
Application Note

DIGITAL CCD CAMERA DESIGN GUIDE

DIGITAL CAMERA BASICS

The Digital Trend

Electronic cameras have often been surrounded by an air of mystery. Certainly we have seen the extensive refinement of the camcorder over the years, with a rich feature set and advanced image processing common in most units. Fundamental to this refinement is the improved quality of the image processing electronics at the heart of these cameras. This same advanced image processing power can now be integrated into a few easily connected IC's to form some of the most capable cameras yet. As the electronics have evolved into system level components, the design of an electronic camera has moved from the arena of black art to mainstream. Designing electronic cameras to meet a wide range of needs becomes a straightforward task using the Crystal Digital Imaging CCD camera chipset.

What Goes On Inside the Camera

Inside the CCD or CMOS imager camera, incoming light from the lens is gathered by a photo-gate or photo-diode array respectively. Light hitting this array is converted into electrical energy, which is then processed by the remaining electronics to create representative brightness and color data for the image. After this basic processing has been completed, timing and synchronization information is added to the image data. The output is then fed to a monitor for display, or recorded on tape or disk for later use.

Video Formats

In most cases the combined video and timing information is output as the familiar analog composite video used in NTSC and PAL based systems. These two video formats have dominated the worldwide broadcast television industry, and have served the industry well for many years. Unfortunately, the analog formats always suffer from image degradation whenever an image is transmitted (encoded/decoded), processed, or stored (as on video tape).

Today a number of digitally oriented video formats are emerging. Some of these include the consumer P1394 or "FireWire" serial standard, the professional video ITU-656 transport (International Telecommunications Union), as well as a number of very high-speed serial digital protocols over coaxial cable. The digital nature of these newer formats is very attractive since transmission, processing, and storage can be accomplished without degrading the image quality.

Modern electronic cameras typically have a CCD (Charge Coupled Device) imager at their heart to collect the light from the lens. In addition to the lens and CCD imager, the camera will make use of a timing generator to clock the CCD imager, as well as control the timing of the analog processing unit which converts the CCD imager output into coded brightness and color signals. This analog processing unit makes use of a number of analog functions, and in the case of a digital camera, ultimately digitizes the CCD imager output for further processing. Fully analog CCD cameras will perform the required image processing using a number

of discrete and integrated components, all controlled by the timing generator.

A digital CCD camera first converts the CCD analog output to digital using an analog-to-digital converter and all subsequent image processing is done digitally. This digital signal processor approach has the added advantage of being highly programmable. This wide range of programmability allows the camera electronics to be configured, and re-configured to meet changing light and environmental conditions. The CCD camera electronics chip-set from Crystal Digital Imaging supports a wide range of camera formats in terms of image size and quality by changing only the CCD imager itself. The design of the camera electronics remains the same across these designs, requiring only minor changes in the programmable setup configuration.

OPTICS AND LENS BASICS

A CCD camera lens system is designed to focus an image on the CCD sensor, allowing the CCD sensor to collect the incoming light correctly. Following are some of the terms that are used to describe a CCD camera lens system: field of view, magnification, overall image brightness and exposure, depth of field, etc. Many of these terms will require some explanation to see how they apply to CCD cameras. As background, the key optical terms and relationships will be discussed with an emphasis on how they relate to CCD cameras. More complete discussions of optics are available from other sources, and we do not attempt more than a basic coverage of this material. Even so, many lenses are simple catalog items, and camera designers should be able to communicate their needs with lens manufacturers based on the included material.

The Pin-Hole Lens

In film photography, pin-hole cameras are the most basic (see Figure 1). They are simply a photographic film material placed in a light-tight container with a very small *pin-hole* directly in front of the

film. Light entering the pin-hole is very nicely focused on the film. The film is exposed by allowing light to enter the pin-hole, while the camera is held stationary. When used with a CCD camera, the mechanics of the pin-hole lens are unchanged. The focused image is formed on the CCD imager, provided the hole is small enough to qualify as a pin-hole. In both the photographic and the CCD camera configurations, there should be no other path for light to enter and strike the CCD imager or film, except the pin-hole itself.




Figure 1. Pin-Hole Lens

The pin-hole lens has several very attractive features: simplicity, low weight, and infinite depth of field (it is always focused). The only significant drawback is that the light gathering capability of a pin-hole lens is relatively small owing to the small aperture or opening size. To be effective, the pin-hole itself should be quite small when compared with its distance from the imager, or film.

