EUROGRAPHICS 2004 / M.-P. Cani and M. Slater
(Guest Editors)

Volume 23 (2004), Number 3

The Design of an Inexpensive Very High Resolution Scan Camera System

Shuzhen Wang and Wolfgang Heidrich

The University of British Columbia

Abstract

We describe a system for transforming an off-the-shelf flatbed scanner into a $200 scan backend for large format cameras. While we describe both software and hardware aspects, the focus of the paper is on software issues such as color calibration and removal of scanner artifacts. With current scanner technology, the resulting camera system is capable of taking black&white, color, or near-infrared photographs with up to 490 million pixels. Our analysis shows that we achieve actual optical resolutions close to the theoretical maximum, and that color reproduction is comparable to commercial camera systems. We believe that the camera system described here has many potential applications in image-based modeling and rendering, cultural heritage projects, and professional digital photography.

1. Motivation

Although the resolution of commercial digital camera systems has steadily improved in recent years, there are still a number of applications in computer graphics, computer vision, and photography, for which the current resolutions are not detailed enough. For example, during the digitization of cultural objects and artwork, millimeter-scale features are often considered important even for objects of large overall size (see, for example, [LPC*00]). Another example is any application where high-quality large format prints (letter size and above) are necessary.

Currently, the kind of resolution necessary for these applications can only be achieved through analog large format photography (with the possibility to subsequently scan the resulting photograph with a flatbed scanner), or through digital scan backends for middle and large format cameras. Unfortunately, analog photography lacks the immediacy and convenience of digital photography, while digital scan backends cost several tens of thousands of dollars and are therefore not readily available to most researchers and photographers.

In this paper, we present a scan camera system in which we transform a $100–$200 consumer-grade flatbed scanner into a very high resolution digital scan backend. When used in combination with a conventional $8'' \times 10''$ large format view camera, this setup achieves resolutions of up to 122-490 million pixels (depending on scanner model). The total system, including large format camera, lens, scanner, and color filters can readily be assembled for about $1200-$1300. An example of the resolution and image quality we can achieve with this setup is shown in Figure 1. The top image represents a downsampled full view image acquired with our camera, while the bottom image shows a detailed portion printed at full resolution.

As an additional benefit, we can use the same camera with an infrared (IR) filter for photography in the near-infrared (NIR) spectrum, which has proven useful for applications such as object recognition and digital matting in cinematography [Fie65, DWT*02].

In the remainder of this paper, we first briefly review the relevant related work. We then describe the optical setup of our system (Section 4), as well as image cleanup (Section 5), NIR imaging (Section 6), and color calibration (Section 7). We describe typical usage patterns for the camera in Section 8 and conclude with some results, including quantitative resolution measurements and a comparison of our camera with a semi-professional digital SLR camera, the Canon EOS D60 (Section 9). These comparisons include both resolution and color reproduction.

Figure 1: *Example image from our camera (top) and detail view for a part at the original resolution (for a 300 DPI print).*

2. Related Work

The very highest resolution commercial cameras available today are digital backends for large format view cameras (typically $4'' \times 5''$ cameras). This has two major reasons: firstly, the image sensors (CCD or C-MOS) that produce these very high resolutions do not fit into smaller devices such as SLR bodies, and secondly, only lenses for medium and large format cameras are of sufficiently high quality to support these resolutions.

The backends exist in two basic flavors: *single shot* backends have a 2-dimensional image sensor, similar to the one used in consumer cameras and digital SLRs. This technology can currently achieve image resolutions of up to 22 million pixels (for example with the Sinarback 54 [Sin02]).

Even higher resolutions can be achieved with *scan backends*. In this technology, a single row of sensor elements (*line sensor*) is swept across the image plane similar to a familiar flatbed scanner. For example, the Betterlight Super 8K [Bet04] scan camera is a backend for $4'' \times 5''$ large format cameras, that can achieve an optical resolution of 8000×10660 pixels (12000×15990 interpolated). The down side of this approach is that the technology cannot be used for animated objects.

