Real-time color imaging with a CMOS sensor having stacked photodiodes

David L. Gilblom* [a], Sang Keun Yoo[b], Peter Ventura[c]
[a]Alternative Vision Corporation, P.O. Box 4055, Los Altos, CA, USA 94024-1055;
[b]HanVision Co., Ltd., KAIST-AVH, 373-1, Guseong-dong, Yuseong-gu, Daejeon, R.O. Korea
[c]Peter Ventura, Foveon, Inc., 2820 San Tomas Expressway, Santa Clara, CA, USA 95051

ABSTRACT

High-performance color image acquisition has heretofore relied on color video cameras using multiple image sensors mounted on spectral separation prisms to provide geometrically accurate color data free of reconstruction artifacts. Recently, a CMOS image sensor has been developed that incorporates three complete planes of photodiodes in a single device to provide color separation without the need for external optical elements. The first commercial device based on this technology has 1512 x 2268 three-color photosites on 9.12 micron centers and includes provisions for combining pixels in X and Y, region-of interest selection and sparse scanning. The camera described in this paper operates the sensor in a variety of scan modes offering tradeoffs between resolution, coverage and speed. In this camera, a 128x128 raster of either a matrix of this size or binned from a large area can be scanned at nearly 150 frames per second and a single 2048-element line can be scanned at 7 KHz. At full resolution, the image sensor will acquire four frames per second. The scan configuration can be reloaded in less than 50 microseconds permitting mode changes on the fly.
.

Keywords: Image sensor, CMOS, CCD, color, multispectral, camera

1. INTRODUCTION

The development of a commercial color sensor incorporating layered photodiodes has brought a new technology to bear on the acquisition of color images. This sensor, described by the authors in a previous paper[1], consists of a two-dimensional array of three-photodiode stacks in which the junction depths of the three photodiodes are placed at depths that exploit the steep wavelength gradient in the optical absorption length of silicon to produce three color-separated signals.

The device is fabricated using 0.18 micron CMOS process, permitting provision of other digital and analog functions on-chip. In the specific sensor described here, the Foveon® X3™ Pro 10M color image sensor, these additional functions include a set of registers controlling scan geometry, amplifiers and switches for pixel buffering and sampling, controls for output multiplexing and averaging and various analog control inputs.

To take advantage of these numerous control functions, a camera, the HanVision HVDUO-10M, was designed incorporating a full suite of software controls for sensor functions and real-time processing of the image data to produce accurate color images.

2. LAYERED PHOTODIODE COLOR SENSING

Semiconductor materials have characteristic photon absorption lengths that vary with the energy of the photons absorbed. Silicon, fortunately, has such a curve that varies by nearly two orders of magnitude over the visible range (Figure 1)[2]. This variation provides sufficient space to stack multiple diode junctions at depths that are both capable of separating photons of various wavelengths and amenable to fabrication using standard CMOS manufacturing processes.

---

* dave.gilblom@alt-vision.com; phone +1-650-625-0318; fax +1-650-240-4005; www.alt-vision.com

2.1 Junction depths

In practice, junctions at depths around 0.2 µm, 0.8 µm and 3.0 µm provide workable spectral separation for true color imaging. With appropriate junction construction, the top diode can collect charge nearly to the surface of the silicon to extend response into the near ultraviolet and the bottom diode can deplete well into the substrate to provide extended near-infrared response.

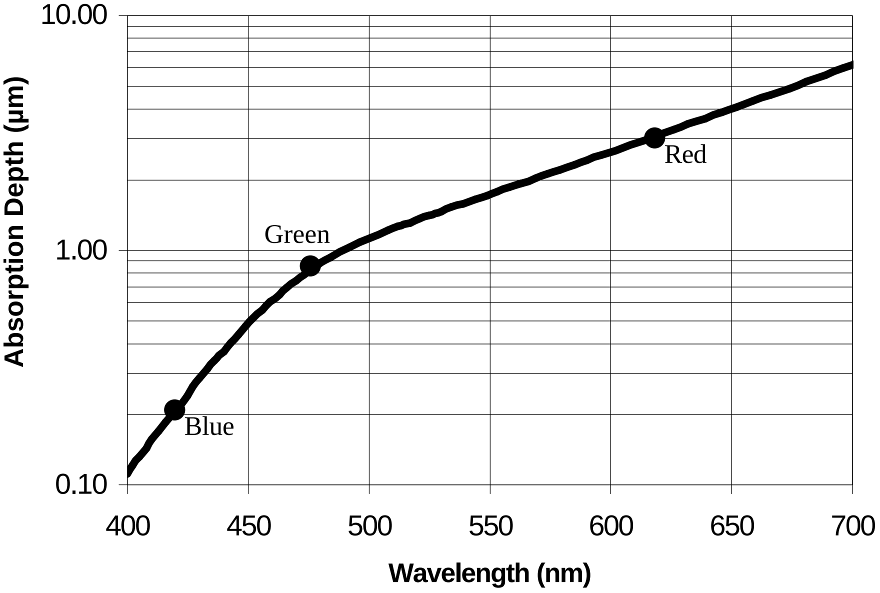

Figure 1 - Silicon photon absorption depth with practical junction depths indicated