Field of View

The field of view for a pin-hole lens can be described by the angle formed between two lines which originate at the edges of the image, and intersect at the pin-hole. For example, the vertical field of view for a CCD imager is the angle formed by one line originating at the top-center of the imaging area, and a second line originating at the bottom-center of the imaging area. The horizontal field of view would be measured as the angle formed from lines originating at the left center of the imaging area, and the right center of the imag-

ing area. The field of view might also be referenced to the diagonal, or corner to corner, measurement of the imager area.

Aspect Ratio	Imager Diagonal Size	Height (mm)	Width (mm)	Area (mm) ²
4:6	35 mm	24	36	864
3:4	1/2"	4.9	6.5	32
3:4	1/3"	3.7	4.9	18
3:4	1/4"	2.7	3.7	10
3:4	1/5"	2.2	2.8	6

Table 1. Basic Imager Dimensions.

The farther the pin-hole is from the imager, the narrower the field of view becomes; the closer the pin-hole is to the imager, the wider the field of view becomes. A 1/3" CCD imager has an imaging area of 3.7 mm x 4.9 mm, so with a 20 mm focal length, the field of view coverage is 10.5 degrees vertically by 14 degrees horizontally, and 17 degrees diagonally. With a 50 mm focal length, the field of view coverage for a 1/3" imager becomes 4.2 degrees, 5.6 degrees, and 6.9 degrees respectively.

Focal Length

Lenses are specified based on the equivalent pin-hole lens image. For example, a 55 mm lens for a 35 mm format film camera forms an image with a 35 degree diagonal field of view. This is roughly equivalent to the image formed by a pin-hole lens placed 7mm away from a 1/3" CCD imager.

Aperture

The aperture is the size of the lens opening when viewed from the front of the lens. The larger the aperture, the greater the light gathering capacity of the lens. Unfortunately, the larger the aperture, the more prone the lens becomes to undesirable optical aberrations (image defects that result from limitations in the lens materials and design). The optimum aperture for a pin-hole lens is listed below.

$$\text{Aperture}_{\text{optimum-pin-hole}} = 36\sqrt{\text{focal length in mm}}$$

f-Stop Number

The f-Stop or F Number of a lens is the focal length of a lens divided by the diameter of the lens aperture. The brightness of the image on the CCD imager is inversely proportional to (f-Stop)². Pin-hole lenses will invariably have very high f-Stop numbers since the comparatively long focal length is divided by the small pin-hole diameter. This means that pin-hole devices have very limited light gathering capacity, which would imply long exposure times for a properly exposed image. By incorporating the refractive capability of glass and plastic lenses, a larger opening can be employed, while still maintaining a typical focal length. The reduced f-Stop numbers will have higher light gathering capacity for shorter exposure times.




Figure 2. Optimum Pin Hole aperture and Resulting f-Number.

Lens Elements

Lens elements bend or refract the light passing through to form images. The number of lens elements varies depending on the applications, but one to five elements are common. The different lenses are used in combination to compensate for the optical aberrations of various materials. By varying the refractive index and other optical qualities of the individual elements, the overall optical quality

of the entire lens assembly can be drastically improved. See Figure 3 for further information.




Figure 3. Basic Lens Description.

Glass vs Plastic

Glass has been the premier lens material since the time of Galileo, and glass lenses are available in a wide range of qualities and optical properties. Glass lens elements are unfortunately more difficult to grind and manufacture than plastic lens elements are to mold. Plastic materials though, are not available with as wide a range of optical properties, which restricts the number of applications where low-cost plastic lenses can be used. Plastics do excel in the area of “spherical” lens elements where they are easily molded into the proper shape. As a rule of thumb, three to five element glass lens assemblies are quite capable of meeting the needs of desktop videoconferencing. Lower cost three element plastic lens assemblies are showing some promise in this area. Applications requiring higher image quality, as would be the case in security cameras, will likely employ one of the high-quality, C-Mount style, glass lens systems.

Depth of Field

Depth of field describes the difference between (1) the shortest distance from the camera where objects appear to be in focus, and (2) the distance farthest away from the camera where objects appear to be in focus. The depth of field varies with the focus setting of a lens assembly. Some settings may have a very shallow depth of field, while other focus set-

tings might have a very deep depth of field. Pin-hole lenses essentially have an infinite depth of field.

Optical Filters

Optical filters are used to perform spectral (e.g. infrared) filtering and spatial filtering (blur or anti-alias filtering). Typically these elements are made of quartz or glass. Optical filters are placed between the lens assembly and the CCD imager. Most filters do not require specific orientation relative to the CCD imager pixel pattern, and could be included in the lens assembly itself. Filters for polarizing light would likely need to maintain a constant position relative to the imager, and are more likely to be secured to the CCD imager itself.

