

Prism-Based Color Separation for Professional Digital Photography

Richard F. Lyon
Foveon, Inc.
Santa Clara, CA USA

Abstract

In the field of silver-halide photography, color separation was explored about a century ago, using filter mosaics, dichroic mirrors, three-shot filtration, and other techniques, before the dominant technology of multi-layered color film emerged. In the field of digital photography, the same techniques are being explored now in conjunction with silicon sensor technology. The three-shot technique, and the related tri-linear scanning technique, can deliver excellent images, but are awkward to employ and unsuitable for subjects that move. The filter-mosaic approach is currently dominant in the digital camera market, but its inevitable aliasing artifacts limit its applicability in the professional end of the business. The use of dichroic mirrors, embedded in prisms as in color-TV cameras, has been attempted by a few camera vendors, so far with limited success. The demand for better digital images, however, has provided a continuing incentive to work on making this three-sensor approach feasible. Several key problems must be solved to make a viable professional camera using this approach: first, sensors with high enough resolution and image quality must be made at a cost that allows three of them to be included in a product; second, a prism must be designed for good color reproduction fidelity and must be economically and precisely manufactured; third, the three sensors must be optically aligned to the prism in a way that avoids ghost reflections and other problems; and finally, a whole new camera system needs to be designed around such a three-channel sensor assembly. This paper recounts some of the ways in which these problems are addressed in the design of the recently introduced Foveon Studio Camera, and provides qualitative assessments of the resulting advantages in image quality and usability.

Introduction

The prism-based, three-sensor, one-shot electronic studio camera and system recently introduced by Foveon came about as a result of understanding the history of approaches to color photography, and through an analysis of the opportunities to improve the quality of electronic imaging. Moving away from the mainstream of CCD sensors with color filter mosaics presented a number of technical

challenges, but has paid off in terms of image quality, especially in delivering an essentially artifact-free low-noise image.

Historical Perspective

The exploration of color separation approaches for electronic photography parallels in many respects the explorations that were done for silver halide photography around the previous turn of the century. But the outcome, in terms of a dominant high-quality technology, is likely to be quite different. Three-shot approaches have enjoyed limited success in both histories, but for general usage, one-shot approaches are required. As Friedman [1] noted in 1944:

To photograph an object in color, an operator has at his command a variety of methods. He can use a one-shot camera in which the lens beam is divided into three parts by any one of the many mechanisms described...he can use one of the integral processes, such as Kodachrome, Ansco Color, Dufaycolor, etc...Another possibility...is to photograph the subject through three filters in succession...only in cases where the object remains perfectly immobile...These are called repeating backs.

So we see the concepts of camera back distinct from color camera, and three-shot versus one-shot, but to see how the filter-mosaic approach fits in, we need to read further, because that approach had already slipped out of use at that writing:

The integral methods for the preparation of color photographs are not confined to the monopacks. Ducos du Hauron suggested a much simpler method almost fifty years before the monopack was even conceived... 1868... alternating lines of red, blue, and yellow... screen plates...

Friedman goes on to note the problems with the screen-plate patterns, the inefficiencies "due to the fact that the individual screen elements transmit but a fraction of all the light falling on them," and the "severe blow" dealt to this approach at the introduction of Kodachrome, Ansco Color, and Agfacolor. In the era of silicon sensors and

microlithography, the ability to pattern fine-scale filter mosaics has given this approach a resurgence. It is now the dominant approach to electronic color photography in spite of its problems of inefficient use of light and a tendency to produce Moiré patterns and other artifacts (errors known collectively as ‘aliasing,’ due to the spatial sampling of the colors being too sparse relative to spatial frequencies in the optical image).

In the digital domain, the screen-plate pattern is realized as a color filter mosaic aligned to individual pixel sensor elements, so that each location senses only one of the three primary colors through its appointed filter. The other two thirds of the color data at each pixel are then interpolated, or guessed, by various schemes that try to use image statistics to make up for the fact that all three color planes are under-sampled, in violation of the Nyquist sampling theorem.

