

Gigapixel Computational Imaging

Oliver S. Cossairt

ollie@cs.columbia.edu

Daniel Miao

dm2701@columbia.edu

Shree K. Nayar

nayar@cs.columbia.edu

Columbia University

120th St, New York, NY 10027

Abstract

Today, consumer cameras produce photographs with tens of millions of pixels. The recent trend in image sensor resolution seems to suggest that we will soon have cameras with billions of pixels. However, the resolution of any camera is fundamentally limited by geometric aberrations. We derive a scaling law that shows that, by using computations to correct for aberrations, we can create cameras with unprecedented resolution that have low lens complexity and compact form factor.

In this paper, we present an architecture for gigapixel imaging that is compact and utilizes a simple optical design. The architecture consists of a ball lens shared by several small planar sensors, and a post-capture image processing stage. Several variants of this architecture are shown for capturing a contiguous hemispherical field of view as well as a complete spherical field of view. We demonstrate the effectiveness of our architecture by showing example images captured with two proof-of-concept gigapixel cameras.

1. Introduction

Today's high-resolution cameras capture images with pixel counts in the tens of millions. When digital cameras can produce images with billions of pixels, they will usher in a new era for photography. A gigapixel image has such a tremendous amount of information that one can explore minute details of the scene (see Figure 2). Gigapixel images are fascinating because they capture orders of magnitude more detail than the human eye, revealing information that was completely imperceptible to the photographer at the time of capture. At present, highly specialized gigapixel imaging systems are being developed for aerial surveillance [5].


To be generally useful to photographers, a gigapixel camera must not differ drastically in terms of form and function from today's consumer grade megapixel cameras. It should have at least a 40° Field-Of-View (FOV). To be useful in both indoor and outdoor lighting conditions, the camera should operate with a small F/#, say between F/4-F/11. To be portable it must not be prohibitively large or heavy. Finally, a design with low complexity will ensure the camera can be manufactured at low cost.

Why are there no gigapixel cameras commercially avail-

able today? CMOS and CCD technologies have improved to the point that imaging sensors with pixels in the $1\mu\text{m}$ range have been demonstrated [9]. It is certainly within the reach of manufacturing technology to produce sensors with 1 billion pixels. On the other hand, it remains a huge challenge to design and manufacture lenses which have the resolving power to match the resolution of such a sensor. This is because the number of resolvable points for a lens, referred to as the Space-Bandwidth Product (SBP) [11], is fundamentally limited by geometrical aberrations. Ideally, all lenses would be diffraction limited so that increasing the scale of a lens while keeping FOV fixed would increase SBP. Unfortunately, SBP reaches a limit due to geometrical aberrations.

There are two common approaches that are taken to increase SBP in the face of this fundamental limit. The first is to just accept the loss in resolution and increase sensor size. As an example, consider the commercially available F/8 500mm focal length Schneider Apo-Symmar lens. If this lens were diffraction limited, it would be capable of resolving a gigapixel image on a 5"×5" sensor. However, because of geometrical aberrations, a sensor size of nearly 12"×12" is necessary to resolve a full gigapixel image.

The second approach taken to increase SBP is to increase complexity as a lens is scaled up. Introducing more optical surfaces increases the degrees of freedom in lens optimization, which can be used to reduce geometric aberrations and achieve diffraction limited performance. Consider the F/4 75mm focal length lens shown in Figure 1. The lens is diffraction limited over a 60° FOV so that a gigapixel image can be resolved on a 75mm×75mm surface, much



(a) An F/4 75mm focal length lens (b) The MTF of the lens in (a)

Figure 1. (a) An F/4 75mm lens design capable of imaging one gigapixel onto a 75x75mm sensor. This lens requires 11 elements to maintain diffraction limited performance over a 60° FOV. (b) The MTF at different field positions on the sensor.