2.2 Spectral characteristics

Even with the large changes in absorption depth with wavelength, the response curves of devices using the semiconductor material overlap considerably (Figure 2)[3]. The steep slope in the silicon curve in the 400-475 nm range provides substantial separation of the blue signal from the red and green below, but the relatively shallow slope above 475 nm results in a significant contribution of longer wavelength illumination to the top two signals. Fortunately, the relatively thin absorption regions of the top two diodes minimize this. In addition, some of the short-wavelength photons will make their way into the middle diode. It is this overlap that makes possible the discrimination of wavelength below 450 nm that is so difficult using color filters. The extended response at both ends of the visible spectrum also makes incorporation of a sharp-cut visible filter essential. The curves in figure 6 include the effects of a filter with cutoffs at 400 and 660 nm.

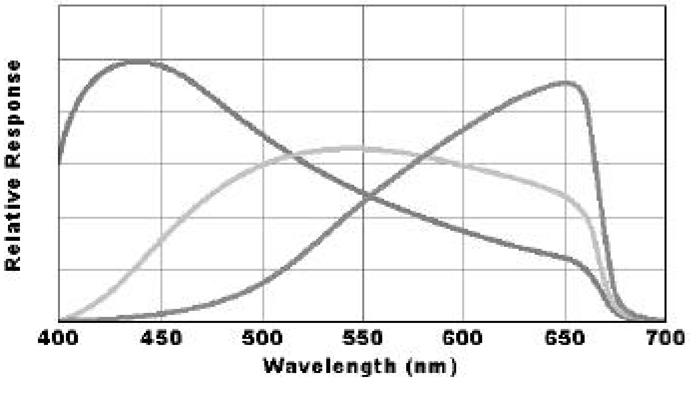

Figure 2 - Relative spectral response of three stacked junctions

2.3 Color transformation

The fundamental requirement for production of color images, that is, images that can be presented in such a way that the eye perceives them as having color characteristics reasonably approximating the colors in the original, is that at least three samples of the original must be taken using detectors with spectral characteristics that can be combined to produce results sufficiently close to the spectral characteristics of three normative curves known as the tristimulus (or color-matching) functions (Figure 3). It is significant that the spectral characteristics of the detectors need not match any actual spectral characteristic of receptors in the eye. Within these constraints a variety of schemes for providing three or more channels of spectrally-separated image data can be employed to produce reasonable color images.

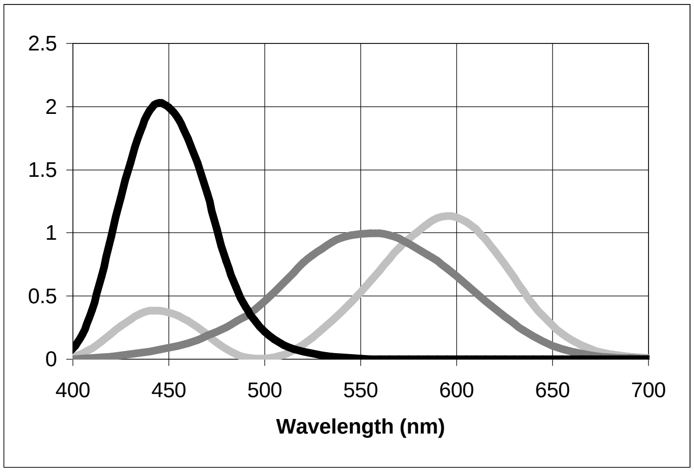

Figure 3 - Tristimulus (color matching) curves (CIE 10° 1964)[4]

If the acquired data can be transformed into a reasonable representation of the tristimulus color space using a simple 3x3 matrix relating the sensed data to tristimulus data, then further simple transforms are possible to permit the sensed data to be displayed by various common techniques producing color results that match to some reasonable degree the colors in the original image. For example, tristimulus data may be transformed by a set of formulas including display gamma correction into the sRGB space[5] used for many computer displays.

3. THE IMAGE SENSOR

To produce a practical layered color imager, Foveon developed in a 0.18 micron, 3.3V CMOS process, a triple-diode structure with junction depths of 0.2, 0.8 and 3.2 µm joined with active pixel signal elements and peripheral timing and readout functions. This image sensor, designated the F7 - now known commercially as the X3™ Pro 10M – (Figure 4), has the characteristics shown in Table 1.