Unfortunately, both technologies are quite expensive (the backends run in the range of several tens of thousands of dollars), which puts digital backends out of reach for most photographers and researchers. In this paper we describe how to turn an inexpensive \$100-\$200 flatbed scanner into a digital backend for $8'' \times 10''$ view cameras. Using only off-the-shelf components, we achieve resolutions surpassing those of the commercial scan backends that are two orders of magnitude more expensive.

Alternative approaches for obtaining very high resolution digital photographs include the scanning of large format analog prints, as well as stitching smaller photographs together into a large one. Scanning of analog prints is currently used by many professional photographers, but it lacks the speed and ease of use of a completely digital process. Panorama stitching (e.g. Szeliski et al. [SS97]) works best for large, distant objects. Close objects often require adjustment of focus and/or viewpoint between images, both of which usually result in visible seams in the stitched image.

A research project related to our work is the light field camera built by Yang et al. [Yan00, YLIM00]. Their system is based on a flatbed scanner with a lens array attached to the glass surface. Their goal is to acquire a 2D array of relatively low resolution images that form a light field. Since their optical setup is quite different from ours, the calibration issues differ significantly.

2.1. Color Imaging

There are several technologies for color imaging with digital photo cameras. Most single shot cameras use a Bayer pattern [Bay], i.e. a grid of differently colored filters in front of the individual sensor elements. Note that this means that only one color channel is actually measured per pixel, while the other channels are interpolated. Scan cameras often have three rows of pixels that obtain true measurements of all color channels in every pixel as the sensor sweeps across the image plane.

Alternatively, one can use a sensor that averages all wavelengths, and combine three images taken with different color filters (usually arranged in a filter wheel) into a single color image. If single shot cameras are used with filter wheels, then they are no longer able to capture moving objects. Other than that, this approach is suitable for both single shot and scan backends. While filter wheels add to the complexity of the acquisition process, they also allow for additional flexibility by using non-standard filters that select specific ranges in the light spectrum. Since our work is based on scanner technology that only has a single sensor row (see Section 4 for details), we make use of a color wheel for our system.

2.2. Image Cleanup

There has been a lot of work on removing sensor artifacts from photographs. The best methods available typically make use of domain-specific knowledge to detect the location of artifacts and repair them. For example, film scanners often scan slides at an infrared wavelength in addition to the visible wavelengths. Since normal slide material is transparent to NIR radiation everywhere, the dark regions in the IR scan present a mask for dust and scratches and other instance where the color channels are invalid. The masked regions can then be replaced with methods such as image inpainting [BSCB00], or other image editing or inpainting methods [DCOY03, PGB03].

Other potential artifacts originate from slight differences in the characteristics of individual sensors in a sensor array. The most important sources of variations are *dark current* and *flat field response* [Jan01]. Dark current results in a non-zero measurement even for pixels that have not been hit by any light at all. It effectively adds a positive constant to all measurements for one particular sensor element. Apart from dark current, the sensor response is very close to proportional to the number of incoming photons. However, the constant of this linear mapping can again vary from element to element in the sensor array. To correct for both effects, one can take one dark (i.e. closed aperture) image and one "white" image (i.e. an image of a white piece of paper, out of focus, and exposed to just below saturation). These two images then define a per-pixel linear mapping that compensates for the differences of the sensor elements. This process is frequently used in astronomy, but also for the calibration of scanners. In practice, both the black and the white image is usually an average of multiple such photographs to suppress random noise.

3. System Overview

At the core of our camera system, we use a scanner with Canon's "LED Indirect Exposure" technology (LIDE [Can04]). As described in the next section, this technology is central to the functioning of our camera system. It is available in a series of products, including the Canon LIDE 30 and LIDE 80 scanners, with 1200 DPI and 2400 DPI resolution, respectively. These are the two scanner models that we have experimented with.

We modify the scanner to work as a backend for a conventional view camera. To keep the cost of the overall system down, we use a view camera kit from Bender photographic [Ben04], and a relatively inexpensive Nikkor-M 300mm f/9 large format lens. It should be noted that any other combination of large format camera and lens would work as well. The complete part list is shown in Table 1.

