the lens is easy to understand in ray-space. The most highly aberrated regions come from the edges of the lens aperture (extremal u values), where the slope is greatest and the sample has a large x extent. Stopping down the aperture prevents these extremal rays from reaching the sensor pixel, reducing image blur and the effect of the aberration. Of course the price is reduced light sensitivity – much fewer rays are captured, so longer exposures are required.

In contrast to these optical techniques, digital correction is an attempt to “straighten the curves” in software. This is possible because, in collecting ray information, the light field camera essentially splits the vertical curved columns into multiple cells. At a high level, digital correction of lens aberrations is simply a repetition of the basic concept used in refocusing: resorting the rays in the recorded light field to where we ideally wanted them to converge. To determine where we want the rays to converge we will ray-trace an idealization of the lens, and




Figure 2. Ray-space illustration of digital correction of lens aberrations.

to determine where the rays actually went in the recorded light field we will ray-trace a model of the real lens. In the latter case, we must accurately model refraction through the lens elements.

Figure 2 is an overview of digital correction in terms of the ray-space, illustrating how to compute a single output pixel in a corrected image. Figure 2A illustrates a set of rays from a single point in the world, tracing into the camera through a double-convex lens. This highly aberrated lens was chosen for illustrative purposes. Figure 2B illustrates the ideal (x, u) ray-space inside the camera, with the aberrated (x', u') light field sampling grid super-imposed. Each cell in the grid represents the rays integrated by a single photosensor pixel inside the camera. The vertical gray strip represents the set of rays shown on Figure 2A. This is the strip in ray-space that we want to integrate to produce the value of an output pixel in an image free of aberrations.

Figure 2C illustrates estimation of the desired vertical strip using the recorded photosensor values. The procedure can be thought of as rasterizing the vertical strip onto the warped grid and summing the rasterized pixels. In contrast, the cells comprising one of the curved columns in Figure 2C represent all the rays collected by a single microlens in the camera. Integrating this column produces a pixel value in a conventional photograph without digital correction. Note that the spatial x extent of this curved column is wider, hence more blurry, than the digitally-corrected estimate in Figure 2C.

The non-linear distortions introduced by aberrations mean that some light field cells pollute the corrected photograph more than others. These observations motivate an optional enhancement in the resampling process for digital correction, designed to further raise the contrast and clarity of the corrected image. The idea is to weight the contribution of each photosensor pixel in inverse proportion to its spatial extent when projected onto the output image plane. This modification means computing a *weighted* average of light field sample values in the final step of the pixel-order algorithm.

In the weighted corrected images shown later in this chapter, we used the following weight function, where Δx and Δy are the projected width and height of the light field cell in the output image. For convenience, the units are in terms of output pixel widths.

$$w(\Delta x, \Delta y) = h(\Delta x) \cdot h(\Delta y), \quad \text{where}$$

$$h(x) = \begin{cases} 1, & x \leq 1 \\ \exp\left(-\frac{(1-x)^2}{2\sigma^2}\right), & x > 1 \end{cases} \quad (1)$$

In words, the weighting function decreases according to a Gaussian fall-off as the projected width of the cell increases beyond one output image pixel. The x and y dimensions are treated separately, with the overall weight being the product of the weights for each dimension. A standard deviation of $\sigma = 2$ is used for the Gaussian fall-off.

4. CORRECTING RECORDED ABERRATIONS IN A PLANO-CONVEX LENS

There were two over-arching goals to the experiments in this section. The first was to visually demonstrate that digital correction could raise contrast and resolution in real images acquired with the prototype camera. The

second goal was to use the prototype camera data to provide a measure of validation to our simulation software for computing raw light fields and digital corrections. This software is used in the last part of the paper in quantitative performance tests of a wider range of lenses at much higher light field resolutions.

The lens tested in this section is a 100 mm plano-convex lens made out of standard BK7 glass. This simple lens was chosen because it produces aberrations extreme enough to be visible and correctable in the relatively low 296×296 photographs produced by our prototype.

An aberrated light field was recorded using the prototype camera by replacing its usual photographic lens with this plano-convex lens (convex side up). A manual aperture was placed against the planar side of the lens, and stopped down to achieve an $f/4$ image-side f -number. The separation between lens and image plane was chosen to focus on a resolution test-chart approximately 24 cm away and focus was tuned by adjusting the height of the target until maximum sharpness was achieved.

A matching raw light field was simulated by Monte-Carlo ray-tracing [13]. The computer model of the plano-convex lens, microlens array and sensor were matched to the manufacturer’s physical specifications. The separation between the main lens and the microlens plane was matched to measurements on the prototype set-up. As with the physical set-up, focus was tuned by adjusting the distance of the virtual resolution chart until maximum sharpness was achieved.

Results

Figure 3 illustrates data from the prototype camera. It visually compares the quality of images computed from the recorded light field with and without correction. Column A, computed without correction, is equivalent to a conventional photograph. It exhibits the classical softness in resolution across the image plane due to spherical aberration. Significant loss in contrast is visible near the edge of the frame, such as in the zoomed view at the bottom of Column A, where regions that should be black and white appear as gray due to pollution from nearby regions of the other color. Column B illustrates that correction raises the contrast, particularly along the edges of the image but less so in the center of the frame where aberrations are less. Column C illustrates that weighted correction raises the contrast and resolution further across the image.