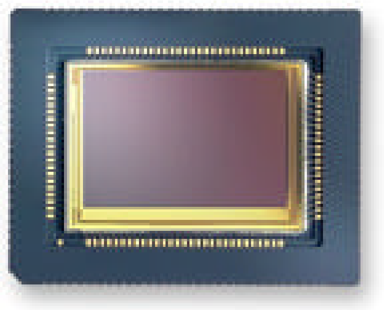

Figure 4 – Foveon X3 Pro 10M CMOS color image sensor

Parameter	Value
Pixel pitch	9.12µm x 9.12µm
Pixel locations	2304 x 1536 (total); 2268 x 1512 (active)
Total photosensors	10.2 million
Active area	20.7mm x 13.8mm
Fill factor	~54%
Output	3 analog
Package	100-pin ceramic leaded chip carrier (CLCC)
Window	Glass with 400-660nm multilayer visible pass filter

Table 1 – X3 Pro 10M image sensor characteristics

In addition, this sensor includes precautions against internal reflection: the inside of the window is coated with a visible range antireflectance layer and the non-active areas between the photosensors are covered with a black mask material. There is no anti-aliasing filter included.

3.1 Operating functions

The use of standard CMOS process affords the opportunity to include extensive control functionality in imagers. In the F7, several useful functions have been implemented.

3.1.1 Power management

Low voltage CMOS and fully-static clocking in the sensor keep the power consumption to 80mW during readout when all counters and signal paths are active. To support battery operation, the sensor also includes a standby mode with 10mW consumption and a power-down mode that maintains register contents while consuming 100µW.

3.1.2 Scan control

Scanning in the sensor is controlled by two sets of counters – one each for horizontal and vertical control. The settings of these counters are loaded through a dedicated serial port as a single data stream. Loading takes less than 50µs. The counters determine which row is to be activated and which pixel location in that row is to be read out. Once the counters are loaded, the counter operation is controlled by external clocks. Each counter has three registers to set the start count, the increment and the stop count separately in the horizontal and vertical directions. These can be set in any combination to define a rectangular region of interest from one pixel location to the entire array. The increment controls permit sparse and reverse scanning.

To facilitate more rapid scanning of any selected area, the sensor supports a function similar to binning in CCDs in which adjacent pixel locations can be read out in groups designated Variable Pixel Size (VPS™). VPS uses two controls in each of the horizontal and vertical directions – one to set the number of pixel locations, in powers of 2, to be grouped and the increment control to match the spacing to the VPS settings. Grouping in CMOS devices is not like binning in CCDs. In CCD binning the charges are added from the binned areas and read out in one packet. This serves to increase the signal but also runs the risk of overloading the readout registers in bright areas. In the sensor, the grouping connects the voltages generated by the grouped pixels to an amplifier node. This reports an analog combination of the applied voltages rather than a sum and reduces the fixed pattern and random noise rather than increasing the signal.

3.1.3 Exposure control

The sensor includes a global reset control that allows the charge stored in all photodiodes to be dumped and a line reset control that dumps only a selected line. The global reset is used primarily for still shot (SS) imaging of the type common in digital still cameras. In this mode, which uses an external mechanical shutter, the exposure cycle is very simple. To start the cycle, the sensor is globally reset, then, scanning is stopped, the external shutter is opened, an external flash may be fired and the external shutter is closed. After completion of the exposure, scanning is initiated. The external shutter may be a blade or electro-optic type, in which the entire sensor is exposed simultaneously, or a curtain shutter, in which a slot is drawn across the sensor, as is typical in 35mm film cameras. An equivalent cycle may be used where the shutter is replaced by a pulsed light source, but, in this situation, care must be taken to assure that the sensor is kept in the dark during scanning.

For flexibility, the sensor includes a separate counter to enable lines for reset, which be used together with the vertical readout counter to set a wide range of integration times between reset and readout of each line. This rolling shutter (RS) scan mode is an electronic equivalent of a curtain shutter. The counters can be configured to wrap around at the last line so that the integration time is constant for all lines. In the RS mode, the global reset is never activated.

Since the reset and readout functions are separately controlled, it is also possible to read the same line repeatedly without resetting it. This non-destructive reading permits monitoring of signal buildup during extended integration times to assure maximum dynamic range without overload. With careful signal management, high-speed binned monitoring can be followed by full-resolution readout without damage to the image data.

3.1.4 Signal control

Because the sensor is essentially three image sensors stacked one upon another, three independent sets of setup voltages may be supplied. External access is provided to these three sets of analog amplifiers and references wherever an adjustment might produce useful results. In signals where channel tracking is most important, the inputs are internally tied. A full discussion of the uses and effects of the analog voltage adjustments is beyond the scope of this paper, however, these do affect well capacity, linearity, antiblooming, voltage output levels, frequency response and reset levels.