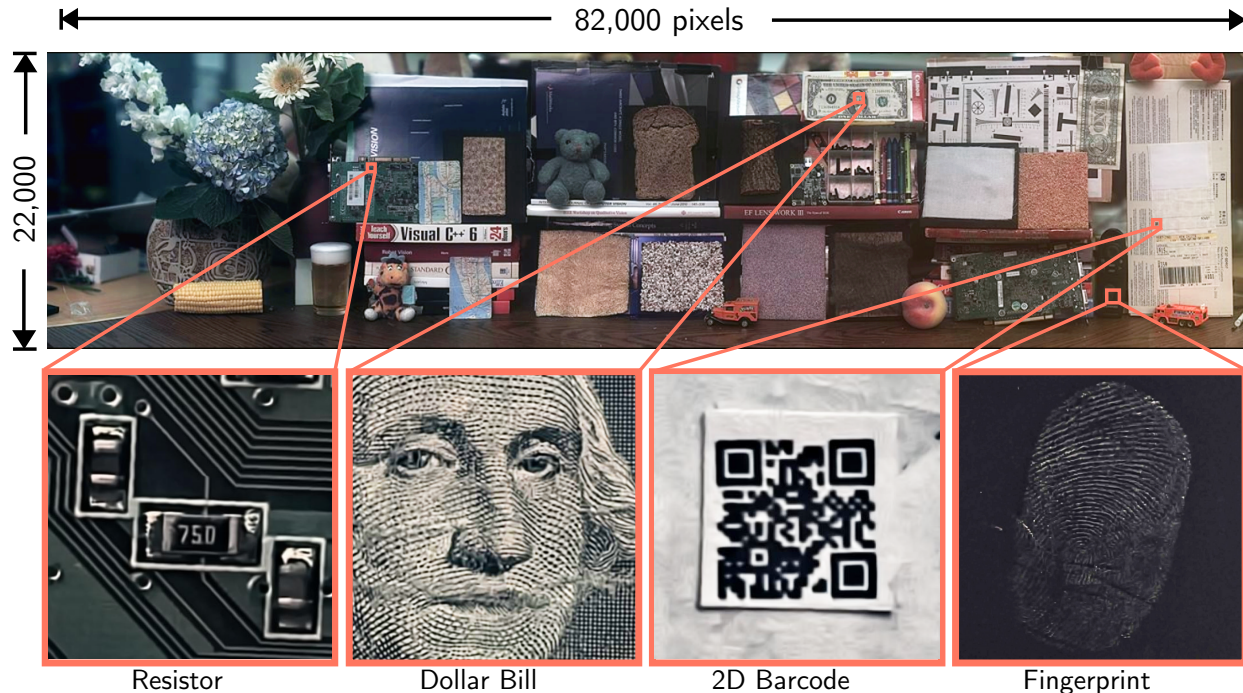


Figure 2. A 1.7 gigapixel image captured using the implementation shown in Figure 9. The image dimensions are 82,000 x 22,000 pixels, and the scene occupies a $126^\circ \times 32^\circ$ FOV. From left to right, insets reveal the label of a resistor on a PCB board, the stippling print pattern on a dollar bill, a miniature 2D barcode pattern, and the fine ridges of a fingerprint on a remote control. The insets are generated by applying a 60x-200x digital zoom to the above gigapixel image. Please see the Supplementary Material to view this example in more detail.

smaller than for the Apo-Symmar. The increase in performance comes at a great cost, however. The design consists of 11 different elements, ranging from 60-100mm in diameter, resulting in a lens that is both expensive to produce and difficult to align.

The following are the main contributions of our paper.

A Scaling Law for Computational Imaging. We present a new approach to increase SBP - the use of computations to correct for geometrical aberrations. In conventional lens design, resolution is limited by the spot size of the lens. For a lens with aberrations, spot size increases linearly with the scale of the lens. For a computational imaging system, resolution is related to deblurring error. *We observe, however, that for a lens with spherical aberrations, deblurring error does not increase linearly with lens scale.* We use this remarkable fact to derive a scaling law that shows that computational imaging can be used to develop cameras with very high resolution while maintaining low complexity and small size.

Two Proof-of-Concept Gigapixel Cameras. We present an imaging architecture that consists of a large ball lens shared by an array of small planar sensors coupled with a deblurring step. Due to our monocentric optical design, field-dependent aberrations are suppressed, and the primary aberrations are spherical and axial chromatic, which are known to code images in a manner that is invertible via post-processing [27] [26] [12] [3]. We demonstrate two

proof-of-concept gigapixel cameras. The first is a single element monocentric camera that is implemented by sequentially scanning a single sensor to emulate an array of tiled sensors. The second is a system consisting of 5 tiled sensors that produce a contiguous FOV.

Several Novel Gigapixel Camera Designs. In addition to demonstrating two proof-of-concept gigapixel cameras, we present two additional designs: one for a single element gigapixel camera with a contiguous FOV, and one for a spherical 2π FOV gigapixel camera.

2. Related Work

Large Format Imaging Systems. A few custom gigapixel imaging systems have been developed using large format lenses. These include systems built with commercial lenses that sequentially scan a large image plane surface [1] [29], as well as a system with a custom lens that is photographed on film and later converted to a digital image [10]. These are special purpose cameras that are extremely large (FL > 500mm) and therefore not suitable for commercial photography.