Mounts and Sizes

The term “lens” is used broadly to describe everything from simple single element lenses without holders, to multi-element combinations of lenses with mounting hardware. The lens assembly is a group of from one to ten lenses with the necessary mounting hardware. There really is no limit to the number of elements, but at some point small imperfections reduce the benefits of adding more elements. Some of the mounting hardware is customized, but much of it falls into a few well-known standards. These include the C-Mount category which is common in CCTV jargon and security cameras, as well as the small format 10 mm diameter screw-in-to-focus units for videoconferencing.

A lens mount connects the lens assembly to the PCB on which the CCD imager is mounted. C-Mount lenses are quite common, and their mounts are readily available. Many of the new small format lens assemblies do not have off-the-shelf lens mounts to position the lens assembly in relation to the CCD imagers. This placement should put the center of the lens perpendicular to and on axis with the center of the imaging area of the CCD. In the

case of the small form factor lens assemblies, the lens mount can be easily molded into the camera's enclosure plastics. This is possible because of the low weight and limited torque that might be applied to one of these small lens assemblies. In the case of a C-Mount lens, which is generally mounted outside of the camera enclosure, a more substantial mounting system should be used to securely mount this larger lens format to the camera without unduly stressing the internal PCB material.

Optical Aberrations

Spherical Aberration:

A distortion caused by the light passing through the edge of a lens being focused at a different distance from the lens than light passing through the center of a lens.

Coma:

The difference in magnification between light which passes through the center of a lens (on axis) and the light that passes through the edge of the lens (off axis). This effect may also cause points of off axis light to form *tails* reminiscent of comets.

Astigmatism:

The blurring of off-axis points in either the radial or tangential directions. The image can be re-focused to correct one or the other of these effects, but not both at the same time.

Curvature of Field:

The error caused when all of the points of an image are not in focus in a plane, but rather over a slightly curved surface. Since the CCD image surface is flat (and non-bendable), the image will appear to be in focus at the center, but not at the edges, or vice versa.

Pincushioning and Barrel Distortion:

The error created when magnification varies with respect to off-axis angle. The effect will be mani-

fested when viewing a square object that will appear to have curved sides.

Chromatic Aberration:

The error caused by a variation of the focal length with the wavelength of the light. The effect will be manifested as the classic rainbow halo around objects. This aberration is compensated for by selecting appropriate lens elements which, when taken together, have very little net chromatic aberration.

Lateral Color:

This is the effect seen when the magnification varies with the color of the light. For example, the red image might be slightly larger than the blue image when displayed on a typical CRT.

CCD IMAGER BASICS

CCD Imager Description

The CCD imager is a group of photo-collectors (photo-gates or photo-diodes) which convert light energy (photons) into electrons. These photo-collectors, or pixels, can be arranged in either a line (linear array), or in a grid pattern (area array). The number of electrons which are generated within the photo-collector is proportional to the intensity of the light falling on it. The electrons are accumulated, and then ultimately transferred out of the CCD, one pixel at a time using a large shift register structure. The shift register is integrated into the photo-diode array.

Linear vs Area Array CCD Sensors

CCD imagers are available in a number of classes from several manufacturers. Area array sensors are commonly used in television studios and camcorders. Linear array imagers are commonly used in manufacturing settings where conveyor lines move products across the linear array's field of view to create a two-dimensional image from a one dimensional array. This document will focus more specifically on area array CCD sensors.

Classes of Imagers and Their Sizes

Area array CCD sensors come in a variety of sizes, each designed for a particular end use. The classifications are tied to both horizontal pixel count and vertical line count. Most CCDs fall into one of two vertical line counts, 494, for 525 line, 60 Hz NTSC systems and 576, for 625 line, 50 Hz PAL systems. The 494 line sensors are often classified as NTSC imagers, while the 576 versions are listed as PAL imagers by many manufacturers. Within both the PAL and NTSC categories there are a number of available horizontal resolutions. For traditional broadcast imagers (those closely associated with a television display device) 512, 768, and 720 are common options, while computer-based systems favor 640 horizontal pixel imagers. In all cases, the image aspect ratio is 4:3. A fixed overall image aspect ratio, and a fixed number of horizontal lines means that the individual pixel aspect ratios are dictated by the horizontal pixel count. High horizontal pixel resolution generates skinny and tall pixels, while low horizontal resolution generates long and flat pixels.




Figure 4. Area Array CCD Image Block Diagram.