3.2 Signal characteristics

The output from the sensor is three trains of negative-going pulses representing the voltage values read from the three layers of photodiodes. The clock rate is nominally 24 MHz. Each line time consists of a transfer interval, during which internal signal stabilization takes place followed by a clocked readout. For 12-bit accuracy, the transfer interval is 49µs. This duration is fixed regardless of the number of pixel locations read out per line. There is no requirement for waiting at the end of the frame; however, scanning can be stopped here to increase the integration time in the RS mode.

The response of the photodiodes to incoming light is sublinear, primarily as a result of the transfer characteristics of reverse-biased diodes used in a voltage-readout mode. Additional suppression of the response curve is intentionally added using an anti-blooming control voltage to prevent accumulation of charge in the photodiodes in excess of that to be used in the final signal.

Analog processing required off-chip is minimal, consisting typically of a buffer amplifier and then a 12-bit digitizer per channel. Output impedance is 200 ohms. Since each column (or the whole frame) is reset simultaneously, access to individual pixel reset signals at the output is not available. The conversion efficiency is 7.14µV/electron and the maximum signal output level is 550 mV.

3.3 Performance

The total quantum efficiency of the sensor at 625nm is approximately 49% including the effects of fill factor. Total quantum efficiency is over 45% from about 530nm to beyond 660nm. Testing is underway to establish the limits of wavelength response. The sensor is expected to have useful sensitivity extending from below 300nm to 1000nm or higher. Well capacity is approximately 77,000 electrons per photodiode but the usual operating point (for restricted non-linearity) corresponds to about 45,000 electrons. Photo response non-uniformity (PRNU) is less than ±1%.

Several fixed-pattern and random noise reduction techniques have been incorporated into the sensor design to realize very good noise performance for the CMOS technology. The total fixed pattern noise from all sources is less than ±1%. The primary contributor to dark noise is ktC noise from diode reset. This noise is approximately 70 electrons. It is possible to reduce this to about 40 electrons by implementing a reset-read-expose-read cycle for the frame and then subtract the first frame from the second. Lag is zero.

The typical dynamic range of the sensor is 61db. This can be increased by the use of a larger portion of the non-linear portion of the transfer curve as long as care is taken to carefully correct the non-linearity. The signal-to-noise ratio is shot-noise limited with an exposure approximately 10% of the nominal maximum signal point.

Dark current is approximately 1.0 nA/cm^2 at 25C, allowing exposures up to several seconds without cooling. Noise contribution from dark current is very small and dark current uniformity is better than ±1%.

4. IMPLEMENTATION REQUIREMENTS

Building a camera to accommodate the x3 Pro 10M image sensor is complex because of the flexibility in both the selection of analog operating points and the variety of potential scan modes. Tools are available to assist with this process. In addition, two measurements must be made for each sensor to assure highest image quality. These measurements correspond to the sensor-based corrections – linearization and color conversion. These parameters can vary, although over fairly narrow ranges, from sensor to sensor. Both linearization tables and color matrices should be determined for each sensor to assure the most accurate color results. Other data that might be measured for correction as required by the application include blemishes, shading and fixed pattern noise. It should be noted that all of these parameters can be influenced to some degree by the sensor operating conditions.

4.1 Optical considerations

Although the stacked photodiode structure eliminates artifacts generated by the offsets of the color receptors in color filter array sensors, some subtle effects remain that should be considered. First, there is a potential with stacked photodiodes to have the sensitivity vary with f-number because lower-angle rays might escape detection by escaping from the side of the diode. However, the diode geometry is designed to minimize this and the effect has not been demonstrated. Next, the presence of low-angle rays might shift the color response due to the variation in ray path length through the silicon. This is possible, but because of the high index of refraction of silicon, incoming rays are bent strongly toward the normal. For most practical optics, the path length variation would be 20% at most. This effect could probably be demonstrated but has not yet been rigorously examined.

Although the sensor optical effects are generally minor, the quality of the optics used with the stacked-photodiode sensors can seriously impact image quality. Most significant is chromatic aberration, which is quite visible in monochrome images taken with the sensor. Generally, color filter arrays mask chromatic aberration because geometrically the effects are on the order of one pixel. However, images taken with this sensor clearly show radially-symmetric color fringing resulting from lenses that suffer from an excess of this problem. Similarly, unsharpness in focus, astigmatism, and other lens defects can have clearly-observable negative effects on images made with stacked-photodiode sensors.

4.2 Processing requirements

In order to optimize dynamic range, noise and uniformity, certain modifications can be introduced into the signal as it is generated on the sensor. These modifications then need to be accurately reversed by the post-processing to provide maximum image quality. In the X3 Pro 10M image sensor, these modifications are non-linear so maintaining the proper order in reversal is important.

4.2.1 Linearization

The non-linearity in the transfer characteristic must be accurately reversed to prevent shifts in color with brightness. Without correct linearization, the color matrix calculations will produce values that depend on the position of each of the color components on the non-linear curve. If gamma correction is to be applied to the image, this must be done after all linear color calculations are completed.