Camera Arrays and Multiscale Optics. Camera arrays have been used to capture high resolution images by tiling multiple sensors paired with a complex lens [30] [24]. However, a camera array for gigapixel imaging would be prohibitively large and expensive because it would require tiling an array of long focal length lenses. A related ap-

proach taken by Brady and Hagen [2] is to use a multiscale optical system consisting of a large single element lens coupled with an array of smaller optical elements, each unique and coupled with a different sensor. The advantage of this approach is that it is a compact design that can correct for geometrical aberrations. The disadvantage is that the system requires a large number of different optical elements, which may be difficult to manufacture and align.

Monocentric Optics and Curved Sensors. Monocentric optical designs are free of field dependent aberrations because they are completely symmetric: the image plane and each lens surface lay on concentric spheres. Monocentric designs date back to the Sutton Panoramic Lens (1859), and later the Baker Ball Lens (1942) [14]. Luneburg proposed the use of a monocentric lens with varying index of refraction to correct for aberrations [21]. Rim et. al proposed a small diffraction limited camera consisting of a ball lens and curved sensor [25]. Krishnan and Nayar proposed the use of a large ball lens and spherical sensor together with deblurring to create a single viewpoint, fully spherical FOV camera [16]. While several researchers have made progress towards developing curved sensors [6] [15] [17], the technology is not yet ready for commercialization.

Recently, Marks and Brady proposed a 7-element large format monocentric lens called the Gigagon [22], which the authors suggest using with a large array of planar sensors. To our knowledge this system has yet to be implemented, but is similar in architecture to some of the designs we propose. Our approach is fundamentally different in that we show how computations can be used to achieve the desired resolution while reducing complexity.

Computational Imaging. In the 90's, Cathey and Dowski proposed a hybrid optical-signal processing system which uses a cubic phase plate to extended depth of field [7]. Later they showed that the same element can be used to reduce the complexity of infrared cameras [8]. Robinson and Stork observed that spherical aberrations are easily invertible via image processing, and proposed the use of simpler lens designs based on this principle [27] [26] [28]. Guichard et. al [12] and Cossairt and Nayar [3] observed that the effects of axial chromatic aberrations can be inverted using a method that is inexact, but produces images that look good.

3. The Curse of Aberrations

Lohmann originally observed that lenses obey certain scaling laws that determine how resolution increases as a function of lens size [19]. Consider a lens with focal length f , aperture diameter D , and image size Δx by Δy . If we scale the lens by a factor of M , then $f, D, \Delta x$ by Δy are all scaled by M , but the F/# and FOV of the lens remain unchanged. If, when we scale the lens, the minimum resolvable spot size has not also increased by a factor of M ,

then we have increased the total number of points that can be resolved. The number of resolvable points for a lens is referred to as the Space-Bandwidth Product (SBP) [11]. SBP is a unit-less quantity that tells us the number of distinct points which can be measured over a given FOV.

The minimum spot diameter of a lens due to diffraction is $\delta_d \approx \lambda F/\#$, where λ is the wavelength of light. Since this quantity is independent of lens scale, the SBP for a diffraction-limited lens is

$$R_d = \frac{M^2 \Delta x \Delta y}{(\lambda F/\#)^2}. \quad (1)$$

The SBP increases quadratically with the scaling factor M (see the red curve in Figure 3). However, the SBP of a lens also depends on the diameter of the blur circle caused by geometric aberrations, δ_g , which increases linearly with the scaling factor M . As a result, when aberrations are considered, the SBP becomes [19]

$$R_g = \frac{M^2 \Delta x \Delta y}{(\lambda F/\#)^2 + M^2 \delta_g^2}. \quad (2)$$

Now the SBP plateaus at $\Delta x \Delta y / \delta_g^2$ when the lens is no longer diffraction limited and $M \delta_g \gg \lambda F/\#$ (see the green curve in Figure 3). Since the geometric blur size can be decreased by stopping down a lens, lens designers typically increase the F/# as a lens is scaled up. A general rule of thumb is that the F/# is increased such that the focal length in mm is approximately equal to $(F/\#)^3$. For instance, the 500mm focal length Schneider Apo-Symmar operates at F/8, and $8^3 \approx 500$. If this rule of thumb is used on a lens which primarily exhibits spherical aberration, it turns out that the geometric blur size δ_g becomes independent of the scaling factor M (see [19] for details). Then (see the blue curve in Figure 3) the SBP becomes

$$R_f = \frac{M^2 \Delta x \Delta y}{\lambda^2 M^{2/3} + \delta_g^2}. \quad (3)$$

Equation 3 is a scaling law that tells us generally how SBP increases with lens size for a conventional lens design.