Figure 4 compares the recorded data with the simulated version. Images A1 and B1 compare close-ups of the raw light field data. Even at this extreme level of zoom, the overall match is quite good, although small differences are visible due to error in calibrating the physical and virtual geometry. Figure 4-A2 and B2 illustrate that these calibration errors cause only small differences in output images. These two images are for correction without weighting, and similarly good agreement is found in uncorrected images and correction with weighting.

Comparison of Weighted Correction with Reduced-Aperture Conventional Imaging

In the plano-convex experiment, weighted correction results in an average weight of 32% for the cells in the light field, which in some sense matches the light gathering power of a conventional camera with an aperture reduced to 32% area. However, stopping down the lens imposes the same sub-aperture on every output image pixel. Weighted correction provides the extra freedom of varying the aperture across the image plane, which results in sharper images.

Figure 5 compares the corrected image using weighting against a conventional image where the aperture is reduced to 32% area (56% diameter). The conventional image was computed by only summing the rays under each microlens that passed through the reduced aperture, without resorting of rays. Although the aberrations in the stopped-down conventional image are reduced compared to the full-aperture version in Figure 3, the weighted correction still provides significantly better contrast. For example, the black bars at the top and bottom are much darker in the corrected image.

Weighted correction produces a superior image for two reasons. First, it re-sorts rays, which results in improvements in the convergence of rays that are used. Second, weighted correction has much greater flexibility in choosing to use rays that converge well on the resolution chart. In the conventional case, stopping down the lens excludes sub-aperture images from the periphery of the lens that tend to contain a larger fraction of blurry pixels, but even the reduced aperture contains some of these artifacts. In contrast, weighted correction can use an effectively larger aperture and discard the worst rays.

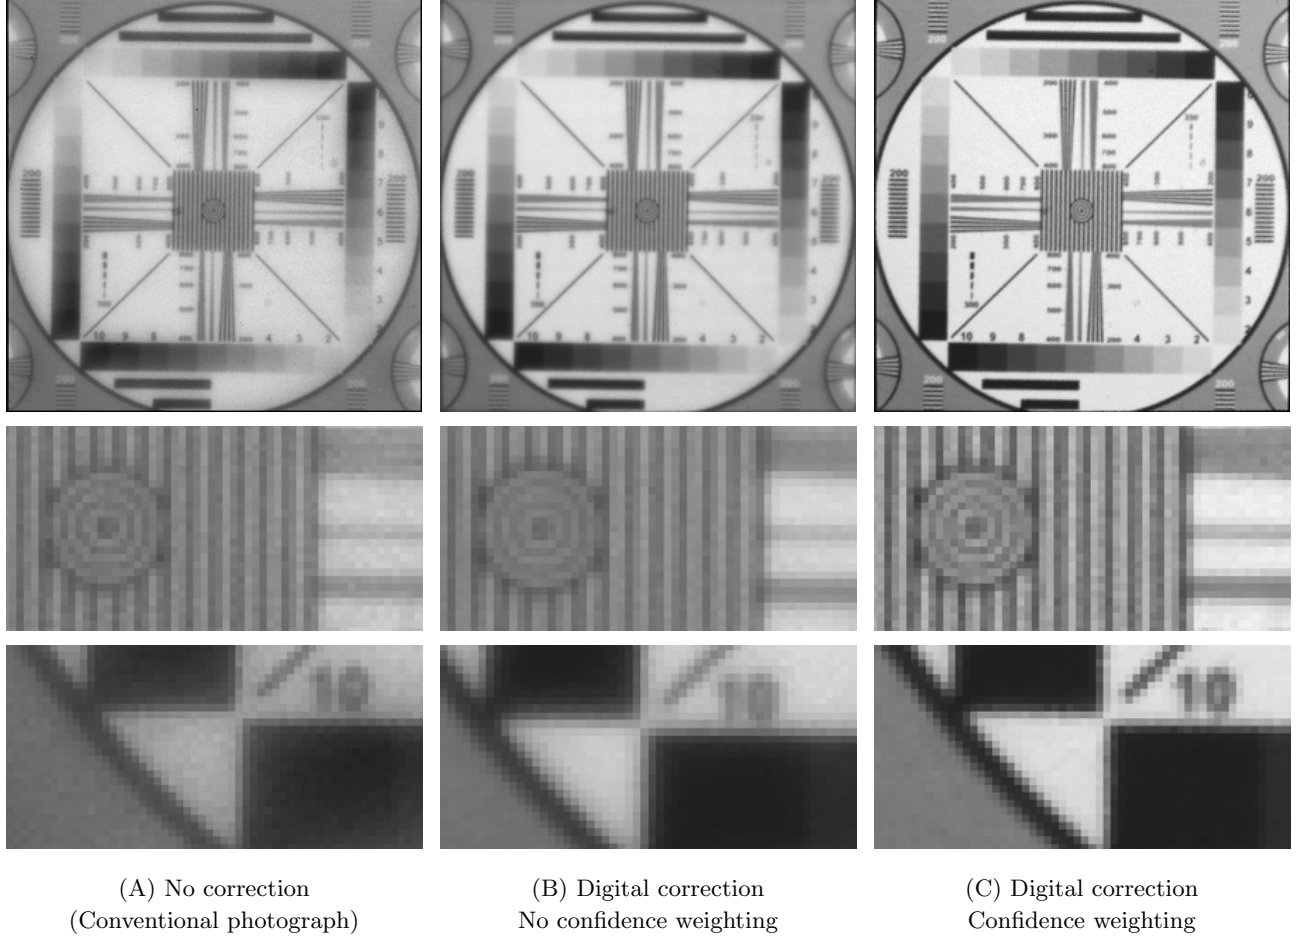


Figure 3. Comparison of uncorrected and corrected images from a light field recorded with the prototype camera.