4.2.2 Dark field subtraction

Several improvements can be obtained by subtracting an image acquired without illumination from the scene image. If the integration time is long enough to produce a dark current signal that is significant, then dark field subtraction can reverse the baseline shift. Shifts in the baseline will affect color matrix calculations. In the still image mode, the dark field should always be subtracted because the dark current will vary linearly from the top of the image to the bottom.

This can cause visible shifts in color vertically in the final image. Dark subtraction can also effectively reduce stationary fixed pattern noise.

In the best case, the dark field should be calculated from a series of frames so that the reset noise can be averaged out. In practice, a single frame must often be used. Still, the increase in noise is usually more than offset by the reduction in non-uniformities. The dark field must undergo the same column filtering and linearization as the scene image to assure accurate subtraction.

4.2.3 Optional steps

Depending on the application, some additional processing steps might be included. For viewed images, compensating for sensor blemishes might be useful. The algorithms for these are the same for this sensor as for monochrome sensors. No special processing is required to accommodate the geometry of color filter arrays. Finally, although this sensor has relatively low non-uniformity, shading correction might be applied, especially if significant optical rolloff exists.

4.2.4 Color space conversion

The actual color space to be used will be determined by the application. If the image is to be directly displayed then a single sensor to sRGB space with gamma correction might be sufficient. If the image is to be used in a color management system, then sensor to XYZ conversion might be appropriate. Understanding the nature of various color spaces and the effect of white point selection is very important to producing images with the desired color fidelity. In many industrial and scientific applications, no consideration is given to color space issues because the narrow filters and the narrow display primaries restrict the range of accurate color representation sufficiently that simple white balance of the output signal is sufficient to realize the full potential of these systems.

4.3 Monochrome imaging

Because devices with stacked photodiodes have co-axial sensors for the entire range of detected wavelengths, the signals from the diodes (after proper first-stage corrections) can be summed to produce a monochrome image. The effect is that of having a good broadband photodiode at each pixel.

4.4 Image samples

Many of the effects mentioned in this paper can be seen in actual images. These do not reproduce well in halftoned images and especially do not survive conversion to black and white. Samples and links to additional images can be found online at http://www.alt-vision.com/r/5210.htm.

5. CAMERA DESIGN

In the HVDUO-10M camera, most of the control features of the X3 Pro 10M image sensor are made accessible to the user and the processing steps essential to producing real-time color images are implemented. This initial model was designed to incorporate all of the desired control and processing functions in a configuration that allows maximum flexibility for design updates. As a result, the package size is not as small as could be realized in a camera with a relatively static design. Even so, a reasonably compact, rugged package was possible including common optical, electrical and mechanical interfaces for ease of installation.

5.1 Electronics complement

The control and output interfaces to the sensor require a combination of analog and digital controls. Basically, the scan control is digital, a combination of settable registers to control scan counting and clocks to provide timing. To control the transfer of image signals inside the sensor, several real-time waveforms that apply particular analog voltages to the right pins at the right time during scanning are also required. Useful changes in sensor operating points can be realized by adjustment of some of these voltages so DACs are provided to facilitate external control. Generally, these voltages would be fixed in a camera designed for a specific application but in this camera the value of allowing the user to experiment is preserved.

5.1.1 Block diagram

The electronic requirements for this camera are similar to those for any three-chip color camera. Since the scanning is all on a single chip and the sensor is CMOS, no scan driver s are required. As a result, all of the digital control can be achieved using standard CMOS logic. In this design, almost all of this logic is contained in FPGAs. To maximize algorithm development flexibility, this design uses a separate FPGA for each function, six in all, but it is possible to combine all of these functions in a single device in cameras designed for narrower application. Figure 5 shows the block diagram of the entire camera.