Figure 3. A plot showing how Space-Bandwidth Product (SBP) increases as a function of lens size for a perfectly diffraction limited lens (R_d), a lens with geometric aberrations (R_g), and a lens whose F/# increases with lens size (R_f).

4. A Scaling Law for Computational Imaging




Figure 4. (a) A singlet lens with strong spherical aberrations. (b) The rayfan shows ray position on the sensor plane as a function of position in the lens aperture. The PSF has a strong peak because rays are concentrated around the center of the image plane.

In Lohmann analysis, it is assumed that the minimum resolvable spot size is equal to the blur size due to geometric aberrations, δ_g . However, the blur caused by a lens with spherical aberrations can be removed via deconvolution [27] [26] [28]. This is because a lens with spherical aberrations produces a Point Spread Function (PSF) that is well-conditioned and easily invertible. Figure 4(a) shows several rays from a distant point source traced through a singlet lens with spherical aberrations. The radial coordinate x of each ray on the sensor plane as a function of radial coordinate u in the aperture plane is given by the cubic curve $x = SA3 \cdot u^3$, where $SA3$ is the spherical aberration coefficient, as shown in Figure 4(b). The PSF is found by projecting this curve along the sensor axis, resulting in a sharply peaked PSF. The modulus of the Fourier Transform of this PSF, called the MTF, is relatively large at high frequencies and has no zero crossings. As a result, a captured image can be deblurred without introducing significant artifacts.

For a computational imaging system (i.e., with deblurring), the resolution is given by the pixel size ξ , and SBP does not depend directly on the geometric blur radius δ_g .




Figure 5. RMS deblurring error as a function of lens scale (M) for a lens with spherical aberrations. As M increases, both the PSF size and the deblurring error increase. While the size of the PSF increases linearly with M , deblurring error increases with $M^{1/3.8}$.




Figure 6. A new scaling law for computational imaging (R_c). Note that R_c not only improves upon the aberration limited curve R_g , it also improves upon the conventional lens design curve R_f without requiring $F/\#$ to increase with M .

A more pertinent quantity for measuring image quality is SNR. In the absence of any noise we can theoretically increase SBP by decreasing pixel size until we have reached the diffraction limit. In order to provide a fair comparison between any two computational imaging systems, we must fix the SNR. Therefore, we fix SNR and then derive a new scaling law for computational imaging.

First, let us consider the effect of deblurring on SNR. The SNR of an image produced by any computational imaging system is simply $1/\sigma_D$, where σ_D is the deblurring error. Figure 5 shows how σ_D scales as a function of lens size. To calculate the deblurring error we follow the procedure used by Zhou and Nayar [32]. A captured image is simulated by blurring a ground truth image with the PSF at the given scale and adding white gaussian noise with standard deviation σ_N . The blurred image is then deblurred using Wiener deconvolution, and σ_D is calculated as the RMS error between the deblurred and ground truth images, averaged over a set of natural images. The deblurring error curves are shown for a variety of σ_N , and the best fit polynomial is empirically found to be in the range from $\sigma_D \propto \sigma_N M^{1/3.6}$ to $\sigma_D \propto \sigma_N M^{1/4}$, which we approximate as:

$$\sigma_D = \sigma_N M^{1/3.8}. \quad (4)$$

Next, we relate pixel size ξ to sensor noise σ_N . Scaling ξ by a factor of M increases the pixel's area by a factor of M^2 . For a fully saturated pixel, assuming a shot noise limited sensor, this will increase the sensor's full well capacity by M^2 and decrease noise by a factor of M^{-1} relative to the signal. The sensor noise is then related to pixel size as:

$$\sigma_N = \frac{k}{\xi}, \quad (5)$$

where k is a constant. If we fix the SNR of our computational imaging system (by fixing σ_D), then, using equations 4 and 5 and defining a new constant $k_1 = k/\sigma_D$, we can relate pixel size and scale as:

$$\xi = k_1 M^{1/3.8}. \quad (6)$$

Finally, assuming that our pixels size is greater than the diffraction limited spot size, the number of resolvable points is just the sensor area divided by the pixel area:

$$R_c = \frac{M^2 \Delta x \Delta y}{\xi^2} = \frac{M^2 \Delta x \Delta y}{k_1^2 M^{2/3.8}}. \quad (7)$$

Equation 7 is a scaling law for computational imaging that says resolution increases faster with M than in the case of Lohmann’s scaling law for conventional lens design, *without* requiring $F/\#$ to increase with M (see Figure 6). This tells us that, for a computational imaging system with a fixed SNR (i.e. fixed deblurring error), SBP scales more quickly with lens size than it does for conventional lens designs.