5. SIMULATED CORRECTION PERFORMANCE

In contrast to the previous section, the results in this section derive purely from computer simulation. The experiments here apply traditional, numerical analyses [12] of the ray-traced point spread function (PSF) to compare the performance of various lenses with and without digital correction. The PSF is the spot of energy that appears in an output image response to a point source of light in the world. Ideal imaging produces a diffraction-limited spot, but in practice aberrations usually result in a larger blur. One of the main goals of this section was to explore how digital correction works across a range of lenses. Is it likely to be a general-purpose technique for improving the quality of lenses in optical engineering? To keep the comparison as simple as possible, this section examines only digital correction without weighting.

5.1. Methods and Image Quality Metrics

The cameras simulated in this section assume the following geometry and resolutions. A 35 mm format sensor is assumed – that is, the microlens array and photosensor both measure 36 mm×24 mm. The spatial resolution, that is the resolution of the microlens array, is assumed to be constant at 1800×1200 (2.1 mp). A range of $N \times N$ directional resolutions are tested, from $N = 1$ (uncorrected), up to $N = 16$. Since the spatial resolution and sensor size are fixed, increasing N assumes increasing photosensor resolution. $N = 10$, requiring 1.9 micron pixels, lies at the limit of technology currently shipping in commodity cameras [1]. 1.7 micron pixels have been

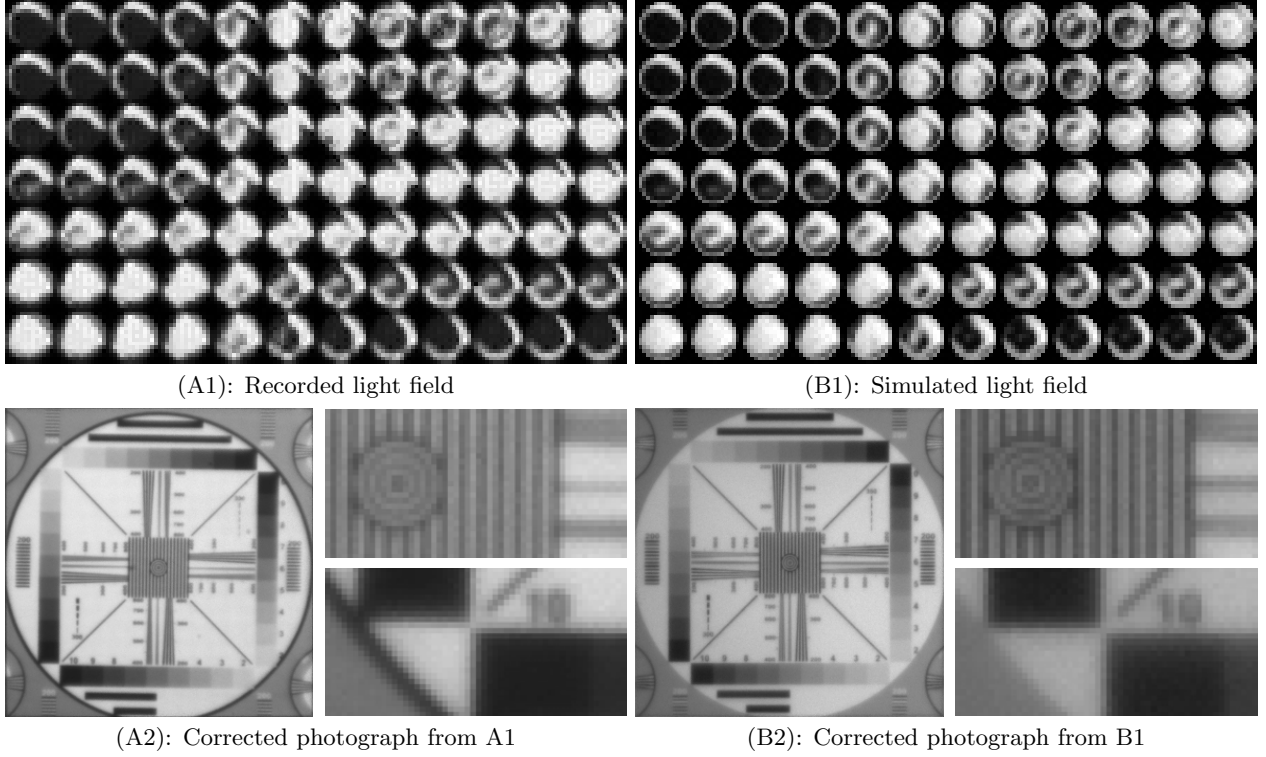


Figure 4. Comparison of recorded and simulated data for digital lens correction.




Figure 5. Comparison of weighted correction with conventional imaging where the lens aperture is stopped down for equivalent light gathering power.

demonstrated in the laboratory [14], and the maximum resolution simulated here, $N = 16$, assumes a further reduction by 26% down to 1.25 micron.

The imaging performance of various lenses on these cameras was quantified by computing PSFs to analyze imaging performance at different points on the imaging plane. Computing a PSF means tracing rays from a point light source in the world through all parts of the lens, down to the imaging plane to produce an image. In the case of simulating a light field camera, the rays are traced through the microlens array down to the photosensor surface to produce a raw light field. Final corrected photographs of the PSF were computed from this raw light field.