Interlace vs. Progressive

Area arrays come in two basic forms, interlace scan and progressive scan. Both of these imager types will typically produce one frame or image for every

two cycles of the ac power system (i.e., 30 fps (frames per second) for a 60 Hz power system and 25 fps for a 50 Hz power system). The progressive scan imager processes the area array from top to bottom in one pass to generate its frame. The interlace scan imagers process the area array from top to bottom in two passes, the first for the odd numbered lines and the second for the even numbered lines. The interlace architecture has its origins in broadcast television where interlaced display technology has been used in conjunction with long persistence display phosphors to reduce screen flicker. Modern computers employ high frame rate (greater than 60 fps) progressive scan display systems to handle highly detailed images without screen flicker.

Many of the current videoconferencing specifications are based on low frame rate progressive scan image formats which are compatible with both broadcast television image rates and computer display systems. Conveniently, the field image size (half of an interline-transfer sensor frame) of many broadcast imagers is very closely matched to the image size of the computer based progressive scan videoconferencing standards. By throwing away every other field of an interline-transfer imager, you basically get a progressive scan image format that can be used with both broadcast and computer based systems. This issue is covered more completely in the later sections.

Color Filter Arrays

Typical CCD imagers are sensitive to light through the visible spectrum up to infrared. On their own, CCD imagers generate monochrome images in proportion to the incident light energy across the entire spectral sensitivity range. Color images can be generated by placing narrowband monochromatic filters over the imager and combining the images. Perhaps the most simple approach would be to combine the information from three separate images of the same scene, using red, green, and blue fil-

ters. The color and intensity information of the scene is identical to the familiar RGB image format used in TV and CRT displays.

Taking three separate images of the same scene, using three separate filters, is tedious to say the least. This procedure is greatly simplified in professional video cameras where the light from the lens is split into three equivalent beams using a prism structure. The three views of the image are passed through a filter and then on to their own imager. Red, green, and blue filters are a common choice for TV systems. Other filter selections could be made, but they would require a matching color system downstream for proper image reproduction and display.




Figure 5. Professional Video Three CCD Imager Beam-Splitter Optics.

Typical consumer CCD cameras do not employ the three-imager/prism approach to color imaging because of its complexity and cost. Instead, a different color filter scheme is used with the filters integrated onto the CCD structure itself. The system employs green plus the three complementary colors, magenta, yellow, and cyan. Very small filters are deposited on the CCD imager in a repeating pattern, one filter color per pixel. Unlike the monochrome CCD output which is basically usable as is, the color CCD output is now a combination of coded brightness and color information (typically

termed *Mosaic Color Data*). This brightness and color information must be separated by the color space processor, and converted into one of the many color space representations used in video imaging. The two most common color spaces are the familiar RGB, and the YCrCb (which is very similar to YUV) color space used in the transmission of most broadcast TV signals.

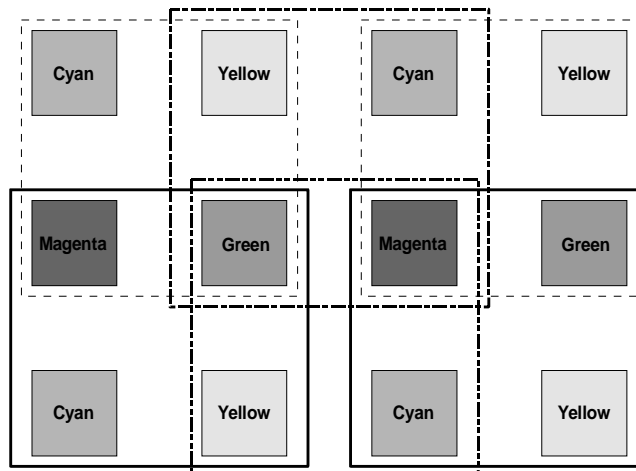


Figure 6. Four Color CCD imagers Employ an Overlapping and Repeating Filter Pattern to Detect Color.

Charge Accumulation

As the individual photons strike the photo-sensitive areas, electrons are created. These electrons fall into a potential well within the material of the CCD imager. The electrons continue accumulating until they are either moved, or until they are drained (erased) away by the electronic exposure reset. The potential wells are created by biasing small regions around the individual pixels such that they are fully depleted of any charge carriers. By varying the bias voltage, the potential well walls can be “dropped,” allowing the well to “move.” The accumulated charge moves with the potential well with very high efficiency and almost no loss of accumulated charge.

Charge Transfer

Once the accumulated charge from the image has reached a reasonable level, the entire image is

moved to the shift register which is built into the CCD imager alongside the individual pixels. Each pixel has its own shift register cell, with the shift register cells typically arranged in columns. For example, if an imager has 512 horizontal pixels, there will be 512 individual vertical shift registers which run from top to bottom on the imager. The number of pixels stored in each vertical shift register will be equal to the number of horizontal lines on the imager (525 or 625 for example).