It was against this backdrop that we realized that the image quality of modern digital cameras is limited by aliasing, not by the number of pixels; using more pixels in a filter mosaic merely pushes the problem to somewhat higher frequencies. We saw that by using three sensor chips on a prism, a given number of individual pixel-sensor elements could be used to produce alias-free images, and that all the light entering the camera could be efficiently used.

Furthermore, we realized that since high-end image sensors are always pushing the state of the art of chip fabrication, they are invariably on the steep part of the chip yield-versus-size or cost-versus-size curve; in this regime, it is much more economical to get more pixels from a set of three imager chips than from a single larger chip. These quality and economy advantages are simultaneously available by using three sensor chips behind a color-separating optical system.

High-Resolution Sensors

Our first major challenge was to come up with silicon sensor arrays that would support high resolution at a reasonable cost (since we need three of them per camera), and would enable the functionality of a fully digital camera, including a live viewfinder with focus capability. We found good reason to expect that the faster learning curve of CMOS chip fabrication, compared to CCD, would continue and would apply to imagers as well. And only CMOS could supply the flexible readout logic needed for our desired functions. So we set about exploring the space of possible image sensor circuits and architectures in CMOS and related processes.

In close cooperation with our semiconductor fabrication partner, National Semiconductor, we have been able to define modifications to a modern 0.25-micron low-voltage CMOS process that allow low-leakage analog functionality, efficient light sensing, and high-speed digital support and timing circuitry, all within a single chip type. These sensor chips are fabricated with a proprietary semiconductor process using mostly standard processing steps.

We have developed sensor and readout circuits that get around the shortcomings of more traditional CMOS imagers, including pixel density, fixed pattern (especially column)

noise, and readout speed, while supporting efficient readout of subwindows of arbitrary size and resolution for viewfinder, focus loupe, and autofocus functionality.

We decided early on to go straight to 2Kx2K sensor chips to make a 12-million-sensor imaging assembly. The sensor chips use 6-micron square pixels, defining the Foveon camera’s 12-mm square format.

The details of our sensors remain trade secret and patent pending until the corresponding patents issue.

Prism Design

The key novel optical component unique to the Foveon camera is the color separation prism that is designed to give a set of three spectral sensitivity curves that closely approximate a set of color matching functions. There has been a long history of color separation prisms in the color television industry, but the demands of professional photography are much more stringent in terms of optical resolution, color accuracy, and alignment between the three sensors.

Hunt [2] describes five strategies for color sensing and processing in color television, distinguishing the degrees and techniques of approximating color matching functions—i.e., sensitivity curves modeling human color—and of transforming the sensed measurements to an output colorspace. At Foveon, we took a more aggressive approach than any of Hunt’s five, and defined color separating dichroic reflectors that allow some violet light into the red sensor channel so that we could get a better approximation to a set of color matching functions that are simultaneously sharper and more nearly nonnegative than the ones that Hunt considered. A spectral shaping prefilter gives us another dimension of control over the spectral sensitivity curves, and we use it to do better than just the conventional IR-reject filtering.

Working with California-based optical fabrication partners, we were able to develop optical quality specifications and opto-mechanical assembly techniques to give us photographic quality imaging at a tolerable cost. A photo of a sample prism is shown in Figure 1 (see the color insert for figures).

Compared to the color accuracy of film, or of CCD filter-mosaic cameras (limited not in theory but in practice by the available patternable dye materials), Foveon’s color has been judged by users to be remarkably superior over a wide range of colored subject materials.

Alignment and Attachment

In tube-television and 3-CCD video technology, sensors have historically been aligned to the prism exit faces with an adjustable air gap. The CCDs usually have their own glass cover windows, and another air gap inside the package. The resulting multiple glass-air interfaces provide multiple opportunities for partial reflections, resulting in ghosting, fuzzing of the image, and loss of light. For the Foveon sensor assembly, we developed a direct-attach technology

with no air gaps, using a precision opto-mechanical alignment system with the camera electronics in an image-analysis feedback loop, and finally secured by an index-matched optical epoxy.

Besides eliminating reflections and light loss, Foveon's proprietary direct optical attachment also eliminates the possibility of dust settling into the focal plane after the camera is built. A corresponding challenge before the camera is built has been to extend the clean-room processing of the silicon into a continuous clean chain through all assembly steps up through attachment to the prism.