Although the PSF can be used to provide very detailed analysis of imaging performance, it can be cumbersome for comparison across lenses because it is very high-dimensional, varying with position on the imaging plane, as well as focal depth.

To compare performance across lenses, we used a series of summary statistics derived from the PSF. The first level of summary is provided by the root mean square (RMS) spot radius. The RMS spot radius can be thought of as the standard deviation of the PSF interpreted as a probability distribution function. This provides a very compact summary of the gross size, but discards information about the shape of the spot, which can greatly affect the qualitative appearance of aberrations in the final image. For example, Figure 7B illustrates an anisotropic PSF exhibiting coma, which in extreme cases is well known to cause objects to appear as if they are “flying out of the field of view” [12]. The overlaid RMS circle clearly does not capture this visual characteristic, giving only the standard deviation.

The RMS spot radius is a large reduction in information, but it is still a 2D function over the image plane. The next level of summary is provided by averaging the RMS measure across the 2D plane, producing a single-number summary of the entire imaging system’s performance. Although it is somewhat crude, this measure does track performance trends faithfully, improving with the quality of the photographic lens and with the amount of directional resolution available in digital correction.

An equivalent, but slightly more intuitive measure of average RMS radius is *effective resolution*. The thought underlying the concept of effective resolution is that the total number of pixels in the output image is irrelevant if the PSF is much broader than one output image pixel. Effective resolution provides a measure of the number of PSF-limited spots that we can discriminate in the output image. We define the effective pixel size as the square that can be inscribed within a circle of one average RMS spot radius. The effective resolution is then just the number of squares that would fit within the sensor. In the experiments below, the effective resolution is therefore defined to be $\frac{24 \text{ mm} \times 36 \text{ mm}}{(\sqrt{2}R)^2}$, where R is the average RMS spot radius (in mm), and $\sqrt{2}R$ is the width of the inscribed square.

The final measure of image quality is the MTF. We find the MTF by computing the Fourier transform of the point spread function. A detail of this implementation is that it incorporates the effective MTF due to the finite sized microlenses and pixels of the sensor [15]. This means that the computed MTF does not exceed the Nyquist sampling rate of the 20 micron spatial sampling grid. Since the PSF varies across the image, we plot the MTF as a function of distance from the optical axis on the image plane.

5.2. Case Analysis: Cooke Triplet Lens

Let us first analyze the PSF of a particular lens in some detail. The chosen lens is an $f/2.9$ Cooke triplet [16]. Figure 6 illustrates the tracing of rays through the lens’ three glass elements in the simulation of the light field camera to calculate PSFs. Columns A, B and C show rays converging on three different positions in the image: at the optical axis, half-way towards the edge of the image, and at the edge.

The middle row zooms in on a 200 micron patch of the imaging plane at the convergence of rays. 10 microlenses are shown across the field of view, and the rays terminate at the bottom on the photosensor surface. These close-up views illustrate the complexity of the ray structure inside the camera. The shape of rays converging on the microlens plane is far from a simple cone as assumed in ideal imaging. In addition, the shape of rays and quality of convergence change across the imaging plane. For example, Diagram A2 shows that the rays in the center of the frame converge onto 2 microlenses. In contrast, Diagram C2 shows worse convergence near the edge of the image, with rays spread over approximately 6 microlenses.

The bottom row of Figure 6 provides a different view of the aberrations, in ray-space. The horizontal, x , field of view is 200 microns, but the vertical, u , field of view spans the entire lens plane. The directional resolution is $N = 16$. The distortion of the sampling grid provides more information about the nature of the aberrations at different parts of the field. In Diagram A3, one can see how optical correction has been used to force rays coming from the edge of the lens to converge reasonably well. In fact, the worst aberrations come from the parts of the lens midway between the optical axis and the edge of the lens. In contrast, Diagram C3 shows that at the right edge of the field, the worst aberrations come from the left-most portion of the lens, since the grid is most distorted near the bottom.

These ray-diagrams highlight the crucial concept of scale relative to image resolution when considering the seriousness of aberrations. If the spatial resolution were less, the curvature would be negligible relative to the spacing of the grid columns. Conversely, any residual aberration will exhibit significant curvature if the spatial resolution is high enough. This means that to correct for aberrations, the number of vertical cuts needed in each column, corresponding to the directional resolution, is proportional to the spatial resolution. This makes intuitive sense: increasing the spatial resolution increases our sensitivity to aberrations in the output image, and raises the amount of directional resolution required to correct for the residual aberrations.

Figure 7 illustrates the uncorrected and corrected PSFs for the triplet lens at the three positions shown in Figure 6. The overlaid circles show the RMS spot. With this amount of directional resolution (16×16), the average RMS spot radius is roughly 14 microns, and the effective resolution of output images is close to the full 2.1 mp of the microlens array. Figure 7 illustrates an important characteristic of digital correction that is not captured by the summary statistics. It shows that digital correction tends to make the PSF more radially symmetric, reducing many of the qualitative problems introduced by aberrations such as coma.