The charge transfer occurs at the beginning of each field for interlace devices and at the beginning of the frame for progressive scan imagers. Once the charge from the individual pixel has been moved into the shift register structure, the pixel's potential well will have no remaining accumulated charge. This effectively resets the pixel charge to allow the pixels to begin collecting the next image, which will itself soon be "transferred" into the shift register.

Vertical Shift

Once the image has been transferred from the pixels to the vertical shift structure, the shift register contents will need to be clocked out of the CCD. This is accomplished in two steps. The first operation involves moving one complete line of imager data out of the many vertical shift registers and into the high-speed horizontal shift register. Each of the vertical shift registers moves one pixel's charge into the horizontal shift register. The horizontal shift register's length matches the number of vertical shift registers for the entire imager. The vertical shift operation generally occurs during the horizontal blanking period for every horizontal line of pixels that are clocked out of the CCD imager. The vertical shift occurs about every 63.5 μ s, which is the roughly 15.75 kHz horizontal line rate of most cameras and TV systems.

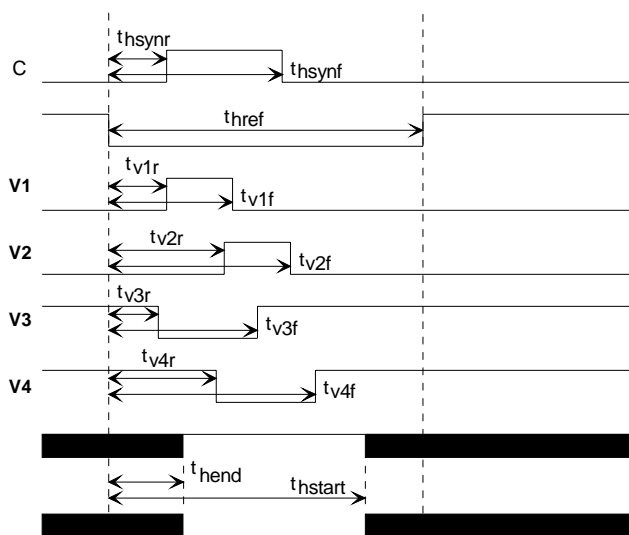


Figure 7. Vertical Transfer Clock Timing.

Horizontal Shift

With the horizontal shift register loaded by the vertical shift operation, the CCD is now ready to clock the pixel values to the output amplifier. One pixel per horizontal shift clock is moved out of the horizontal shift register at the pixel clock rate. The active video portion of most camera systems is 52 μ s per horizontal line. The entire horizontal line will need to be clocked out of the CCD within this interval, which dictates the frequency of the pixel clock. An imager with 512 horizontal pixels will require a pixel clock of just under 10 MHz, while one of the newer ITU.601 compliant imagers with 720 active horizontal pixels will require a pixel clock of 13.5 MHz.

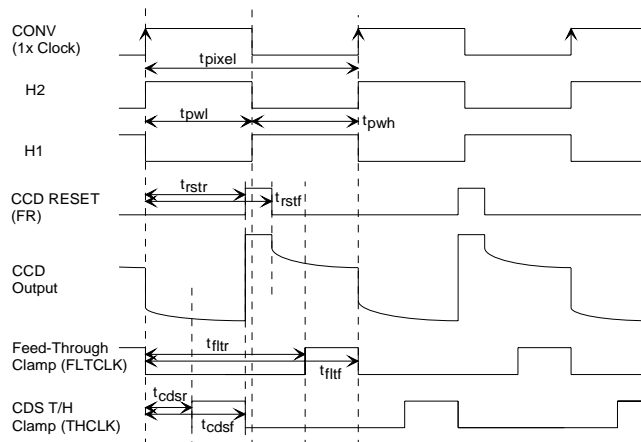


Figure 8. Horizontal Transfer Clock Timing.

Reset and the Output Amplifier

The CCD output amplifier converts the charge, which was clocked out of the horizontal shift register, into a voltage output level. This output voltage is then used by off-chip electronics to generate a usable image with the appropriate synchronization signals.

Any residual charge from the previous pixel can corrupt the amplifier output level. To eliminate this effect, the input of the amplifier is reset before each pixel is shifted in. During this reset period, the output amplifier presents its characteristic reset level, which might be as high as +12 V. Once the reset signal is released, the output amplifier will be at the zero reference, or feedthrough level, for the next pixel. The feedthrough level varies slightly from pixel to pixel, but is generally a little less positive than the reset output level. The pixel charge will then be shifted into the amplifier just after the feedthrough level. The amount of charge collected by that pixel is measured relative to the feedthrough level. The pixel level is generally a little less positive than the feedthrough level.

