

Foveon Technology and the Changing Landscape of Digital Cameras

Paul M. Hubel, Foveon, Inc., Santa Clara, California, USA

Abstract

Increasing competition and asymptotic image quality improvements in point-and-shoot digital cameras has led to two late-blooming but profitable markets: the affordable digital SLR and inexpensive camera modules. In this talk I will describe this new digital camera landscape and how the color imaging conference has helped this industry evolve. The talk will focus on recent advances in Foveon's X3 sensor technology and new image processing concepts that have helped point the technology into the evolving digital camera market. I will discuss simulation work that we undertake at Foveon to help with the design of sensors and image processing algorithms including the theory of how silicon can be used for color separation, sensor noise modeling, and a hyperspectral imaging model. Lastly, the advantages and importance of color resolution will be addressed.

Introduction

As the camera market has changed, Foveon has been working hard to follow the trend to higher quality and lower cost sensors. Through a rigorous program with our partners and in the design of or sensors we have achieved significant improvements aimed at both the high-end large area camera sensors and those aimed at lower-end markets.

Modeling of X3 silicon color separation

The mechanism of color separation used in the Foveon X3 sensors¹ relies on the absorption of photons at different wavelengths and at different depths. The higher energy photons, those at the blue end of the spectrum, are absorbed at the surface whereas the lower energy photons penetrate deeper into the silicon substrate before they are absorbed. The wavelength-dependent absorption coefficient of silicon, and corresponding mean penetration depth, are plotted in Figure 1². In the Foveon X3 sensor, regions within the depth of the silicon are formed by transitions between different doping gradients and used to separate the electron-hole pairs that are formed at different depths by this naturally occurring property of silicon. The depths of these transitions are the key variables that determine the spectral sensitivities of such a device. In the talk I will show a demonstration of a Matlab model that computes quantum efficiency or spectral sensitivity curves of multiple layers at a chosen set of thicknesses. Figure 2 shows the theoretical spectral sensitivities computed by the model alongside actual sensitivity data from an early X3 sensor measured using a monochrometer.

In addition to the color separation model we have been working on models of all of the assorted noise sources that contribute to the overall noise in the image capture process. From a long list of quantities that includes variables such as photodiode and sense

node capacitance, well depth, pixel size, fill factor, and read out time we can generate estimates of the noise levels that we expect to correspond with the quantum efficiency curves generated from the color separation model. We use these color and noise models, alongside a model of ISO speed^{3, 4} and a model of metamerism index⁵ to compare the tradeoffs between sensor design parameters and these two figures of merit.




Figure 1 Absorption coefficient and penetration depth in Silicon vs. wavelength.




Figure 2. Dashed curves show the color separation computed with a simple model of light absorption in silicon and the solid line show measurements measured from an early Foveon X3 image sensor.

Combining the two Foveon X3 specific models of color separation and sensor noise sources with hyperspectral image data⁶ and an illumination model we can simulate the camera capture process and calculate raw camera data. Color images will be shown in the

presentation that will show the similarity between a simulation of an image containing a Macbeth ColorChecker chart created with the model described here alongside a captured and processed image taken with a Foveon X3 sensor.

This combination of sensor and noise models along with a hyperspectral image model has been extremely useful for both the design of sensor characteristics and optimization, but also in developing new image processing routines (discussed below), where having images with typical noise levels and ideal noise-free raw data can form the basis of the research process.

Although there are many factors that have an impact on the sensitivity curves of X3 sensors, the basic shape of the curves – what matters for color reproduction – remains stable and reproducible. This consistency is in part due to the extremely high degree of accuracy that can be obtained in epitaxial silicon growth – especially compared to the tolerances of color filter dye concentrations and deposition thickness. In markets where camera calibration is too expensive, this stability is one of X3’s key advantages. As described in our previous paper at this conference¹, the color accuracy that can be obtained with sensors with this kind of spectral overlap can be extremely high such that metamerism is negligible.

New Image Processing Methods

As with every aspect of life, the technological elegance of the color separation method and the high color accuracy come at a price: the color correction transformation matrix required to convert the native sensor signals to a standard color space is aggressive compared with conventional camera sensors and this amplifies noise. This noise amplification necessitates the use of advanced image processing algorithms that lower noise without adversely affecting other aspects of image quality. Many noise reduction techniques have been reported in the image processing literature and several are appropriate for our application. The desire of lowering computational complexity, however, becomes most important when entering into markets where image processing hardware must be simple, small, and fast.