Figure 8 illustrates the improvement in correction performance as the directional resolution ($N \times N$) increases from $N = 1$ (uncorrected conventional imaging) to 16. The histograms for each value of N illustrate the distribution of RMS radii across the imaging plane. The vertical axis of the graph is measured in terms of output image pixels. For example, for $N = 1$, 68% of the imaging plane has an RMS spot radius of approximately 2 pixel widths. These histograms were computed by sampling the PSF at over 1000 different positions distributed evenly but randomly (using stratified sampling) across the $36 \text{ mm} \times 24 \text{ mm}$ imaging plane. 65,536 rays were cast in estimating the RMS radius for each PSF.

The histograms illustrate how the corrected RMS spot size converges as the di-

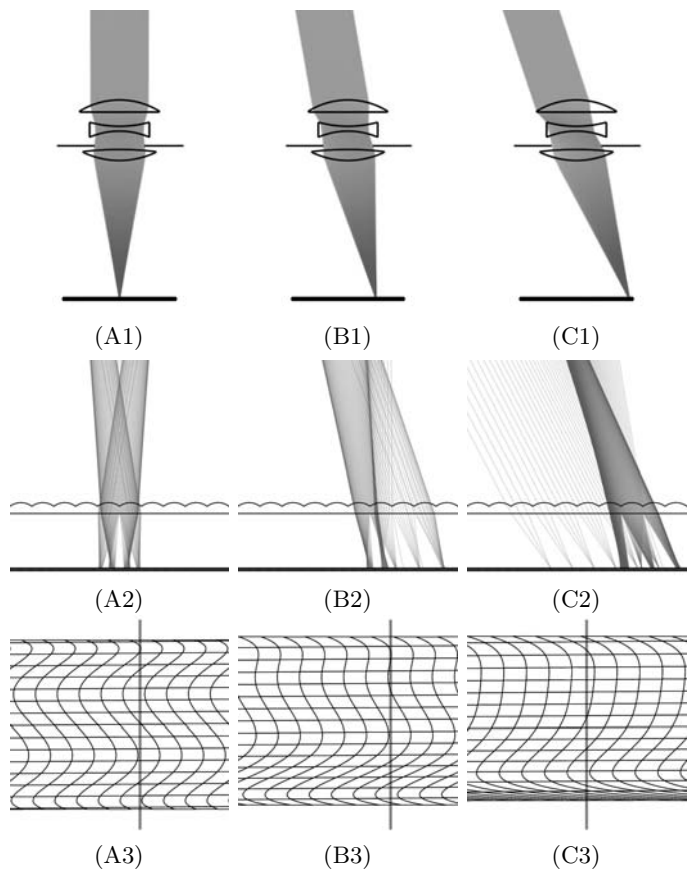


Figure 6. Aberrated ray-trace and ray-space of a Cooke triplet lens.

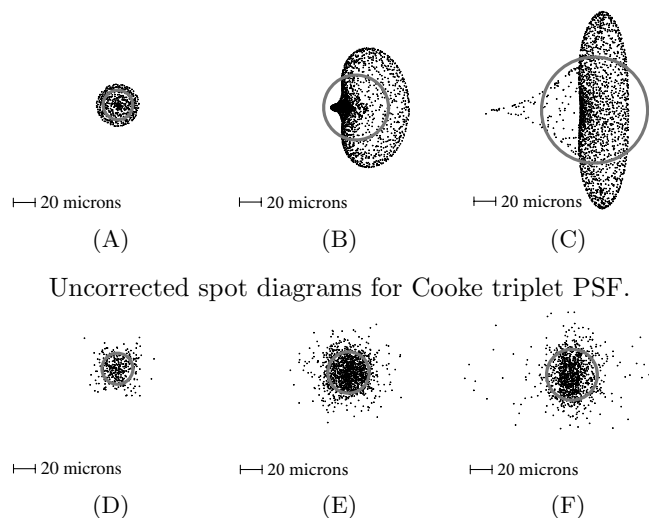


Figure 7. Digitally corrected spot diagrams for Cooke triplet PSF at field positions shown on Figure 6.

rectional resolution increases. The rate of convergence depends on how much aberration is present in the lens, with greater distortions requiring larger amounts of directional resolution for accurate correction. The effective resolution of the output image grows from 0.28 mp to the full 2.2 mp as the directional resolution increases from $N = 1$ to 16. The starting, uncorrected resolution may seem surprisingly low. One factor is that we are using the lens with its aperture wide-open, where aberrations are worst. Second, effective resolution decays rapidly as the spot size increases. If the effective resolvable spot covers just 3 pixels, as it almost does in the triplet lens without correction, then the

The histograms give a sense for how the average performance increases across the image, but not how the performance varies from the center to the edge. Figure 9 measures variation along that axis, in the form of MTF plots. The Nyquist spatial sampling rate is 25 cycles per mm, and the graphs plot the MTF at three spatial frequencies: 1/2, 1/4 and 1/8 the Nyquist rate.

Figure 9A shows that, even without digital correction, the contrast of the triplet lens focused at infinity is quite good across the imaging plane, dipping only slightly at the very edge of the image. Sharpness, however, is relatively poor at these resolutions, especially out towards the edge. Digital correction improves sharpness and contrast across the field of view, and makes performance more even.

Figure 9B illustrates the triplet focused at a distance of 100 mm for 1:1 magnification (i.e. macro photography). The lens was not designed for such close focusing, and its performance is quite poor at this depth. The corrected MTF in Figure 9B illustrates dramatic improvement in the contrast of the system. This example illustrates that digital correction *extends the useful focal range* of a lens. However, the improvement at macro focus is less than for infinity focus. There are at least two reasons for this phenomenon. First, the light field distortion is much greater at the macro focal depth. Figure 10 illustrates this fact by comparing the ray-space for infinity and macro focus at the three image positions shown in Figure 6. From the discussion previously in this chapter, it is evident that a higher directional resolution would be required to correct the greater aberrations at the macro focus.