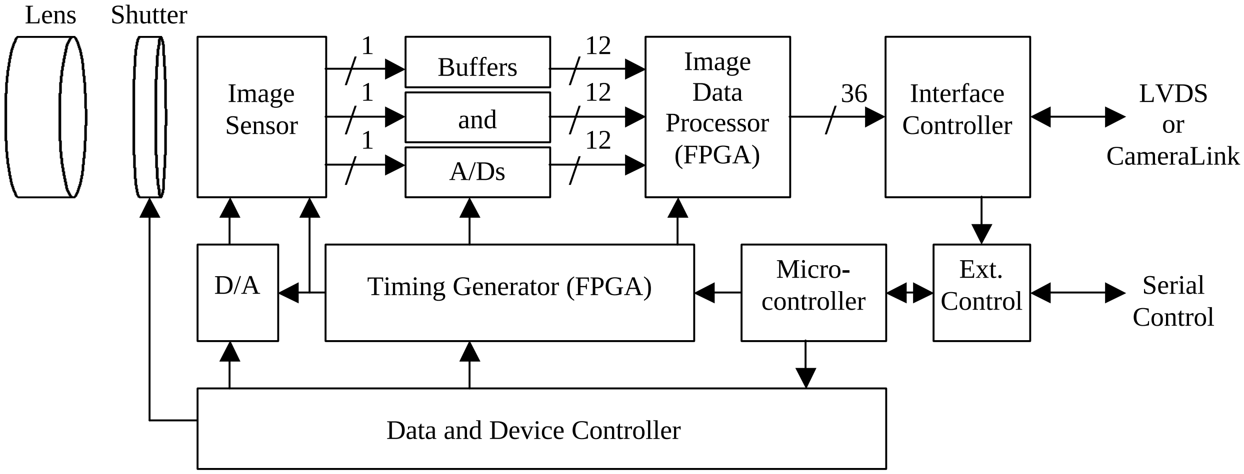

Figure 5 – HVDUO-10M Block Diagram

5.1.2 Control and timing

Overall camera control is managed by a stored-program 8-bit microcontroller. This receives configuration information from an external serial port (either separate, in the LVDS version, or embedded in the CameraLink port) connected to the host computer. Based on this information, the microcontroller sets registers wherever required. The registers control both the state of various devices and the settings to be used in timing generation. The data and device controller (DDC) coordinates the operation of various devices requiring real-time command and response.

The timing generator, a stored-program FPGA, generates the timing waveforms that drive the sensor and coordinate the processing of data from it. Some outputs of the timing generator switch, in real time, voltage levels set by a group of D/A converters that allow voltage control of eight analog input sensor pins. Some pins require as many as three different operating voltages to accommodate the various operating modes.

5.1.3 Image data path

The output of the sensor is three analog signals with inverted polarity. These are digitized to 12 bits and provided to the image data processor, which carries out the three steps described earlier – linearization, dark frame subtraction and color space conversion. Linearization is accomplished with three loadable 4096-element lookup tables that also restore the data to positive polarity. Both dark subtraction and color conversion use signed 16-bit integer arithmetic with careful control of intermediate results to assure that the data is not truncated.

The data for the linearization tables are stored in non-volatile memory under the control of the microcontroller. This permits loading of new tables in the event of a sensor change or linear tables to permit examination of the raw sensor data. The dark frame data is stored from an image that can be taken on command. A new dark frame is acquired automatically whenever any change is made to the camera settings that renders the current dark frame data invalid. The user may also take a new dark frame at any time. Dark frame subtraction may be disabled on command.

The coefficients for the three color space conversion equations are supplied to the processor from the host. These are stored in non-volatile memory and may be read back to the host on command. Users may notice that coefficient values read back may change slightly from the values sent. This results from conversion to a fixed-point 16-bit number. The range of permissible values is –7.000 to +7.000 to offer a wide range of digital gain without risking the generation of contours in the image. The output of the processor is truncated to 12 bits per color for transfer to the output interface.

5.1.4 External interfaces

The HVDUO-10M is offered in both parallel LVDS and Camera Link models. In the LVDS model, the image data is output to three 8-bit channels at the same pixel clock rate used at the sensor. When the camera is set to the 12-bit output mode, the data is transmitted in two clock cycles at twice the pixel clock rate. The serial control is incorporated into the multipin connectors on the camera rear panel but split out to a separate D connector to mate with a standard serial COM port.

In the CameraLink model, the standard 3 x 8 bit RGB Base mode is used. Double-speed transmission is used for 12-bit data as in the LVDS model. The serial control is embedded in the CameraLink MDR-26 connector.

For flash or other pulsed illumination applications, the camera also includes a flash sync connector.

5.2 Optical layout

The standard models of the camera incorporate a mount compatible with Nikon F-type bayonet lenses. These have a 46.5 mm flange focal distance. Commercial T-mount (42 mm x 0.75 mm pitch thread) adapters may be used with this mount to permit mounting on optical instruments. Since the sensor active area is nearly 14 x 21 mm, C-mount lenses are not suitable. The shutter is mounted behind the lens. It should not be used to control exposure because the exposure time with shutters in this position is not uniform. In addition, care must be taken not to exceed the shutter rated lifetime. The camera is available with shutters rated for 300,000 and 1 million operations.

5.3 External control

Since this camera was designed as an evaluation tool for potential users of these sensors, providing extensive access to camera settings was very important. To facilitate use of these controls, a program, the HVDUO Configuration Control, was developed. This program combines a complete set of camera controls with a real-time viewer. It runs under Windows 2000 or Windows XP. The command set was implemented using only simple ASCII text strings so that users could easily incorporate the commands into other programs under various operating systems.