One new method we are using at Foveon involves the use of two parallel pipelines of data (such as those shown in either figure 3b or 3c) rather than the conventional single pipe technique (figure 3a). In higher-end cameras or in raw processing software, one can afford the luxury of image buffers and multi-pass methods. The challenge here was to design a fast completely pipelined process that uses a minimum of line buffers so that data could be processed continuously at video rates. In cameras designs that depend on the use of a rolling shutter, data must be read out and processed quickly to avoid motion artifacts. For video, the on-chip binning of

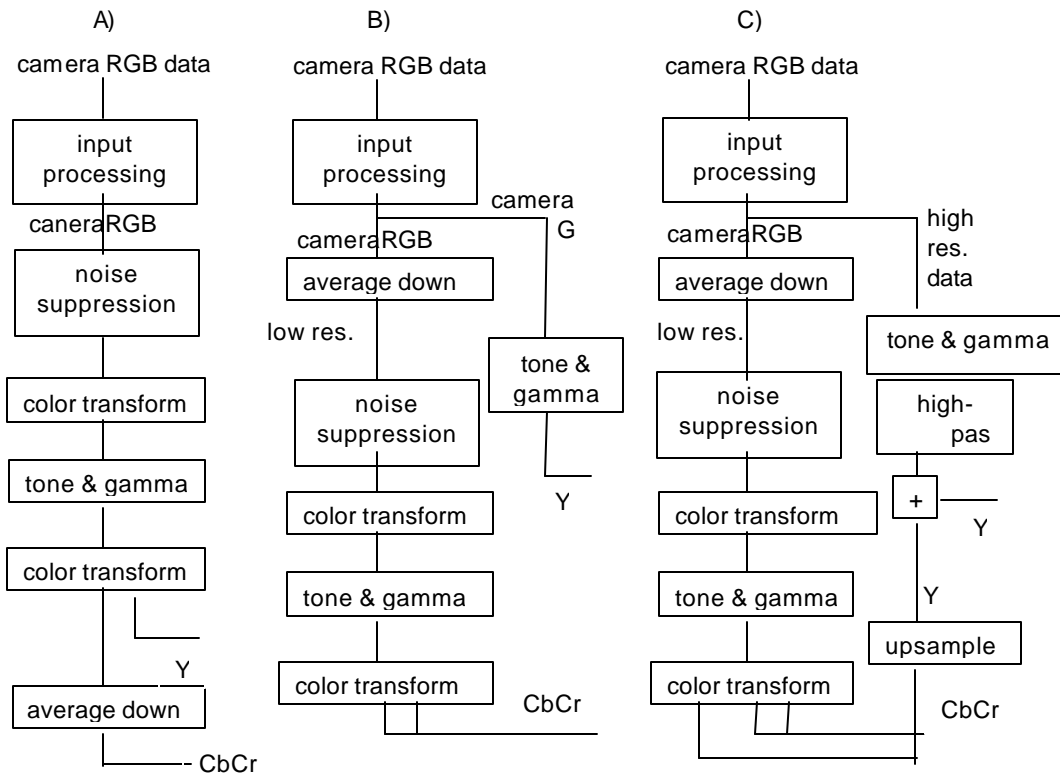


Figure 3. A) a normal single processing pipeline where all data passes through all blocks at full resolution. Matrix noise amplification can be minimized by using a two-pipe processing scheme such as: B) a pipeline where only a low resolution copy passes through noise suppression and color transformation and C) a pipeline similar to B) but where an upsampled copy of the color corrected Y channel is combined just before final output of the Y channel

groups of pixels help reduce the bandwidth requirements (In Foveon X3 sensors we use a unique variable pixel size (VPS) technology to bin neighboring pixels⁷).

The main insight of the method shown in figure 3 came first from the realization that the level of detail and low noise image quality of our raw camera data was extremely high compared to the results we would get when processing through a conventional digital camera pipeline (such as figure 3a). The recognition that most of the luminance and chrominance noise that was noticeable in the final image was due to the color transformation matrix magnifying the noise, led us to develop this two-pipe process that avoids sending high frequency details through the color transformation at all. In fact, given that most image data is compressed using some kind of method that averaged down the chromatic components of the image data, the branch of the processing pipeline could optionally be done at lower resolution than the high frequency branch. This idea, which follows from the early design of television broadcast systems, also significantly reduces the level of computation; the most expensive noise reduction and color transformation computations can be done on 1/4 or less of the data depending on the resolution difference between the two branches of the process. The result of the lower resolution pipe can then lead directly to the CbCr chroma channels and the high resolution branch leads to the Y channel of the output image.