Format Equivalences

The Foveon camera format is a 12 mm square. Compared to the 24x36 mm format of a 35 mm camera for which the lens family was designed, our format is smaller by a factor of 2 in one dimension and a factor of 3 in the other, so the field of view is correspondingly narrower. Therefore, if you like to think in terms of focal length equivalent on a 35 mm camera, the lenses are effectively longer by a factor of 2.0 to 3.0. For an equivalently framed image with an 8x10 aspect ratio, the factor is 2.5. Alternatively, compared to lenses on a medium-format camera with 56 mm square format, the lens focal length is equivalent to a lens 4.67 times longer. But the game of equivalents is dangerous if one does not factor in the other implications of the different imaging format size.

In studying different format sizes, we discovered that an analysis of ray cones on the subject side of the lens leads to a simple format-independent invariance: the depth of field for a given field of view and subject distance depends only on the absolute aperture diameter. With the subject-side ray cones thereby fixed, cameras of different formats image identical circles of confusion from subject space into the camera focal plane space. Furthermore, the subject-referred diffraction blurring also depends on the absolute aperture and not on the format size.

Therefore, when comparing formats in terms of equivalent focal length, one can also use an 'equivalent f-number' that captures the fact the smaller-format camera gets a better depth of field for a given f-number, or same depth of field at a lower f-number. The higher equivalent f-number of the smaller format may be easily computed as the longer equivalent focal length over the actual aperture diameter. Furthermore, completing the format equivalence translation requires that to get a correct exposure, the smaller-format camera needs to be represented with a higher 'equivalent ISO speed.' The equivalent focal length and f-number scale with the format size, and the ISO scales as the square. Shutter speed and subject lighting do not vary.

The Foveon format's f-number can be said to be higher by a factor of 4.67, compared to medium format, and the effective ISO higher by the square of that factor, so shooting the Foveon at ISO 32 is like shooting medium format at ISO 700. That is, the small Foveon format can get as good sharpness, diffraction limit, and depth of field at f/5.6 as the medium-format camera needs to stop down to f/26 to achieve, and can do so at ISO speed 32, while the medium

format camera would need ISO 700. Correspondingly, the Foveon format's 35 mm equivalent ISO speed would be about 200. Since it is typical to shoot medium format around ISO 800 and 35 mm around ISO 200, it will be possible to capture equivalent images shooting the small Foveon format at ISO 32.

The actual ISO speed rating for the Foveon camera has not yet been measured according to the standard. Preliminary estimates put it at around 50, but since the Foveon sensors have considerable latitude for overexposure, we usually shoot using an ISO speed setting around 32 to get a better SNR than the standard requires.

Camera System

Just as we had decided to make a complete sensor system, rather than adapt a video sensor technology to do still photography, we decided to make a complete camera system, rather than adapt a computer to a film camera body. Many high-end digital capture systems are sold as camera backs that must be hooked to a conventional camera and tethered to a computer to form a complete camera; even then the result more resembles a computer with a peripheral capture head than it does a camera, and it usually takes computer experience to operate.

In examining the needs of a self-contained camera for control, storage, display, and communications, we did appreciate the advantage that a fully integrated modern computer could bring to the problem. So as the basis for a camera design we decided to use an embedded notebook computer to provide most of those functions, and to integrate it closely into a camera with our sensor system, rather than making the sensor separately as a peripheral. The corresponding challenges, of getting a good industrial and mechanical design to integrate the notebook with the rest of the camera, and of getting reasonable real-time performance from an available standard channel (SCSI via PCCARD in our case), entailed significant time and effort.

The camera components, including the notebook, are integrated using a curvaceous cast magnesium frame that defines the overall look of the Foveon camera.

The other significant part of the camera system design was the selection of a lens family to support. We elected to go with 35 mm format lenses instead of designing our own. We chose the EOS family of lenses available from Canon, Sigma, Tamron, Tokina, and others, due to the wide availability of excellent optics at reasonable prices, and due to the all-electronic control interface for aperture and focus.