Track-and-Hold circuits (T/H) are generally used to determine the pixel value. The first T/H is used to hold the feedthrough level, while a second T/H grabs the difference between the held feedthrough level and the imager output amplifier level during the pixel portion. This technique is known as "Correlated Double Sampling" (CDS). The CDS output varies from one CCD type to another, but typical ranges are from a 200 mVpp to 1.6 Vpp for most consumer CCD imagers. Note: in most cases the CDS is performed downstream of the CCD. However, some CCDs provide CDS as an on-chip function.

Timing Signals

A typical CCD imager has several classes of timing signals. These include the vertical shift clocks, the horizontal shift clocks, the output amplifier reset, and the electronic exposure control. The horizontal

shift clocks and the output amplifier reset operate at the pixel clock rate. The vertical shift clocks and the electronic exposure control operate at the horizontal line rate. Many of today's CCD imagers accept CMOS/TTL inputs for the pixel clock rate control signals, however, the vertical shift clocks and the electronic exposure controls tend to require fairly high voltage levels (on the order of 23 Vpp).

Vertical Drive Circuit

Modern CMOS based electronics are not capable of generating the high voltages required for the vertical shift clocks and exposure control signals. This has led to the development of special circuits which switch the required high voltages under the control of standard CMOS/TTL level input controls. These special circuits are known as vertical drive units. Most CCD manufacturers provide a vertical drive chip to be used with their CCD imagers. In general, these vertical drive units are very similar to one another. Some include charge pump circuitry for generating the required bias voltages, while others include circuits that simplify the electronic exposure control.

Electronic Exposure

Once the accumulated charge has been transferred into the vertical shift registers, the pixels begin collecting the next image. In 60 Hz camera systems the image is transferred once every 1/60 of a second (16.6 ms) for interlace scanning transfer devices, and every 1/30 of a second (33.3 ms) for progressive scan imagers. The image exposure times are 1/60 and 1/30 respectively. These times are 1/50 and 1/25 of a second for 50 Hz power systems (20 ms and 40 ms). Exposure time can be extended to improve low-light imaging by acquiring fewer frames per second as is done in slow-scan cameras. Exposure time can be shortened by acquiring more frames per second, but this is not very practical since the pixel clock is already a high-speed clock, commonly in the 10 MHz range.

A more practical means of shortening exposure time is to drain away all of the accumulated pixel charge for some prescribed portion of the time between two successive charge transfers. The interval between the draining of the charge and the charge transfer is the exposure time. If the charge is drained away for a short period of time after the last frame transfer, the pixels will have a relatively long interval to accumulate charge. If the charge is drained away just before the image transfer, the pixels will have a very short interval to accumulate charge and exposure times will be very short.

The voltage on the overflow drain input of most consumer CCDs determines if charge is accumulated, or if it is drained away. This is the high-voltage electronic exposure control mentioned above. Many of the consumer CCDs will drain away all accumulated charge if the overflow drain voltage is taken up over +22 V or so, i.e. at or above the vertical shift clock level. It is fairly straightforward to pulse the overflow drain input once per horizontal line until the number of horizontal lines remaining before the next field transfer matches the desired image exposure time. Pulsing the overflow drain with only one horizontal line interval before the next frame transfer will result in an exposure time of about 1/16,000 of a second. Pulsing the overflow drain with 80 lines at 62 μ s per line would generate an exposure of 1/200, and so on.

Both the electronic exposure control and a mechanical aperture are capable of adjusting the total charge accumulated by the pixels. For particularly well-illuminated pixels, the charge can build up to the point where it spills over into adjacent potential wells. Exposure control, whether mechanical or electronic, is usually set to limit charge accumulation to a level where this type of over exposure is minimized.

In situations where there is a great deal of motion in the scene, the electronic exposure is usually set to a very short interval. This has the affect of “stop-

ping” the motion with minimized blurring of the image. The downside of very short electronic exposure is that there may not be enough ambient light to collect a highly detailed image with sufficient contrast. This can be minimized by opening the mechanical aperture to a wider setting, or using a lens with larger aperture to begin with. For larger lenses, this is not usually a problem; however, many of today’s smaller low-cost lenses may require an additional light source for indoor scenes involving significant motion. Applications like videoconferencing, where there is generally very little motion content, can easily employ these low cost lenses with the longer exposure times and still be free of blurring effects.