5. Proposed Architecture

According to Equation 7, a computational imaging approach can enable a greater resolution to be achieved with a smaller camera size. To demonstrate this principle, we show results from two proof of concept systems that utilize a very simple optical element. By using a large ball lens, an array of planar sensors, and deconvolution as a post processing step, we are able to capture gigapixel images with a very compact camera.

The key to our architecture lies in the size of the sensors relative to the ball lens. Together, a ball lens and spherical image plane produce a camera with perfect radial symmetry. We approximate a spherical image plane with a tessellated regular polyhedron, such as an icosahedron. A planar sensor is placed on each surface of the polyhedron. Note that because sensors are typically rectangular, a different polyhedron, such as a truncated icosahedron, may provide more optimal sensor packing. Relatively small sensors are used so that each sensor occupies a small FOV and the image plane closely approximates the spherical surface. As a result, our camera produces a PSF that is not completely spatially invariant, but comes within a close approximation.

Single Element Monocentric Gigapixel Camera. The first system we demonstrate consists solely of a ball lens and an array of planar sensors. We use a 100mm acrylic ball lens and a 5 megapixel 1/2.5” Lu575 sensor from Lumenera [20] (see Figure 7(a)). We emulate an image captured by multiple sensors by sequentially scanning the image plane using a pan/tilt motor. With this camera, a 1 gigapixel image can be generated over a roughly $60^\circ \times 40^\circ$ FOV by tiling 14×14 sensors onto a $75 \text{mm} \times 50 \text{mm}$ image surface. When acquiring images with the pan/tilt unit, we allow a small overlap between adjacent images.

The PSF as a function of field position on each individual sensor is shown in Figure 7(b). Note that the PSF shape remains fairly consistent across the FOV of each sensor. The MTF (shown in in Figure 7(c)) avoids zero crossings up to

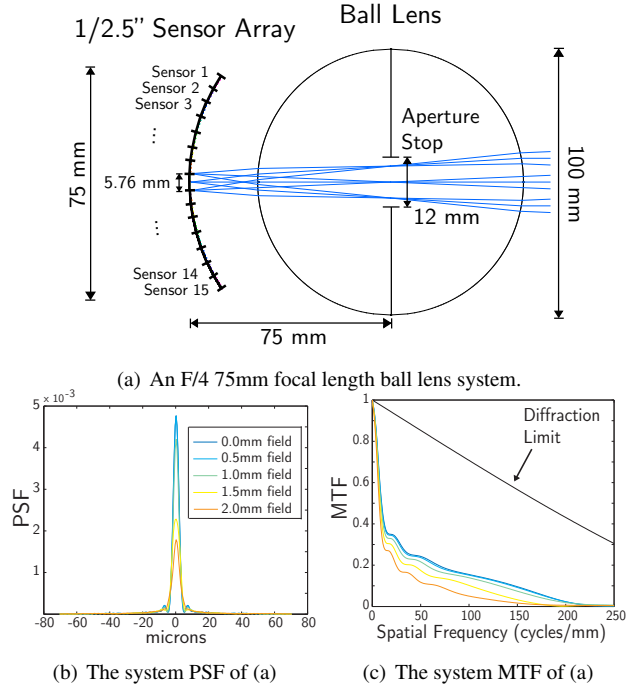


Figure 7. (a) Our single element gigapixel camera, which consists solely of a ball lens with an aperture stop surrounded by an array of planar sensors. (b) Because each sensor occupies a small FOV, the PSF is nearly invariant to field position on the sensor. (c) The PSF is easily invertible because the MTF avoids zero crossings and preserves high frequencies.

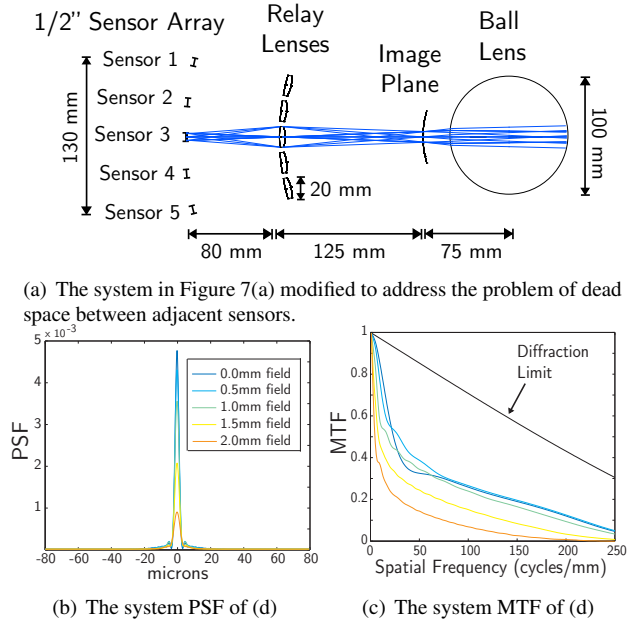


Figure 8. (a) A multiscale design based on our proposed architecture. An array of relay lenses modifies the system magnification so that the FOV of adjacent sensors overlaps slightly. The PSF (b) and MTF (c) are similar to the system in Figure 7(a).