A photo of the Foveon Studio Camera is shown in Figure 2.

Performance Comparisons

We have started down the path of applying various ISO standards and draft standards to characterize the Foveon in comparison with other cameras. Preliminary results indicate a larger dynamic range but lower ISO speed (by one to two stops) compared to several different 2Kx3K filter-mosaic

cameras. The primary reason for the ISO speed difference is the smaller format and smaller pixel sensors of the Foveon, since ISO speed is computed from a pixel signal-to-noise ratio (SNR) and since the number of photons seen by a sensor is proportional to its area. Normalizing for pixel area, the Foveon is found to measure photons as efficiently or better than CCD sensors, and as pointed out above, the lower ISO speed on a smaller format is not really much of a disadvantage.

When exposed for equal SNR criteria in dark areas (i.e. each according to its ISO speed), the Foveon shows better SNR in the lighter regions than other cameras do, indicating that we were successful in reducing our sensor's fixed-pattern noise to below the level of CCD noise.

Resolution is difficult to compare according to the standard, since the filter-mosaic cameras have so much aliasing and such different results with luminance and color. Moreover, the Foveon's advantages of good resolution of color structure and relative freedom from aliasing are not captured by the standard resolution measures, so we also compare on new resolution charts that we have created with various different contrasting color pairs.

Figure 3 compares a portion of a color resolution test chart that we use, as shot by the Foveon (Fig. 3a) and by a 2Kx3K filter-mosaic camera (Fig. 3b). For these shots, the 2K dimension of the Foveon's square format was framed to the same subject dimension as the 3K long dimension of the other camera's format, so the comparison puts the Foveon at a worst-case disadvantage in terms of resolution and sharpness. The portion shown corresponds to 2.67 by 3.25 inches from a 15x15 inch or 10x15 inch enlargement. If instead both cameras framed the same subject with a 4:5 aspect ratio, the Foveon's 1600x2000 would go up against the other's 2000x2500; i.e. the Foveon could be framed 5/6 tighter for a more fair and favorable resolution comparison. The highest frequencies shown at the left and right edges correspond to 1000 cycles, approximately equal to the Nyquist frequency of the 2Kx2K, below the Nyquist frequency for 3K pixels, but above the 750 cycle Nyquist frequency for the 2x2 pixel repeat unit of the CCD.

From the color chart photos it is clear that the image captured using full color measurement of three prism-separated color images is almost totally free of aliasing artifacts, because the spatial frequency response drops off gradually toward the Nyquist frequency (at the outermost edges of the image portion shown, approximately). The image captured with a filter mosaic and interpolation, on the other hand, shows a variety of artifacts that depend on the particular color combinations involved, and a spatial resolution that is also very dependent on the color combinations. The image remains 'sharp,' though not clear, beyond the Nyquist rate of the sensor repeat unit, where the sharpness is 'manufactured' by the interpolation algorithm and has its own nonlinear and jaggy artifacts. Since the processing is optimized for best response on black-and-white resolution charts, it is generally supposed that the luminance response is good enough and that blurring the chrominance will reduce the artifacts. When we apply correction filters

supplied by the CCD camera manufacturers, presumably based on these assumptions about the luminance component, we find that the chart image becomes nearly monochrome except near the center, that there is a severe further loss of resolution, and that the large-scale artifacts remain. After lots of realistic comparison shots involving fine-textured clothing, for example, we find that dealing with the artifacts of the filter-mosaic sensor is simply not feasible in general. For many of our customers, this difficulty is the single biggest reason for selecting the prism-based technology.

Figure 4 compares images of a region of blue sky as a way to show the different noise statistics of the two cameras; blue sky is an easy source of a uniform subject with less structure than typical gray or white charts. Both images are enhanced to show the noise in the processed file, preserving the relative SNRs and pixel scales. In each image, the blue channel is shown above the diagonal, and the full RGB is shown below. Note that the Foveon image (Fig. 4a) has a small uniform-looking noise, whereas the image captured by a camera using a bare multi-block 2Kx3K CCD with filter mosaic in the focal plane (Fig. 4b) shows an oriented noise, sensor block boundaries, shadows of dust particles in the focal plane, semi-periodic dark spikes, and other patterned variations peculiar to that technology. These visible defects, since they are very localized, may not have much impact on standard deviation measurements and SNR calculations, especially when the measured area misses the block boundaries and dust particles—but they tend to dominate the visual difference between the camera noise patterns. Foveon's one-block image sensor chip architecture and sealed focal plane preclude such problems.