Depending on the characteristics of the sensor and the IR filter, the *camera-Green* channel can be use directly as the resultant Y channel (because, as discussed in our previous publication, the broad bandwidth of X3's green channel can be designed to give an accurate match the human eye's luminous efficiency curve - the Y tristimulus function). On the other hand, if the high resolution path is run through a high-pass filter, this data can be combined with the low resolution Y data from the chroma path to form the final Y plane. In this case, there is not any special requirement concerning the spectral sensitivity of the high frequency data path.

In the presentation I will show a magnified area of an image where the data is processed with a single full resolution pipeline compared with the result from the two-path pipeline. The significant savings in computation of the latter method allows for a computationally inexpensive method that can be implemented in a hardware or software pipeline for use in middle and low-end digital camera products.

Raw Camera Processing

The methods described in the previous section apply well to consumer photography where the appropriate use of image compression make digital photography manageable in terms of storage and transmission requirements and yet maintain image quality at or above consumer silver halide photography. For professionals and advanced amateurs, however, there is increasing demand for uncompressed and RAW mode camera data storage and software that processes these images. The advantages of RAW mode operation are numerous, including the ability to control white-point, color transformations, noise suppression, tone control, sharpening and the ability to apply advanced powerful

algorithms such as Foveon's X3 FillLight technology (described below). Having all of the information that was recorded in the camera at the time of capture – at least beyond the point of analog to digital conversion – gives the utmost in flexibility for the rendering stage

X3 FillLight.

X3 FillLight⁸, which is a form of high dynamic range rendering, allows for quick one slider adjustment of a local tone correction method that controls the shadow and highlight relationship. Having all of the bit depth of the captured scene allows the photographer to bring up the shadow areas without introducing quantization noise that would come if the method were applied to previously rendered image. The method retains expected color and texture which has been shown to be a difficulty in previously reported methods. The photography community has welcomed this kind of raw mode processing and there have been many examples sited where FillLight has save an otherwise hopeless photograph. Although this kind of image recovery is more common with amateur users who do not have the best understanding of photographic lighting, it has also been embraced as a new tool for professionals and allowed them to change their shooting style. In the presentation I will shows a demonstration of the kind of powerful image correction that can be done using the X3 FillLight technology.

Color Aliasing

Foveon's unique X3 technology allows a significant advantage in optical resolution per unit area of silicon and at the same time avoids the color aliasing issues that occur when using color sub-sampling filters.

In cameras that employ color filter array (cfa) mosaics aliasing occurs in the three color planes at different locations. This artifact, referred to as color aliasing, appears as brightly colored fringes in locations of image detail at a particular spatial frequency. This artifact has been compensated for in most digital cameras by using an optical blur filter (also called an optical low-pass filter or anti-aliasing filter) but as the name suggests, the use of such a filter can lower the camera's resolution. Fabricating these filters for the large image sensors found in DSLR cameras is expensive and often professionals prefer to use the sharper un-blurred optics by avoiding photographing fabrics that are prone to color aliasing. (Many studio photographers instruct their clients to avoid finely striped clothing – example.) In the very low-end markets, the cost of these blur filters is prohibitive and the current quality requirements are not such that these filters are used. There is a common misconception about blur filters and cheap cameras that I would like to address: a poor lens does not help avoid color aliasing. If we consider the MTF of a typical camera, color aliasing occurs if there is significant modulation energy at or near the Nyquist frequencies for the color planes of the camera. When a birefringent filter is added to the optical path, it can remove the modulation energy at one position – usually chosen to be just below the Nyquist frequency. The advantage of this kind of filter is that it does not reduce the contrast at lower frequencies. A poor quality lens, on the other hand, lowers contrast at lower

frequencies resulting in an unsharp image and allows plenty of modulation at the Nyquist frequency that will still result in color aliasing – the worst of both worlds!

Using a direct image sensor such as Foveon's X3 technology, reduces the requirement of a blur filter because aliasing occurs in all color channels equally giving no brightly colored aliasing fringing. Unfortunately, this un-susceptibility to color aliasing is difficult to quantify in terms of a resolution measurement technique (such as ISO 12233) but clearly it can have significant impact on the reproduction of fine details.

Color Resolution

One aspect of higher-end photography that Foveon has been active in is characterization and measurement of color resolution. Although it is well known that luminance resolution is more important, especially in bandwidth limited situations such as television, the ability to accurately render colored details, color textures, and colored fabrics cannot be overlooked. This includes the ability to accurately render single-pixel color details as well as avoiding color aliasing. It is hard to imagine why all resolution measurements, through ISO standards and several image analysis software packages, all measure resolution using a completely black and white target. Usually these targets do not even have any grayscale information – so they really are *only* black and white. Yet essentially all digital cameras on the market today record in color and the scenes people are photographing are usually color.