Bias Voltages

CCD imagers use potential wells to hold the accumulated charge generated when photons strike the pixel elements. These potential wells are formed by establishing fairly large voltage potential differences around each of the pixel elements. Some of the biasing voltages commonly seen are +15 V combined with -8 V, while other CCD employ +18 and 0 V to form the potential wells. The bias voltages required by a given CCD imager may be quite unique when compared to the supply standards of today +3.3 V, +5 V, or +12 V. DC-to-DC converters are generally employed to generate the necessary CCD bias voltages from the available supply feeds.

In addition to the correct formation of potential wells, the CCD bias voltage can affect pixel conversion efficiency, and noise levels present in the CCD output. The +12 V to +18 V substrate bias voltage will have the largest affect on overall image quality. In many cases all of the other bias voltages can be allowed to vary slightly, while the main +12 V to +18 V may need to be trimmed using a small potentiometer. Fortunately, many of the most recent CCD imagers do not seem to be as sensitive

to slight variations in this bias level, which allows the bias voltage to be fixed with no trimming at all.

PROCESSING THE CCD OUTPUT

Converting the CCD Signal to Digital

The main objective of the analog processing unit is to convert the analog CCD output signal into a digital format. Most four-color consumer CCDs generate a “mosaic” output data stream which encodes both brightness and color pixel data. The analog processing section of the camera uses several track-and-holds, and a number of other circuits, to precondition the CCD output prior to being digitized by the ADC. Once the CCD output signal has been digitized, it will be processed by the “color-space converter” to generate a standardized digital video format.

Shape of the CCD Output

The common CCD output signal is made up of three main levels; these are reset, feedthrough, and the pixel level (see the *Reset and the Output Amplifier* section under *CCD Imager Basics*). The reset level is the most positive, at about +12 V for many devices. The reset is performed to clear out any residual charge from the last pixel, before the next pixel is output from the CCD. Following the reset level, the CCD output will be at the feedthrough level, which is often referred to as the pixel black level reference. The feedthrough level is a few tenths of a volt below the reset output level. The final CCD output level is that of the pixel itself. The pixel level can be from a few hundred millivolts less than the feedthrough level to as much as two volts below. The maximum pixel voltage size is very much dependent on how the CCD imager was designed. The individual pixels will vary from this maximum pixel level (saturated level) to almost no voltage difference from the feedthrough level. Output levels are directly dependent on the intensity of the light incident on the specific pixel.

Correlated Double Sampling

Correlated Double Sampling (CDS) is based on using two successive track-and-holds (T/H) to reduce the low frequency noise inherent in standard CCD imager outputs. The individual pixel output level is defined relative to the CCD feedthrough level. The feedthrough level is roughly equivalent to a standard pixel output level with no light incident on the pixel itself. Although the amount of accumulated charge on a *dark* pixel is very nearly zero, several leakage and noise sources contribute a finite charge accumulation even under dark conditions. Theoretically the feedthrough or dark level is the same for all pixels, but in practice there is a statistically significant level of variation in the feedthrough level across a CCD imager. Additionally, low frequency noise sources affect pixel levels and feedthrough levels equivalently by introducing very small voltage offsets in the CCD output. By storing the feedthrough level in a track-and-hold, the pixel level can be subtracted from the held feedthrough level for a very accurate measurement of the pixel. This difference is typically held in a second track-and-hold, which is the basis of the double sampling in the CDS. Some ADC topologies do not require the second T/H to maintain the difference signal over an extended period, and will simply eliminate the second T/H to allow the ADC to perform this function directly.

Pre-Amp and Gain Settings

The analog processing front-end for a digital CCD camera might include a buffer stage to isolate the CCD output from any analog input anomalies of the analog processor, the most common one being high input capacitance. If the input capacitance exceeds the drive capacity of the CCD imager output, settling time will suffer. To combat this problem, an emitter-follower type amplifier stage is often used to increase drive capacity at the expense of a very slight reduction in output level (simple followers typically operate with slightly less than uni-

ty gain). Modern high-speed electronics based on very fine line geometry's present a very small input capacitance (<25 pF) to the CCD imager output, and may not need this follower stage. These same devices are not generally able to withstand a +12 V analog input level, and will require a dc blocking capacitor to allow the high-speed CCD output signal to pass, but not the high-voltage dc bias level.