5.3.1 Control Program

A screen shot of the HVDUO Configuration Control (HDCC) program is shown in Figure 6. A full description of the controls is available but is beyond the scope of this paper. Only an outline of its capabilities will be given here. There are several functional areas simultaneously displayed on the screen. At the top are the normal Windows menu functions, tailored to the needs of this program. There are a few processing functions available here, which may be applied to frozen images, but serious processing or analysis should be done with software intended for that purpose. Below this is a button bar containing controls for managing files, initiating and ending acquisitions, toggling data displays and setting the view port magnification.

Below the button par is the viewport, in which real-time sensor images will be displayed. In general, the images here will not be ideal because what will be viewed is linear camera data, which appears dark on monitors for which that data has not been gamma-corrected. If a true representation of the original gray scale is desired, the display system gamma must be set to 1. This will wash out typical computer-generated graphics, which expect a gamma around 2. The viewport can also display the RGB intensity profiles of any line, selected with a right mouse-click.

To the right of the viewport is a set of tabs providing access to all adjustable camera settings. The “Process” tab shown includes commands for acquiring dark images, toggling dark subtraction, managing linearity tables, and designing, loading and storing color conversion matrices. At the bottom of the tab is a progress bar that operates whenever communication is in progress with the camera

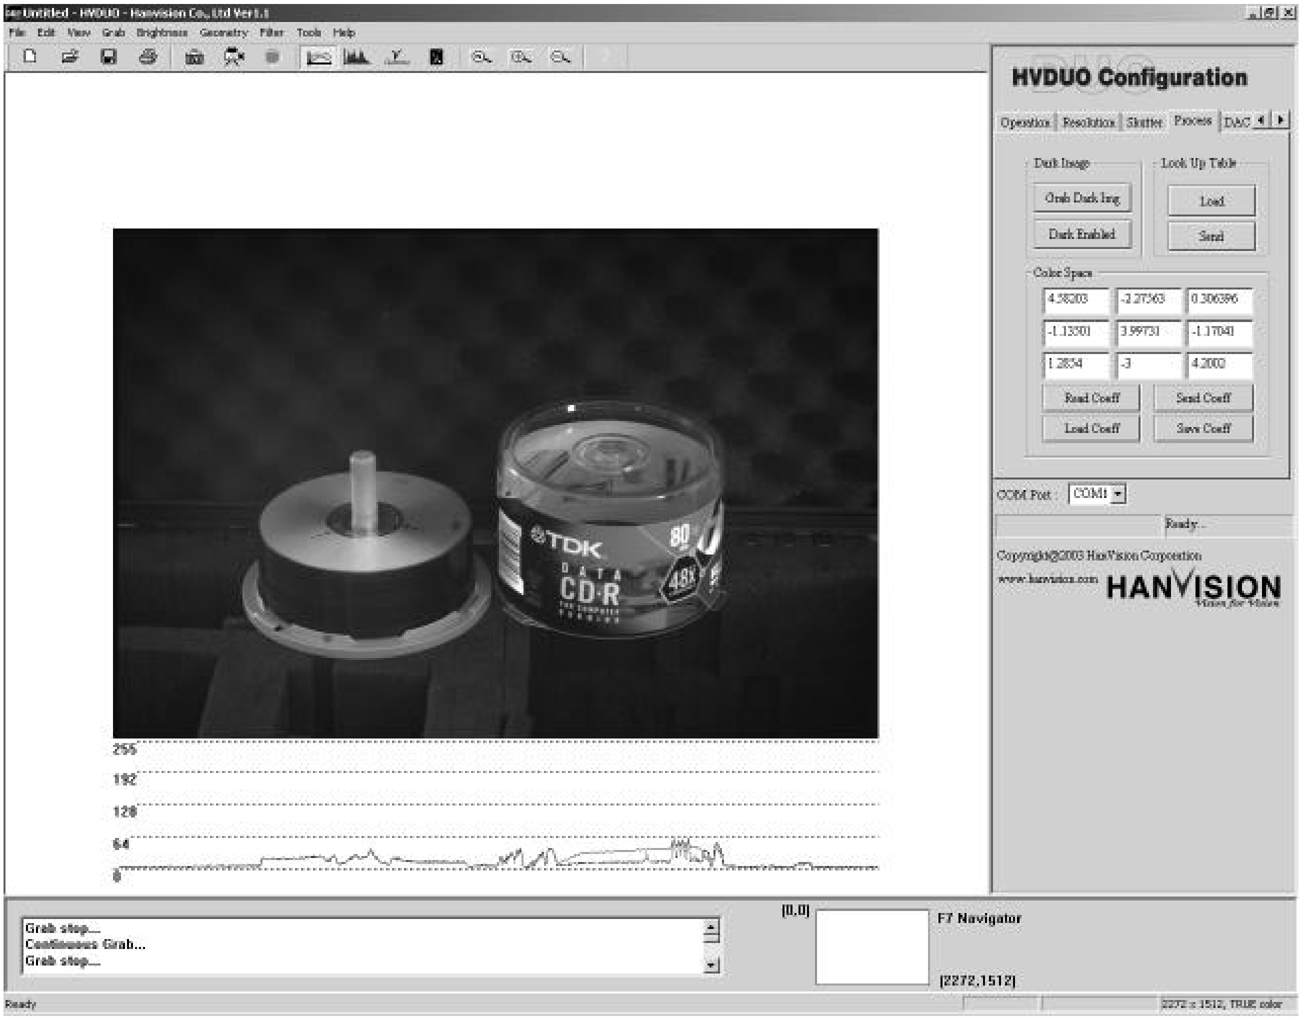

4.58203	-2.27563	0.306396
-1.13501	3.99731	-1.17041
1.2834	-3	4.2002

COM Port : COM1

Ready...

```
Grab stop....
Continuous Grab...
Grab stop....
```

[0,0]

[2272,1512]

Ready

2272 x 1512, TRUE color

Figure 6 – HVDUO Configuration Control Screen Shot

At the bottom are, left to right, a running record of actions executed, a display showing which part of the sensor is currently used for imaging, and, just a bit lower, the color data for the pixel under the mouse pointer, blank here because the pointer is not in the viewport.

5.3.2 Resolution selection

It is an unfortunate fact that the vast majority of framegrabbers must be told in advance the size of the image they are to capture. This information is generally provided to the framegrabber by the host computer in the form of a setup file. Since a separate setup file is required for each configuration, it is not feasible to allow unlimited flexibility in setting the size of the active pixel matrix in the camera, even though the sensor does permit this. As a result, the selection of raster sizes has been limited to about 10 selections on each axis, covering a reasonable variety to permit some variation in acquisition speed. These run from the full sensor down to 128 x 128 pixels. Pixel grouping sizes in powers of 2 from 1 to 16 are also provided, matched with similar row and column skip settings. The raster can be set to start on any pixel.