The dimensions of the view camera and the flatbed scanner result in a $8.5'' \times 10''$ exposed "film" area, which corresponds to approximately 122 million pixels at a scan resolution of 1200 DPI (LIDE 30), and approximately 490 million pixels at 2400 DPI (LIDE 80). Figure 2 shows a photograph of the overall camera system. We would like to point out that,

Description	Part	Cost
Scanner	Canon LIDE 80	$200
Scanner (alternative)	Canon LIDE 30	$100
View camera	Bender $8'' \times 10''$ kit	$400
Large format lens	Nikkor-M 300mm f/9	$600
UV/IR cutoff filter	Schneider Optics B+W #486	$40
Red filter	Lee filter 25	$15
Green filter	Lee filter 58	$15
Blue filter	Lee filter 47B	$15
IR filter	Lee filter 87	$15

Table 1: *Part list for our camera system with two different scanner options.*

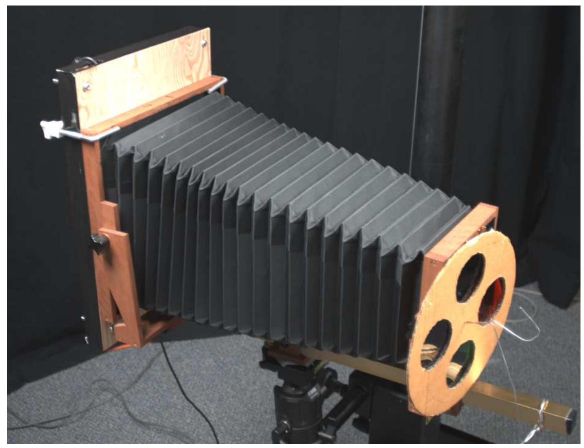

Figure 2: *A photo of our camera system.*

while we have verified that our system works in principle with the LIDE 80 scanner, memory and software restrictions have so far prevented us from acquiring a full resolution image with that scanner – the raw data for a 490 million pixel RGB image with 16 bits per channel is almost 3GB. The full-size photographs used for comparison in Section 9 are therefore all 122 million pixel images taken with the LIDE 30 scanner.

4. Optical Setup

As mentioned above, the central part of our camera is a scanner with Canon's LIDE technology [Can04]. With this technology, the sensor covers the complete width of the paper (about $8.5''$), and focusing is achieved through an array of lenses directly on top of the sensor (rather than a combination of mirrors and a single lens). The conceptual design is depicted in Figure 3. The sensor in a LIDE scanner consists of a single row of sensor elements that are sensitive to all visible frequencies, as well as near-infrared. In normal operation, the scanner sequentially illuminates the paper with a red, green, and blue LED for every scanline to obtain the different color channels.

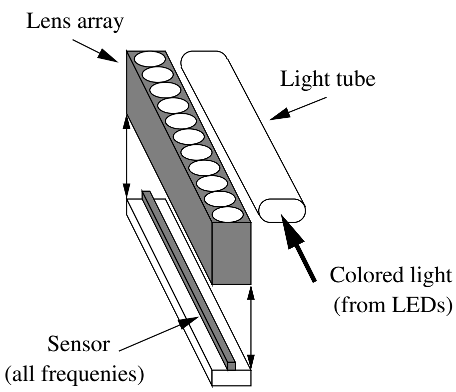

Figure 3: *Design of a scanner with the LIDE (LED Indirect Exposure) technology.*

For our purposes, we remove the filter array in front of the sensor, because we want to use the lens of the large format camera for creating a focused image instead. We also disable the LEDs and remove the light guide, so that only light from the lens of the view camera arrives at the sensor.

For color photography, we have to acquire multiple scans with different color filters. To this end, we built a manually operated filter wheel from cardboard and photographic polyester filters. Commercially available, computer controlled filter wheels would work as well, but would increase the cost of the system. Also important for color photography is an UV/IR cutoff filter that suppresses the near-infrared light that would otherwise make color calibration impossible. Alternatively, a polyester infrared filter can be used for NIR photography.

Note that scanners based on different technology would require much more dramatic changes. This is because typical flatbed scanners use a catadioptric optical system (i.e. an optical system based on lenses and mirrors) to map the full paper width onto a much more narrow sensor. The field-of-view for these catadioptric systems is very limited, so that they would have to be completely redesigned for our purposes, and possibly the mechanical setup would have to be changed as well.