Unfortunately, the second factor is that the directional resolution is halved at macro focal depths. This can be seen in Figure 10 by the fact that there are only half as many vertical cuts in the grid columns of D-F as there are in A-C. The underlying cause of the reduction is that the separation between the main lens and the imaging plane increases by a factor of 2 when changing focus from infinity to macro. As a result, the aperture




Figure 8. Histogram of triplet-lens PSF size across imaging plane.




Figure 9. MTF of triplet lens for infinity and macro focus, with and without correction.

appears half as large radially from the perspective of the microlenses, and the images that appear under each microlens span only half as many pixels.

These observations suggest that in designing lenses to be used with a light field camera, the designer should optimize optical image quality for close focus distances rather than infinity. Higher directional resolution in the recorded light field further focal distances can be used to offset slightly worse optical performance in that range.

5.3. Correction Performance Across a Database of Lenses

The previous section illustrated how digital correction could be used to improve the performance of one multi-element lens. It should be clear from the ray-traces diagrams and visualizations of the aberrated ray-space that the amount of improvement will depend on the exact formula of the lens and the shape of the distortions in the recorded light field.

To evaluate the performance of digital correction across a range of lenses, we conclude with a summary of simulated correction performance for a database of 22 lenses. The optical formulas for these lenses were obtained by manually extracting every fixed-focal-length lens in the Zabase [17] database for which the description implied that photographic application was possible.

The initial set of lenses selected from Zabase was modified in two ways. First, the set was pruned to exclude lenses that contained aspheric elements, because the simulation system does not currently support analytic ray-tracing of such aspheric surfaces. Second, the remaining lenses were linearly scaled in size so that their projected image circles matched the diagonal length of the 35 mm format sensor in our virtual camera. The issue is that the scaling of the original lens formulas was quite arbitrary. For example, many formulas were normalized in size to a nominal focal length of 100 mm, regardless of the effective focal length in 35 mm format. Figure 11 lists the final set of lenses used in this experiment, and some of their basic properties.

The database spans an eclectic range of design forms, including triplets, double Gauss variations, telephotos and wide-angle lenses. They also span a fairly wide range of focal lengths and f -numbers. Many of these lens recipes originated in classical optical engineering textbooks, or from expired patents. Nevertheless, many of the designs are quite old, and this may best be thought of as an experiment on variations of some classic design patterns. There are thousands of lens design forms used in many diverse imaging applications today, and this database does not attempt to sample the breadth of current art.

In spite of these caveats, the 22-lens database provides useful experimental intuition. The results imply that digital correction will be quite generally applicable at resolutions that are achievable in present or near-term technology.

Figure 12 presents the experimental data. Each lens appears on a different row, sorted in order of increasing uncorrected performance. For each lens, effective resolution with unweighted correction is plotted across the row as a function of directional resolution. The directional resolution is $N \times N$, beginning at $N = 1$ (conventional uncorrected), and increasing up to $N = 16$, as shown by the key near the bottom of the figure. Note that the horizontal axis is plotted on a log scale.

The most striking feature of the data is the regular improvement provided with digital correction. Effective resolution rises steadily with directional resolution for all the lenses, indeed as one would expect from the

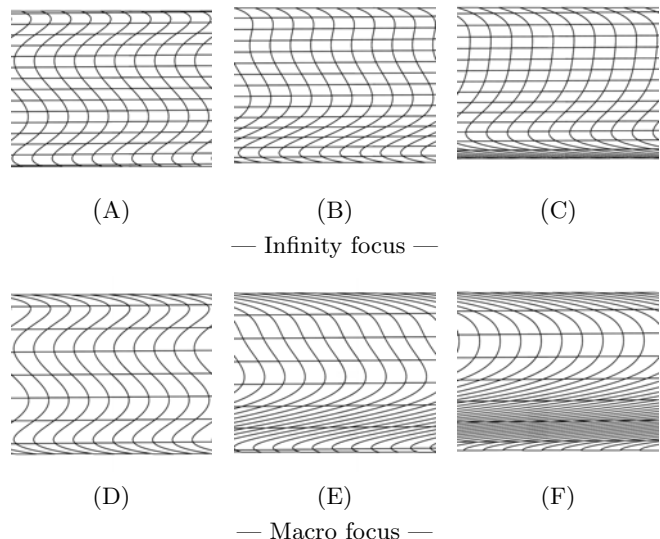


Figure 10. Ray-space of Cooke triplet lens at two focal distances, for field positions shown in Figure 6.