In a custom camera, additional flexibility can be offered, including the capability to have raster position or size switched rapidly under real-time control to support true dual-mode applications.

6. PERFORMANCE

The performance of the HVDUO-10M camera (Figure 7) conforms to the basic capabilities of the sensor. In the 8-bit mode, the dynamic range is limited by the number of available bits and the S/N is shot noise limited throughout the digitized range. In the 12-bit mode, the noise floor is about 8 DN for typical color matrices. The noise consists almost entirely of KtC noise resulting from the lack of the capability for correlated double sampling (CDS) in CMOS imagers. Methods for canceling this noise exist but these carry penalties in either acquisition time or in substantial additional sensor complexity that were judged too burdensome to implement at this time.

Figure 7 – HVDUO-10M Camera

Acquisition speeds are as calculated. No changes in image quality have been seen as a result of changes in scan size. Combining pixels produces an effect somewhat different from that observed with CCD binning. Since the result of combining pixels is to produce an average, in general the apparent average brightness of the image will drop when pixels are combined. This results from the suppression of small bright areas by the combining process.

7. FUTURE DEVELOPMENT

Currently, the Foveon sensor is available in two versions, one, used here, with a visible-pass filter on the window and another, not discussed, with a clear glass window. In a camera with the clear glass window, images over the band extending from the near UV to the silicon cutoff at 1100 nm is possible. This suggests that dual-mode cameras are possible – color with the visible pass filter and UV or IR with an appropriate cutoff filter. Cameras with more than one sensor could be built, taking advantage of the three simultaneous detection bands to produce multispectral images with a reduced sensor count. Also feasible is the addition of a spectral tailoring filter to produce six-band visible imaging to produce images with very accurate color characteristics.

Since the device is silicon, it is also responsive to cooling to permit longer integration times. Finally, simce this is a CMOS device, increasing readout speed or improving noise performance are directly amenable to known design solutions. This should, as the use of Foveon X3 devices grows, enable design and production of a variety of devices suitable for use in high-speed instrumentation applications.

8. REFERENCES

1. David L. Gilblom, Peter Ventura and Sang Kuen Yoo, “Operation and performance of a color image sensor with layered photodiodes”, Proc SPIE Vol 5074. in press
2. Michael H. Jones, Stephen H. Jones, “Optical Properties of Silicon”, Virginia Semiconductor, Inc., August 2002
3. Foveon, Inc., “Foveon X3 Pro 10M CMOS Image Sensor”, January 2003
4. Color and Vision Research Laboratories, http://cvrl.ucl.ac.uk/ciexyz64.txt
5. Michael Stokes,, Matthew Anderson, Srinivasan Chandrasekar, Ricardo Motta, “A Standard Default Color Space for the Internet – sRGB”, Version 1.10, November 5, 1996, from http://www.w3.org/Graphics/Color/sRGB.html