5. Black&White Imaging and Calibration

After attaching the modified scanner to the view camera in place of the usual film holder, images can be acquired with the help of any scanner software. In our experiments we typically use the shareware program VueScan [Ham02] instead of vendor-provided scan software, since it gives us access to the uncalibrated 16 bit raw sensor data. This helps us implementing our own calibration steps, which is necessary due to the changes we performed on the scanner optics.

The raw sensor data suffers from a number of artifacts. Firstly, there are the usual differences in the sensor response (dark current and flat fielding). These effects are independent of the position of the scan head, and therefore constant across a row in the image. In addition, however, there are also artifacts that depend on the position of the scan head. For example, there can be small scratches or dust on the protective coating that covers the sensor. In normal scanner operation, the scanner optics does not move with respect to the image sensor, so that light always hits the sensor from the same direction, and the effect of scratches and dust is again constant across every image row. In our modified setup, however, the sensor moves relative to optics (i.e. the camera lens), which means that light arrives at the sensor from different directions depending on the position of the scan head. This means that a dust particle, for example, can now block light only for certain image columns but not for others. Figure 4 shows an (enhanced) scan of an homogeneous surface to illustrate these effects.

Figure 4: *Enhanced raw scan of a white surface illustrating both position-dependent and position-independent sensor artifacts.*

To compensate for these deficiencies, we take a two-stage approach. First, we define a per-pixel linear mapping using a black image and a white image, as is usually done for dark current and flat field correction. Since we define a linear mapping for every pixel, and not just for every row, we also compensate for most position-dependent artifacts in this way. However, there are usually a few dozen or so of rows (out of about 20,000 for a 2400 DPI scanner) that are not adequately corrected by this process. This includes sensor elements that are severely scratched or just plain broken, as well as sensors that have suffered from additional dust particles since the calibration image was taken. In a second step, we therefore detect these artifacts, and use image inpainting [BSCB00] to repair them.

In order to compute a mask for the individual artifacts, we

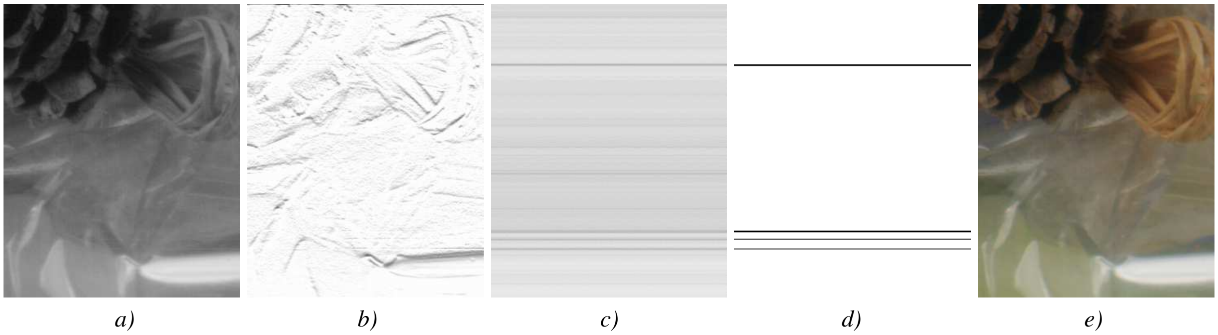

Figure 5: *Process of the image repair algorithm. a) portion of a raw image. Note the horizontal artifacts in the bottom third of the image. b) Sobel edge detector applied to the same image. c) a 1D Hough transform picks up the horizontal edges, in particular the three artifacts at the bottom of the image and one near the top. d) the thresholded Hough transform image serves as a mask for the image inpainting algorithm. e) repaired image, now with all color channels.*

tile the photo into smaller parts (typically about 1000×1000 pixels), and run a Sobel edge detection operator. The resulting edge image is passed through a one-dimensional Hough transform that detects the strength of horizontal edges only. Thresholding of that image finally yields a mask for the actual artifacts (see Figure 5). As mentioned before, we use image inpainting [BSCB00] to replace the color values of these masked regions.