Lens File	Description	Glass elements	Focal length	Aperture	Row*
E-001	Cooke triplet anastigmat	3	135 mm	$f/3.5$	13
E-002	Triplet	3	130 mm	$f/2.0$	1
E-006	Triplet objective USP 2,453,260	3	130 mm	$f/2.7$	5
E-009	Cooke triplet USP 3,176,582	3	52 mm	$f/2.9$	15
F-020	Wide angle USP 3,976,465	10	23 mm	$f/1.8$	12
F-025	Wide-angle objective	6	36 mm	$f/3.0$	17
G-001	Telephoto	5	130 mm	$f/5.8$	18
G-002	SLR telephoto	9	191 mm	$f/2.8$	22
I-014	Photographic objective	5	63 mm	$f/3.4$	20
L-001	Double Gauss	6	63 mm	$f/2.5$	16
L-002	Double Gauss	7	63 mm	$f/1.8$	7
L-008	Photographic lens USP 2,350,035	7	85 mm	$f/1.4$	2
L-009	Photographic objective USP 2,379,392	7	85 mm	$f/1.5$	9
L-010	Large aperture objective USP 2,701,982	8	85 mm	$f/1.0$	4
L-011	Modified Gauss USP 3,743,387	7	58 mm	$f/1.2$	6
L-012	Large aperture photo lens USP 3,938,884	7	58 mm	$f/1.2$	3
L-013	Lrg. aper. photographic USP 4,364,643	7	54 mm	$f/1.2$	8
L-014	Gauss type objective USP 4,390,252	6	85 mm	$f/2.5$	21
L-015	Double gauss objective USP 3,552,829	7	63 mm	$f/1.6$	14
L-020	8-glass double Gauss USP 2,701,982	8	63 mm	$f/1.1$	11
L-033	Momiyama USP 4,364,643	7	54 mm	$f/1.2$	10
L-035	Double Gauss objective	6	63 mm	$f/2.5$	9

* In Figure 12.

Figure 11. Database of lenses.

interpretation of the algorithm in ray-space. In other words, the experiment provides evidence that digital correction is a robust and general-purpose technique for improving image quality.

A second important point is that the data show that in many cases the amount of improvement is roughly proportional to the recorded directional resolution. In particular, the experiment indicates that relatively moderate amount of directional data can provide significant improvements in image quality. For example, a directional resolution of 8×8 , which is viable in current VLSI technology, produces increases in effective resolution by an average factor of 5.4 across the database, and by up to a factor of 15 (for rows 1 and 7). Taking a step back, this is an important validation of the required scale of digital correction. From first principles it was not immediately clear whether correcting for aberrations across the lens aperture would work without extremely high resolution images of the lens aperture. The experiment here shows that it does, although of course more resolution provides greater effective resolution up to the limiting resolution of the microlens array.

Another feature of the data is that the effective resolution without digital correction (the left-most data point for each row) is quite low for many lenses. This corresponds to the performance of conventional photographic imaging. For 14 out of the 22 lenses the effective resolution is actually below one quarter of the assumed 2.2 mp output image resolution. Although some lenses are unusually low, and probably reflect a prioritization of maximum light gathering power (aperture size) over image quality, it is not surprising that many lenses would suffer from low effective resolution with their apertures wide-open, as tested here. It is common photographic wisdom that in order to obtain the sharpest images from any lens, one should stop the lens down to half its maximum aperture or less [18]. Of course this reduces the light gathering power, requiring longer exposures. The results here re-iterate that digital correction can raise the image quality without having to pay as great a penalty in light gathering power.