Usability

The Foveon camera was designed primarily for studio use, where it fits into standard lighting setups and is used like any other tripod-mounted camera. It is often found wired to power, a network (for downloading images to process or review at another station), a flash sync cord, and a shutter release cable. In studios, some photographers prefer to operate with radio devices for both shutter release and flash, eliminating those wires. Since the camera contains its own batteries, and does not need to be networked or tethered to any other devices to function as a complete camera, some photographers remove the power and network connections, too, for better mobility. The camera is sometimes used outdoors and on the road this way, for several hours at a time, with as much portability as is typical of a view camera.

Our software team was continually challenged during the development process with the slogan "It's a camera, not a computer." The resulting interface minimizes the use of on-screen buttons, cryptic key combinations, etc., in favor of hard-labelled keys. White labels call the photographer's attention to understandable photographic functions, such as aperture and shutter speed adjustment, new roll creation, auto focus, and shutter release. A wired shutter release is of course also supported, as is a sync jack that works with all electronic flashes. Most photographers find that they are

able to walk up to a Foveon and operate it without instruction.

The live viewfinder on the left side of the camera's screen, in combination with the movable magnifying loupe subwindow, allow easy composition and real-time focusing. Focus can be done manually (by turning the lens focus ring), electrically (by moving the focus position using the touchpad while holding the Focus key), or automatically (by pressing the Auto Focus key).

The instant preview that shows up on the right side of the camera's screen after a snap provides immediate confirmation of the captured pose, composition, focus, and exposure. A histogram shows the range of tone values captured in the image, and an over-exposure warning calls attention to regions with loss of highlight detail. The same pointing action that moves a focusing loupe around the viewfinder window can move a review loupe around the instant preview window, so that full-resolution fine details can be examined even on the limited-resolution screen. No separate meters or sensors are needed for setting exposure or focus, since the real result from the sensor is immediately available for review. Like a Polaroid image used to preview lighting and focus, it gives quick feedback; but there is no distinction between the instant preview and the final capture, so you don't have to worry about losing your best pose to a test shot.

The software that runs on the Foveon Processing Station (a part of the bundled Foveon Studio System) is also designed to be very simple and easy to use, with all processing options visible, not buried deep in menu and dialog structures. Since this program can very rapidly move through and display the preview images computed at the camera, many customers find it to be useful as a sales presentation tool, as well as using it for final processing.

Conclusion

We expect that due to its inherent advantages of one-shot alias-free color-accurate sensing, the prism-based color separation approach will become the dominant approach for high-end professional photography. The Foveon Studio Camera is the first electronic instrument to successfully apply this approach at a resolution and quality level competitive with medium-format film.

References

1. Joseph S. Friedman, History of Color Photography, American Photographic Publishing Co., Boston, p. 135, 197 (1944).
2. R. W. G. Hunt, The Reproduction of Color, 2nd ed., John Wiley & Sons, p. 335 (1967).

Reprinted from the Proceedings of IS&T's PICS 2000 Conference:
Image Processing, Image Quality, Image Capture Systems Conference
March 26-29, 2000
Portland, Oregon
pp. 50-54

©2000 The Society for Imaging Science and Technology

extra reprints including the original color plate may be obtained from the author:

Richard F. Lyon
Foveon, Inc.
3565 Monroe St.
Santa Clara, CA 95051
(408) 350-5110
richard.lyon@foveon.net

more information may be found at <http://www.foveon.net>




Fig. 1. A color-separation prism.



Fig. 2. Foveon Studio Camera.




Fig. 3. Color resolution/alias test; a: Foveon image (above); b: CCD image (below).




Fig. 4. Enhanced blue-channel/RGB sky noise test; a: Foveon image (above); b: CCD image (below).