The final step in the black&white calibration is to gamma-correct the image for specific displays. In this step, we also reduce the 16bit raw data to 8bits.

6. NIR Imaging

Due to the infrared sensitivity of our scanner, we can use the camera system for NIR photography, simply by using the appropriate filter to block out all visible light. The calibration issues for NIR images are the same as for black&white images, as described in the previous section. Figure 6 shows a comparison of an infrared and a black&white image of a painted ceramic figurine. Note that the different color paint used for the figurine is equally reflective in the infrared spectrum, so that pigmentation effects are mostly removed.

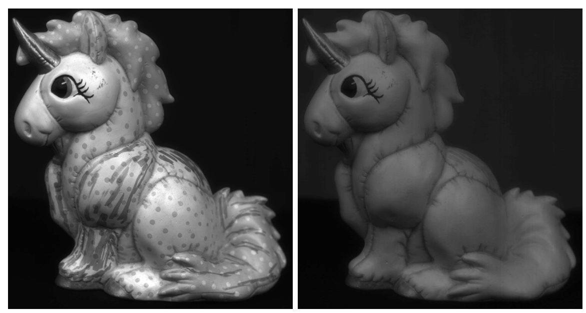

Figure 6: *Comparison of a black&white photograph (left) with a near infrared photograph (right). Note that except for the black color of the eye, the effect of different paints is mostly removed in the NIR image.*

7. Color

For color imaging we have to acquire the different color channels in individual scans using the different color filters of our filter wheel. Since these filters do *not* by themselves block infrared light, it is also mandatory to use the UV/IR blocking filter at the same time. We were unable to determine whether the scanner is also sensitive to UV light, but since our large format camera lens is made of glass, UV radiation should not be an issue even without the UV/IR cutoff filter.

After the scans for the individual color planes have passed through the black&white calibration process described in Section 5 (except for the final gamma correction step), they need to be aligned to each other. We find the scans to be aligned very well in the columns (indicating that the sensor elements repeatably move along the same line in every scan), but the scanlines can be off by about 20-30 pixels, since the initial starting position can vary slightly from scan to scan. To compensate for this, we manually select an offset by examining a small region of the image that has sufficiently rich detail. We have also experimented with automated alignment methods, such as a 1D search on a region of the image after edge detection. However, we find the manual alignment to be by the far easiest and fastest method in practice.

A simple merging of the color channels in this way produces images with a strong red shift (see top left of Figure 7), since our red color filter is overall more transparent than the blue and green filters. Consequently, some form of color calibration is always called for. The easiest form of color calibration is a simple white balance (i.e. relative scaling of the color channels), followed by gamma correction like in the black&white case. As shown on the top left of Figure 7, this simple method already produces quite convincing results.

A more reliable color reproduction can be achieved with so-called ICC (International Color Consortium) profiles, which describe the characteristics of both input and output devices [Wal00]. A profile can be used to convert an image from the color space of the input device into a common profile connection space (PCS) which is either CIE $L^*a^*b^*$ or CIE XYZ with a defined white point (illuminant D50). An output profile for the printer or monitor (which includes the appropriate gamma correction) converts image data from the PCS to the color space of the output device before it is displayed or printed. Profiles for many input and output devices are available from the individual manufacturers, but they can also be created from images of color targets with both professional and public domain color management software.

We use the public domain program SCARSE [Fro02] together with a Kodak Q60 color target to generate an ICC profile for our camera under a given illumination condition. The Q60 target (depicted in Figure 7) is based on the ANSI standard IT8.7/2 [ANS03]. It has more color fields than some other color targets (such as the Gretag McBeth color checker), providing a denser sampling of the color gamut for an input device. The precise color of every color field can vary slightly due to tolerances in the production process, but the precise calibration data for every batch is available from the Kodak FTP server (ftp://FTP.Kodak.com/GASTDS/Q60DATA). The bottom left of Figure 7 shows the color calibrated version of the target.