Figure 9. A system used to verify the performance of the design shown in Figure 7(a). An aperture is placed on the surface of the ball lens. A gigapixel image is captured by sequentially translating a single 1/2.5", 5 megapixel sensor with a pan/tilt motor. A final implementation would require a large array of sensors with no dead space in between them.

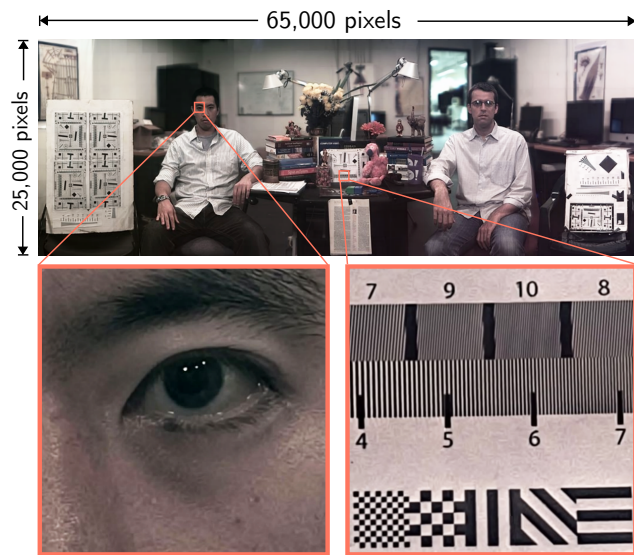


Figure 10. A **1.6 gigapixel image** captured using the implementation shown in Figure 9. The image dimensions are 65,000 x 25,000 pixels, and the scene occupies a $104^\circ \times 40^\circ$ FOV. The left inset reveals individual hairs and pores on the person on the right. The right inset reveals the extremely fine features of a resolution chart. Please see the Supplementary Material to view this example in more detail.

the Nyquist frequency of the sensor. The plots were generated using Zemax Optical Design Software [31].

An implementation of this design is shown in Figure 9. Figures 2 and 10 show two gigapixel images captured with this system. Note the remarkable level of detail captured in each of the photographs. Zooming in to Figure 2 reveals the label of a resistor on a PCB board, the stippling print pattern on a dollar bill, a miniature 2D barcode pattern, and the extremely fine ridges of a fingerprint. Closeups in Figure 10 reveal pores on the skin of one of the persons photographed.




Figure 11. Implementation of the design of Figure 8(a) capable of capturing a 15 megapixel region of a gigapixel image. A full gigapixel camera requires 66x as many sensors and relay lenses.




Figure 12. A 14 megapixel image captured using the video ready implementation shown in Figure 11. The image dimensions are 10,000 x 14,000 pixels, and the scene occupies a $15^\circ \times 2^\circ$ FOV. The insets show a closeup of a person's face (left) and a resolution target (right).

Multiscale Monocentric Gigapixel Camera. Our first camera system is extremely compact, but it assumes there is no dead space between adjacent sensors. Sensors require at least some packaging around the active pixel area, which renders this solution impractical. A solution to overcome this problem is to introduce a secondary optic for each sensor, changing the system magnification so that the FOV of adjacent sensors overlaps slightly. This approach is taken in the implementation shown in Figure 8(a).

A photo of the video ready implementation, currently under development, is shown in Figure 11. The system consists of custom optics manufactured by LEI Optics [18], and 5 NET-USA Foculus Sensors [23]. A 14 megapixel image captured with this system is shown in Figure 12.

Color. Because our cameras do not include any color correcting elements, they suffer from axial chromatic aberrations. For the 100mm diameter ball lens that we use, the chromatic focus shift is about 1.5mm over the visible wavelength range. However, most of the image blur caused by chromatic focus shift is in the chrominance channel of captured images [12] [3]. Since humans are less sensitive to blur in chrominance channels, axial chromatic aberrations do not cause a significant degradation in perceived image quality. We use the deblurring technique from Cossairt and Nayar [3], which is inexact but produces images that look good.