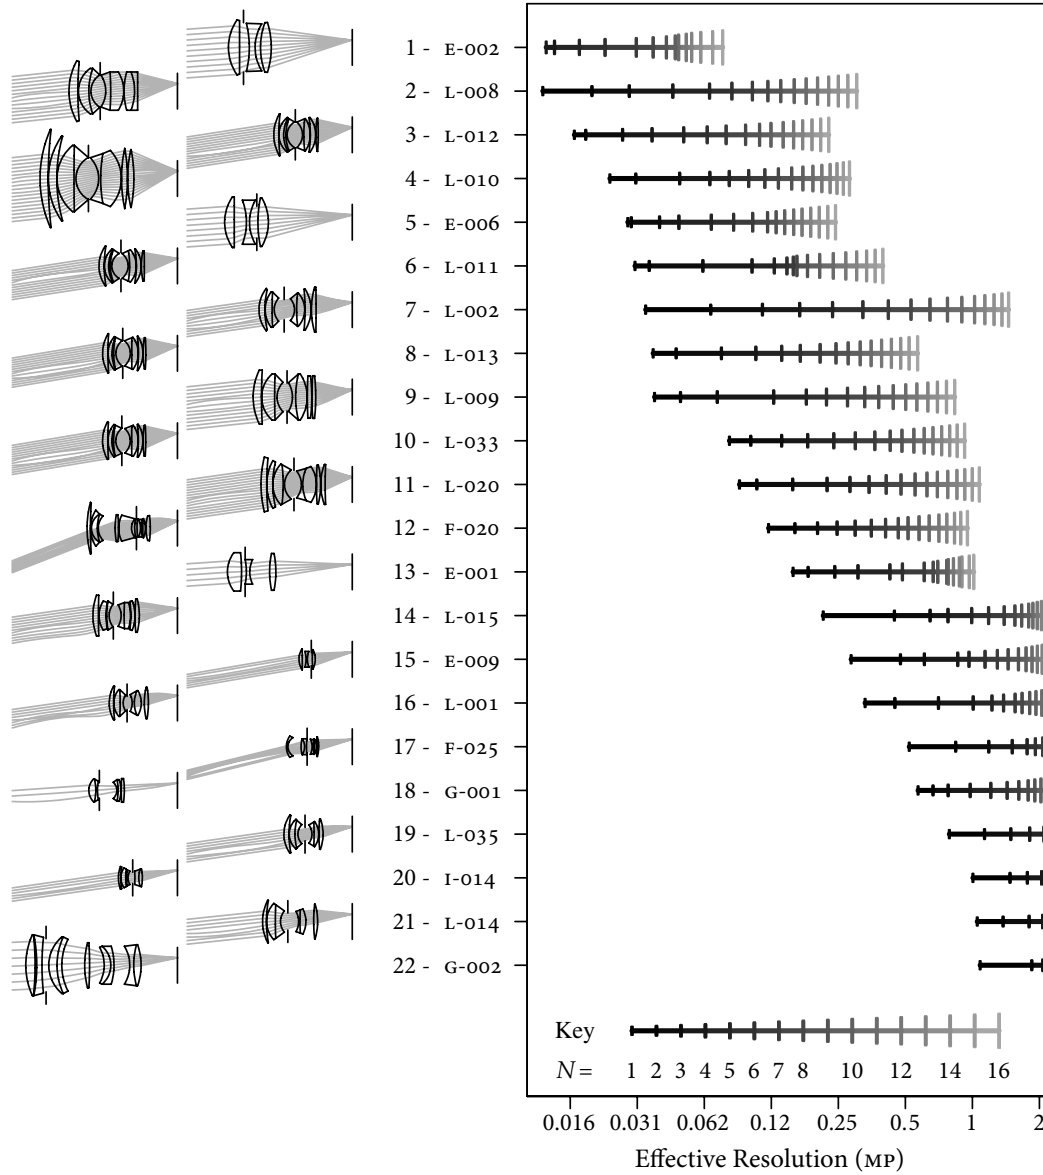


Figure 12. Improvement in effective resolution with digital correction.

In Closing

The experiments reported in this paper are just the beginning. To exploit digital correction to maximum advantage, it must be considered during the optical design phase of the lens, while determining its optical requirements. For example, digital correction will expand the viable range of f -numbers for a desired output image resolution. Alternatively, it may enable simplification of the optical formula (e.g. a reduction in the number of glass elements) required to attain a target image quality.

In order to reach the full potential of the technique, further study is needed into methods for joint optimization of the the optical recipe and digital correction software. As one example of the complexity of the design task, rows 5 – 9 in Figure 12 illustrate very different growth rates with directional resolution, even though they all start with similar uncorrected effective resolutions. From a qualitative examination of the distortions in the ray-space,

it seem plausible that a major limitation on the rate of improvement are regions of extremely high slope that are localized in u , such as uncorrected spherical aberration. On the one hand, this suggests that unweighted digital correction would work best on lenses without this characteristic. On the other hand, localized regions of high curvature in the ray space are especially amenable to weighted correction, which can eliminate them without greatly reducing the collecting power.

These fundamental non-linearities are what make the study of aberrations and their correction such a challenging subject. Continued study of digital correction will have to address much of the complexity of traditional lens design, and add to it the additional nuances of weighting function design.

REFERENCES

1. P. Askey, "Digital cameras timeline: 2006," 2006. Digital Photography Review, <http://www.dpreview.com/reviews/timeline.asp?start=2006>.
2. B. W. Keelan, *Handbook of Image Quality : Characterization and Prediction*, New York: Marcel Dekker, 2002.
3. R. Ng, "Fourier slice photography," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2005)* **24**(3), pp. 735–744, 2005.
4. R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, "Light field photography with a hand-held plenoptic camera," Tech. Rep. CSTR 2005-02, Stanford University Computer Science, 2005.
5. R. Ng, *Digital Light Field Photography*. PhD thesis, Stanford University, June 2006.
6. J. C. Maxwell, "On the general laws of optical instruments," *The Quarterly Journal of Pure and Applied Mathematics* **2**, pp. 233–246, 1858.
7. J. Dickerson and G. Lepp, *Canon Lenses*, Sterling Pub. Co. Inc., 1995.
8. E. H. Adelson and J. Y. A. Wang, "Single lens stereo with a plenoptic camera," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**, pp. 99–106, Feb 1992.
9. G. Lippmann, "Epreuves reversibles donnant la sensation du relief," *J. Phys.* **7**, pp. 821–825, 1908.
10. H. E. Ives, "Optical properties of a Lippmann lenticulated sheet," *J. Opt. Soc. Amer.* **21**, pp. 171–176, 1931.
11. J. Arai, M. Okui, M. Kobayashi, and F. Okano, "Geometrical effects of positional errors in integral photography," *J. Opt. Soc. Am. A* **21**, pp. 951–958, June 2004 2004.
12. W. J. Smith, *Modern Lens Design, 2nd edition*, New York: McGraw-Hill, 2005.
13. M. Pharr and G. Humphreys, *Physically Based Rendering: From Theory to Implementation*, Morgan Kaufmann, 2004.
14. Micron, "Micron Technology, Inc. demonstrates leading imaging technology with industry's first 1.7-micron pixel CMOS sensor," 2005. http://www.micron.com/news/product/2005-05-16_first_1-7_pixel_cmos_sensor.html.
15. S. K. Park, R. Schowengerdt, and M.-A. Kaczynski, "Modulation-transfer-function analysis for sampled image systems," *Applied Optics* **23**(15), 1984.
16. E. Tronnier and J. Eggert, "Three lens photographic objective." United States Patent 3,176,582, 1960.
17. Zemax, *Zemax, Optical Design Database, Version 5.0*, San Diego: Zemax Development Corp, 2003.
18. A. Adams, *The Camera*, Bulfinch Press, 1995.