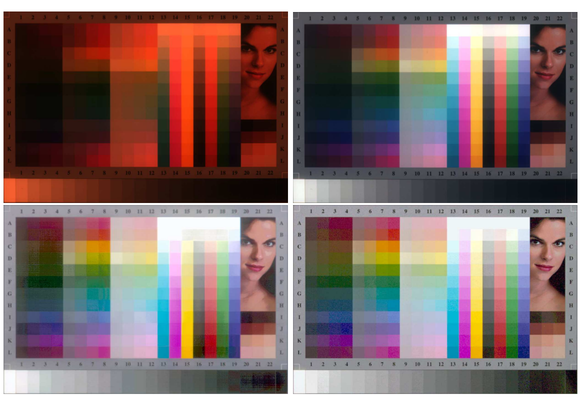

Figure 7: *A picture of the color calibration target photographed with our scan camera. Top left: no calibration. Top right: with simple white balance. Bottom left: calibrated with ICC profile. Bottom right: result for the Canon EOS D60 after ICC profile calibration.*

8. Workflow

Since the usage patterns of scan cameras are different from regular one-shot digital cameras, we give a brief overview of the workflow in the following. Note that the points made in this section apply to all scan cameras and are not specific to ours.

The typical imaging process starts with scene composition and focusing. With analog view cameras this is done by replacing the film with a frosted glass plate and using a focusing cloth. With a digital backend like ours, we have the option of using the same method, since the backend easily detaches from the view camera and can be replaced with focusing aids. One slight complication is that the commercially available focusing aids do not line up with the focus plane of our digital backend. This could be solved rather easily by building a special focusing aid with the appropriate distance. However, we find it much more convenient to use low-resolution scans directly for scene composition, and high-resolution scans of selected image regions for supporting the focusing process. This allows us to keep the digital backend attached at all times, but more importantly we feel we can determine the exact location of the focus plane much more precisely in this way. The individual scans for focusing and image composition take only about 10-15 seconds each, while a full-resolution scan with USB 2.0 takes about 3 minutes.

After scene composition and focusing, we need to adjust the exposure. Just as in normal photography, we can modify the lens aperture to change both exposure and depth-of-field at the same time. In single-shot photography we also have control over the exposure time, which allows us to change exposure *without* affecting depth-of-field. By construction, scan cameras do not have this possibility. However, we can control exposure by adding neutral density (gray) filters in front of the lens so that we retain control over the depth-of-field.

Another difference between photography with scan cameras and other cameras is lighting. Flashes or flickering light sources such as fluorescent lights cannot be used due to the extended capture period. Fortunately, there is a wide selection of stable light sources for video or film that can be applied without problems. We ourselves employ a pair of HMI (metal halide) lamps that have a bluish color similar to daylight. Color calibration should be performed with the same light sources used for taking the actual photographs.

9. Results

In order to obtain quantitative results with our camera, we compared it against the Canon EOS D60 together with a Canon EF 28-105mm zoom lens. That camera is a 6 million pixel digital SLR camera for the high-end consumer and semi-professional market. Although the technology behind both cameras is quite different, the comparison gives some reference for the quality of color reproduction and the resolution we achieve. Also note that the D60 alone is almost twice as expensive as the complete camera/lens system described here. As mentioned above, we have so far been unable to acquire full resolution scans with the LIDE 80 scanner due to memory constraints. For that reason, all comparisons in this section are performed with the LIDE 30 scanner at a resolution of 122 million pixels.

As a first step, we tested the light sensitivity of our scan camera. To this end, we set both our system and the D60 to an aperture of f/9 (the largest aperture of our large-format lens). We then set the film sensitivity on the D60 to 100 ASA, and adjusted the exposure time until we got an image of a brightness comparable to the one obtained with our scan camera. This way we determined that our camera has roughly the sensitivity of a 100 ASA film operated at an exposure time of 1/60 s. Like many digital cameras we can simulate a higher sensitivity at the expense of increased noise levels by linearly scaling up the raw data (i.e. turning up the gain value).

To compare the color reproduction of our camera, we acquired an ICC profile of the EOS D60 using the same software that we also used for the scan camera. The color corrected Q60 chart for the EOS D60 is shown at the bottom right of Figure 7. When comparing this image to the chart for the scan camera after calibration (bottom left of the same figure), it can be seen that the colors are very similar. Another comparison of the color reproduction can be seen in Figure 9, which compares photos of a Chinese ornament taken with both cameras (the full size image for the scan camera is shown in Figure 11). Again, the colors produced by both cameras after calibration are very similar. The calibration software also reports statistics about the number of pixels that were outside the color gamut and could not be mapped one-to-one. For all images that we tested, the numbers were very similar. These results are not surprising, since we combine an off-the-shelf sensor with color filters that are widely used in photography.