Post Processing. The post processing for captured images follows several steps. First, a transformation from RGB to YUV color space is applied to each captured image. Next, Wiener deconvolution is applied to the luminance channel only, and the image is transformed back to RGB color space. A noise reduction algorithm is then applied to the image to suppress deblurring artifacts. We found the BM3D algorithm [4] to produce the best results. Finally, the set of captured images are stitched to obtain a high resolution image using the Microsoft Image Composite Editor [13].

6. Discussion

A Single Element Design. The design in Figure 7(a) is extremely compact, but impractical because adjacent sensors must be packed without any dead space in between them. The design in Figure 8(a) is practical enough for an implementation using off-the-shelf components (see Figure 11), but is much less compact. The size of this system is limited by the package size of the sensor relative to the active sensor area. Sensors with a package size that is only 1.5x larger than the active sensor area are currently commercially available. With these sensors, it is possible to build a gigapixel camera that uses only a single optical element,




Figure 13. A single element design for a gigapixel camera. The design is a hybrid between the two implementations introduced in Section 5. Each sensor is coupled with a lens that decreases focal distance, allowing FOV to overlap between adjacent sensors.

as shown in Figure 13. In this design, each sensor is coupled with a smaller acrylic relay lens that decreases the focal length of the larger acrylic ball lens. The relay lenses share a surface with the ball lens, which means that it is possible to combine the entire optical system into a single element that may be manufactured by molding a single material, drastically simplifying the complexity (and hence alignment) of the system.

Capturing the Complete Sphere. All the designs proposed in this paper use a ball lens. A great advantage of using a ball lens is that, because it has perfect radial symmetry, a near hemispherical FOV can be captured. In fact, it can even be used to capture the complete sphere, as shown in Figure 14. This design is similar to the one in Figure 8(a) with a large gap between adjacent lens/sensor pairs. Light passes through the gaps on one hemisphere, forming an image on a sensor located on the opposite hemisphere. As a result, the sensors cover the complete 2π FOV at the cost of losing roughly half the incident light.

Limitations of the Scaling Law. Lohmann’s scaling law tells us how the resolution of a conventional lens design will scale with lens size, but the law is based on a generalization, and there are many exceptions to this rule. According to Equation 7, with the aid of computations, the resolution of a lens with spherical aberrations will, in general, scale more quickly than for a conventional lens design. However, there is no free lunch. A lens which requires deblurring will have a smaller SNR than a diffraction limited lens of the same scale. Implicit in our new scaling law is a trade-off between SNR and complexity. In our proposed architecture, we have chosen a design that favors simplicity, and as a consequence, also results in a lower SNR.

Balancing Complexity and Computations. Any computational imaging system poses an inherent trade-off between complexity and SNR. In the context of developing a compact high resolution camera, we are forced to choose




Figure 14. A design for a gigapixel camera with a 2π radian FOV. The design is similar to the implementation in Figure 8(a) with a large gap between adjacent lens/sensor pairs. Light passes through the gaps on one hemisphere, forming an image on a sensor located on the opposite hemisphere.

a small pixel size. Small pixels make the above trade-off more pertinent. In practice, exploring this trade-off requires a carefully designed measure for complexity.

A good complexity measure must take into account many different factors: the number of surfaces, the degree polynomial of each surface, etc. While we have not given a complete treatment to lens complexity in this paper, the scaling law of Equation 7 does give new insight into the potential benefit of computational cameras. Given an appropriate measure for lens complexity, it may be possible to design a lens that achieves an *optimal* balance between complexity and computations. Rather than favoring simplicity, an optimal design may consist of more elements than the designs discussed previously in this paper. One can imagine, for instance, a lens of moderate complexity, such as a Schneider Apo-Symmar, that is optimized for deblurring quality. By finding an optimal trade-off between complexity and computations, such a lens might be capable of imaging a gigapixel with a much smaller package size while still maintaining a satisfactory SNR.

7. Acknowledgements

This research was supported in part by DARPA Award No. W911NF-10-1-0214. Oliver Cossairt was supported by an NSF Graduate Research Fellowship. The authors thank Ravi Athale of Mitre Corporation for his comments, Keith Yeager of Columbia University for his help with machining, Lauren Kushner for her 3D modeling expertise, and Tai-Hsu Lin for programming assistance.