In recent work⁹ we have discussed several techniques for measuring resolution using colored targets and from this work we can see why this has been avoided by the digital photography community. The Nyquist limits for the individual color components of color filter array cameras are significantly lower than if the color filters were ignored. Despite some pushback, we are working to at least include an informative annex into the revision of ISO 12233 Resolution Measurements for Electronic Still cameras that will include one of the two methods we have proposed. The first method would use a colored target that would include a slanted edge between two colors such as red and blue, green and blue, etc. or a target that would include a colored sinusoidal Siemens Star pattern (similar to the monochromatic target being proposed for inclusion in the next ISO12233 revision). The analysis of images taken with either of these targets could employ one of the available analysis routines. (Several of the color resolution targets we have used will be shown during the presentation). Another proposed method would use the standard black and white targets and their existing analysis methods but have these photographed through color separation filters¹⁰. This method would be easily implemented and controlled and would be an easy way to isolate a camera's ability to render modulation of patterns formed by primary colored objects. By simple specification of the appropriate color separation filters a method could be suggested that avoids the problems and variation involved in fabricating a colored resolution target. Discussions on this topic are currently being held within the ISO 12233 experts group.

Summary

As pixels become smaller the engineering of the depth of the sensor is becoming more and more important. Foveon's unique X3 direct image sensors have been working into the depth of silicon material and avoid the problems of plastic color filter array deposition on top of the sensor. A model of silicon's ability to separate color components is shown and the results are compared to experimental data from an existing X3 sensor. This model along with models of sensor noise, ISO speed, metamerism index, and illumination are used with hyperspectral image data to simulate full camera operation. In addition to many other activities, the results from the simulations have been used to develop new image processing methods including a two-pipe method which has helped apply X3 technology to lower-end markets.

The other exciting market for both Foveon and digital photography is the digital SLR market. Raw processing software that includes Foveon's X3 FillLight for high dynamic range rendering is described and shown during the presentation. Finally, the importance of color aliasing and color resolution is discussed and several new measurement techniques are described.

The exciting growth in both high-end and low-end digital photography we hope will hold new opportunities for the color imaging work at Foveon and for the continued wealth of new ideas presented at the Color Imaging Conference.

References

1. R. Lyon, P. Hubel, "Eyeing the Camera: Into the Next Century", IS&T/TSID 10th Color Imaging Conference Proceedings, Scottsdale, AZ, USA; 2002 pp. 349-355
2. A.J.P. Theuwissen; Solid-State Imaging with Charged-Couples Devices, Kluwer Academic Publishers, Dordrecht, 1995
3. ISO 12233:2000 Photography-Electronic still picture cameras-Resolution measurements.
4. R. Baer, J. Holm: A Model for Calculating the Potential ISO Speeds of Digital Still Cameras based upon CCD Characteristics. PICS 1999, Savannah, Georgia, USA, April 25-28 1999: 35-38
4. ISO 17321 Graphic Technology and Photography – Colour characterization of digital still cameras: <http://www.i3a.org/wg20>.
5. D. H. Brainard, Hyperspectral Image Data, <http://color.psych.ucsb.edu/hyperspectral/>.
6. <http://www.foveon.com/article.php?a=71>
7. <http://www.foveon.com/article.php?a=73>
9. P. Hubel, et. al., "Spatial Frequency Response of Color Image Sensors: Bayer Color Filters and Foveon X3." Proceedings of SPIE, Vol. 5301, EI'04, San Jose, CA, 2004; p. 402-407
10. P. Hubel, "Resolution for Color photography" to be presented at SPIE-EI, San Jose, CA, January 2006

Author Biography

Dr. Paul M. Hubel has been a Principal Scientist at Foveon, Inc. since 2002. Before joining Foveon Dr. Hubel worked for ten years as a Principal Project Scientist at Hewlett-Packard Laboratories working on digital cameras, photofinishing, scanners, copiers, and printers. Prior to this, Dr. Hubel worked at the Rowland Institute for Science and then at the MIT-Media Laboratory. Dr. Hubel received his B.Sc. in Optics from The University of Rochester in 1986, and his D.Phil. from Oxford University in 1990. Dr. Hubel has published over 30 technical papers, book chapters, and authored 25 patents. Dr. Hubel is a member of IS&T and SPIE, and has served as the technical and general chair of the Color Imaging Conference.