Figure 9 also shows a comparison of the relative resolutions of the two cameras. We took the photographs in such a way that the full image covered roughly the same field of view. We then selected a comparable subregion from both images and printed it at the same size. That size was chosen such that every pixel in the scan camera image corresponds to a printer dot when printed at 300 DPI (the full image from Figure 9 fills a $34'' \times 40''$ poster when printed at that scale). Figure 11 shows a scaled down print of the full photograph.

To obtain quantitative information about the actual resolution of the complete optical system, we performed a measurement of the modulation transfer function (MTF). The MTF is a standardized way of evaluating the resolution of digital imaging devices, that is based on analyzing how different spatial frequencies are attenuated by the combination of optics and sensor. The most common way to compute the MTF is by doing a frequency analysis of a slanted, focused edge [ISO, Bur00, WB01]. In the case of scanning devices, the behavior in the rows and in the columns is analyzed separately [WB01]. We computed the MTF for different locations in the image, but the differences are relatively small. Figure 8 shows the results for the image center and the image corner.

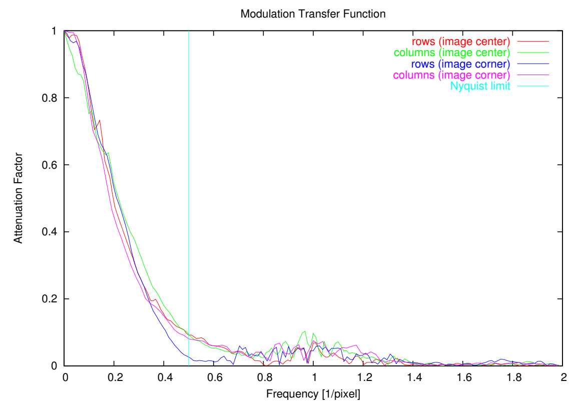

Figure 8: *The modulation transfer function for the rows and the columns of our optical system.*

Due to the Nyquist limit, the highest useful frequency any imaging system can have is 0.5 cycles/pixel. For optimal resolution, the MTF should drop to 10% or so around the Nyquist limit. If the function drops too early, the image appears slightly blurred, and the full resolution of the sensor is not used effectively. If, on the other hand, the MTF is large beyond the Nyquist limit, the imaging system has a tendency toward aliasing. Note, however, that transitions that are too abrupt (like a step function) are also not desirable, as they would introduce the well-known ringing artifacts of sinc filters.

As can be seen in Figure 8 the 10% cutoff for our camera lies at about 0.36 cycles/pixel for the rows (i.e. within a scan line), indicating a small amount of blur. For the columns, the cutoff is at about 0.54 cycles/pixel. Together, these values indicate that the optical system is very good, and operating at close to the nominal resolution we can expect from our scanner.

10. Discussion and Future Work

In this paper we described a 122-490 million pixel backend built from an off-the-shelf flatbed scanner. The backend itself costs about \$100-\$200, a complete camera system can be built for about \$1200. This compares to several tens of thousands of dollars for commercial digital scan backends. An additional benefit is that the proposed approach will scale well with improvements in scanner resolution.

We have in particular discussed software issues, such as image denoising and the removal of scanning artifacts, as well as color calibration. We compared our system to a commercial digital SLR camera that is twice as expensive as our system, and found the color reproduction to be comparable, but the image resolution to be vastly superior. Although it takes scan cameras longer than a single shot camera to acquire an image, this added time is tolerable for applications where setup and lighting take up a significant amount of time.

We think that the camera will be useful for scanning cultural objects and artwork for this reason. We are currently evaluating the use of the camera for this purpose with some professional arts photographers and a museum. The museum is also interested in acquiring NIR images of some of its exhibits for research purposes. This should be an interesting application for our camera system.