References

- [1] M. Ben-Ezra. High Resolution Large Format Tile-Scan - Camera Design, Calibration, and Extended Depth of Field. In *ICCP*, Mar 2010.
- [2] D. J. Brady and N. Hagen. Multiscale lens design. *Opt. Express*, 17(13):10659–10674, 2009.
- [3] O. Cossairt and S. Nayar. Spectral focal sweep: Extended depth of field from chromatic aberrations. In *ICCP*, Mar 2010.
- [4] K. Dabov, A. Foi, V. Katkovich, and K. Egiazarian. Image denoising with block-matching and 3 D filtering. In *Proc. SPIE*, volume 6064, pages 354–365. Citeseer, 2006.
- [5] Darpa at 50. www.darpa.mil/Docs/1-25013846.Eprint.200811141152151.pdf, 2010.
- [6] R. Dinyari, S. Rim, K. Huang, P. Catrysse, and P. Peumans. Curving monolithic silicon for nonplanar focal plane array applications. *Applied Physics Letters*, 92:091114, 2008.
- [7] E. Dowski and J. Cathey. Extended depth of field through wave-front coding. *Appl. Opt.*, 34:1859–1866, 1995.
- [8] E. Dowski Jr, R. Cormack, and S. Sarama. Wavefront coding: jointly optimized optical and digital imaging systems. In *Proc. SPIE*, volume 4041, pages 114–120. Citeseer, 2000.
- [9] K. Fife, A. El Gamal, and H. Wong. A 3MPixel Multi-Aperture Image Sensor with 0.7 um Pixels in 0.11 um CMOS. In *IEEE ISSCC Conference*, 2008.
- [10] The Gigapixel Project. <http://www.gigapixel.org/>, 2007.
- [11] J. Goodman. *Introduction to Fourier optics*. Roberts & Company Publishers, 2005.
- [12] F. Guichard, H.-P. Nguyen, R. Tessières, M. Pyanet, I. Tarchouna, and F. Cao. Extended depth-of-field using sharpness transport across color channels. volume 7250. *Proc. SPIE*, 2009.
- [13] Microsoft Image Composite Editor (ICE) website. <http://research.microsoft.com/en-us/um/redmond/groups/ivm/ICE/>, 2010.
- [14] R. Kingslake. *A history of the photographic lens*. Academic Press, 1989.
- [15] H. Ko, M. Stoykovich, J. Song, V. Malyarchuk, W. Choi, C. Yu, J. Geddes Iii, J. Xiao, S. Wang, Y. Huang, et al. A hemispherical electronic eye camera based on compressible silicon optoelectronics. *Nature*, 454(7205):748–753, 2008.
- [16] G. Krishnan and S. Nayar. Towards A True Spherical Camera. In *SPIE Human Vision and Electronic Imaging*, Jan 2009.
- [17] L. Lee and R. Szema. Inspirations from biological optics for advanced photonic systems. *Science*, 310(5751):1148, 2005.
- [18] LEI Optics. <http://www.leioptics.com/>, 2010.
- [19] A. W. Lohmann. Scaling laws for lens systems. *Appl. Opt.*, 28(23):4996–4998, 1989.
- [20] Lumenera Corporation. <http://www.lumenera.com/>, 2010.
- [21] R. Luneburg. *Mathematical theory of optics*. University of California Press, 1964.
- [22] D. Marks and D. Brady. Gigagon: A Monocentric Lens Design Imaging 40 Gigapixels. In *Imaging Systems*. OSA, 2010.
- [23] NET GmbH. <http://www.net-gmbh.com/>, 2010.
- [24] Y. Nomura, L. Zhang, and S. Nayar. Scene collages and flexible camera arrays. In *Proc. EGSR*, 2007.
- [25] S. Rim, P. Catrysse, R. Dinyari, K. Huang, and P. Peumans. The optical advantages of curved focal plane arrays. In *Proc. SPIE*, volume 5678, pages 48–58, 2005.
- [26] M. Robinson and V. Bhakta. Experimental validation of extended depth-of-field imaging via spherical coding. In *Computational Optical Sensing and Imaging*. OSA, 2009.
- [27] M. Robinson, G. Feng, and D. Stork. Spherical coded imagers. In *Proc. SPIE*, volume 7429, page 20, 2009.
- [28] M. Robinson and D. G. Stork. Extending depth-of-field: Spherical coding versus asymmetric wavefront coding. In *Computational Optical Sensing and Imaging*. OSA, 2009.
- [29] S. Wang and W. Heidrich. The design of an inexpensive very high resolution scan camera system. In *Computer Graphics Forum*, volume 23, pages 441–450. Citeseer, 2004.
- [30] B. Wilburn, N. Joshi, V. Vaish, E. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy. High performance imaging using large camera arrays. *ACM Transactions on Graphics (TOG)*, 24(3):776, 2005.
- [31] Zemax Optical Design Software. <http://www.zemax.com/>, 2010.
- [32] C. Zhou and S. Nayar. What are good apertures for defocus deblurring? In *ICCP*, pages 1–8. IEEE, 2010.