11. Acknowledgments

We would like to thank Michael Goesele for helpful pointers to MTF measurement literature and software. Thanks also to Dave Martindale for several discussions, and to Ted Gierschner from Lens&Shutter for help with selecting camera equipment. This work was supported by the BC Advanced Systems Institute and the Natural Sciences and Engineering Research Council of Canada.

References

[ANS03] ANSI: Standard IT8.7/2-1993: Color Reflection Target for Input Scanner Calibration, 2003.

[Bay] BAYER B. E.: Color imaging array. U.S. Patent 3,971,065.

[Ben04] BENDER PHOTOGRAPHIC:. http://www.benderphoto.com, 2004.

[Bet04] BETTERLIGHT:. http://www.betterlight.com, 2004.

[BSCB00] BERTALMIO M., SAPIRO G., CASELLES V., BALLESTER C.: Image inpainting. In *Proc. of ACM SIGGRAPH 2000* (July 2000), pp. 417–424.

[Bur00] BURNS P.: Slanted-edge MTF for digital camera and scanner analysis. In *Image Processing, Image Quality, Image Capture Systems Conference (PICS)* (Apr. 2000), Society for ImagingScience and Technology, pp. 135–138.

[Can04] CANON INC.: LED Indirect Exposure (LIDE). http://www.canon.com/technology/t_seihin/tech_scan.html, 2004.

[DCOY03] DRORI I., COHEN-OR D., YESHURUN H.: Fragment-based image completion. In *ACM Transactions on Graphics (Proc. of SIGGRAPH '03)* (July 2003), pp. 303–317.

[DWT*02] DEBEVEC P., WENGER A., TCHOU C., GARDNER A., WAESE J., HAWKINS T.: A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (Proc. of SIGGRAPH '02)* (July 2002), 547–557.

[Fie65] FIELDING R.: *The Technique of Special Effects Cinematography*. Hastings House, 1965.

[Fro02] FROLOV A.: SCARSE: Scanner CAlibration ReaSonably Easy. http://www.scarse.org/, 2002.

[Ham02] HAMRICK SOFTWARE:. http://www.hamrick.com/vsm.html, 2002.

[ISO] ISO: ISO 12233:2000 Photography – Electronic still picture cameras – Resolution measurements.

[Jan01] JANESICK J.: *Scientific Charge Coupled Devices*, vol. PM83. SPIE PRESS, 2001.

[LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D.: The digital michelangelo project: 3d scanning of large statues. In *Proc. of ACM SIGGRAPH 2000* (July 2000), pp. 131–144.

[PGB03] PÉREZ P., GANGNET M., BLAKE A.: Poisson image editing. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2003)* (July 2003), 313–318.

[Sin02] SINAR AG:. http://www.sinar.ch, 2002.

[SS97] SZELISKI R., SHUM H.-Y.: Creating full view panoramic mosaics and environment maps. In *Proc. of ACM SIGGRAPH '97* (Aug. 1997), pp. 251–258.

[Wal00] WALLNER D.: Building ICC Profiles – the Mechanics and Engineering. available at http://www.color.org/iccprofiles.html, 2000.

[WB01] WILLIAMS D., BURNS P.: Diagnostics for digital capture using MTF. In *Image Processing, Image Quality, Image Capture Systems Conference (PICS)* (Apr. 2001), Society for Imaging Science and Technology, pp. 227–232.

[Yan00] YANG J.: *A Light Field Camera For Image Based Rendering*. Master's thesis, MIT, 2000.

[YLIM00] YANG J., LEE C., ISAKSEN A., MCMILLAN L.: A low-cost, portable light field capture device. In *Siggraph Conference Abstracts and Applications* (2000), p. 224.

Figure 9: *Top: detail print of a small region from the Chinese ornament photograph in Figure 11, as taken with our camera. Bottom: the same region taken from the Canon EOS D60 photo for comparison (Figure 10).*

Figure 10: *The full photograph of the Chinese ornament taken with the Canon EOS D60.*

Figure 11: *A scaled-down print of a Chinese wall ornament. The full image was taken at 122 million pixel resolution and would fill a* $34'' \times 40''$ *poster if printed at 300 DPI. Also see Figure 9 for a detail image.*