Matching LSB's and CCD Noise (SNR)

The CCD imager output voltage range needs to be matched via a gain stage to the analog processor's ADC input range. The two key issues which affect the correct gain setting in a digital camera are maximizing the portion of the ADC input range used to digitize the image, and setting the gain such that the CCD output noise dominates the image, not ADC quantization noise. If the CCD output voltage level is larger than the ADC input range, the input signal should be attenuated enough such that the input signal is inside the ADC input range. If the CCD output signal is smaller than the ADC input range, the CCD signal should be amplified until the CCD noise level is greater than the LSB (Least Significant Bit) of the ADC while not exceeding the overall ADC input range; this guarantees that the CCD noise dominates the image rather than the ADC noise. When the CCD noise dominates, the ADC injects very little additional noise into the image capture process.

Automatic Gain Control

A variable gain amplifier (VGA) is often included between the two CDS T/Hs. Gain applied to the feedthrough level minus the pixel level will scale the apparent pixel intensity. Because the difference signal is likely to be ground referenced (a result of the CDS operation), the variable gain amplifier is able to operate in a more linear region of its range when compared with applying gain to the +12 V referenced CCD output signal directly. An Automatic Gain Control (AGC) circuit is formed when the VGA is controlled in a closed-loop manner

based on the image brightness. The loop can be closed in hardware for rapid response, or in software using camera control ports. AGC loops are generally operated with very low loop bandwidths (well below Frame rate 60 Hz) to avoid visible image artifacts as the loop adapts to its desired value.

Since many image compression techniques are based on frame differencing algorithms, a continuously operating AGC loop may not be desirable. The frame-to-frame intensity difference caused by slightly different AGC settings will often be seen as motion by compression algorithms, even if the image content has not changed between the frames. In image compression-centric applications like video-conferencing, the AGC should be used on a periodic basis to keep the image properly exposed, but to avoid detecting false AGC-based motion. For example, a system might enable the AGC loop initially to adjust for room lighting. Once set, most exposure settings can be held at their current values for long periods, with only sporadic updates, while maintaining normal image exposure and quality.

Black Level Clamping

Black level clamping is a process which has its origins in composite analog video, where it was used to separate the synchronization signal from the video portion of the signal. Two small sections of the composite horizontal line, one before and one after the active video, were defined as the *black* video levels. Signals greater than the black levels would be interpreted as brighter image data, while levels less than the black level would be interpreted as synchronization clocking. The simplest means of separating these two signals is to clamp the designated black sections of the horizontal line to ground (0 V), which leaves video above ground and sync data below ground.

Digital video is transferred as numbers within a range ... 0 to 255 for 8-bit systems, or 0 to 1023 for 10-bit systems. Black level clamping establishes the image black level several codes above code 0 to

allow for a few codes to be used to encode synchronization information. For example, defining digital code 16 as the black level allows codes 0 to 15 to be used for synchronization and timing data. The ITU 601 recommendation for digital component video and the ITU 656 recommendation for digital video transport, both define code 16 as black. Codes greater than 16 represent video information.

Black clamping inside a digital camera can be performed in either analog or digital circuits. If done using analog processing, the black level clamp adds an offset to the output of the CDS such that a dark pixel generates code 16 (if code 16 is used to define the black level) at the output of the ADC. If the black clamping is done in the digital processor, the black code offset is added to the image luminance data. Adding an analog offset after the CDS is likely to reduce the ADC dynamic range, while adding a digital offset might cause some clipping of brighter image details. Since digital cameras do not derive their synchronization queues from the mosaic output data, black level clamping is necessary only to simplify compliance with various digital video standards.

Clock Sequence Overview

Processing the image data from a CCD imager is done to the repeating beat of system clocks. The clocks are separated into a number of groups, each with a given purpose. Three groups of clocks are closely associated with the CCD imager, and one clock group is tightly coupled to the analog processing. Most CCD cameras derive all of these clock groups from one timing generator. Many older cameras have been built around imager-specific timing generators which cannot be easily interchanged. Modern CCD cameras will employ programmable timing generators which can be easily reconfigured for different CCD imagers.

The three CCD clock groups are the frame-transfer, the vertical transfer, and the horizontal transfer

groups. The frame transfer clocks will cycle at the image frame or field rate depending on whether the imager is progressive or interlace transfer. For common consumer CCD imagers, the frame transfer clocks are often designated as VH1 and VH3. The vertical transfer clocks cycle at the horizontal line rate of around 15 kHz; the vertical transfer clocks are often designated as V1, V2, V3, and V4. The second group of CCD imager clocks are the horizontal transfer clocks which include the actual horizontal clocks, H1 and H2, as well as the CCD imager output amplifier reset clock. All three of the horizontal rate clocks will cycle at the pixel rate, which ranges from 1 MHz to 13.5 MHz.

The processing group of clocks includes the feedthrough clock, the T/H clock for the CDS, the ADC convert clock, and the black level clamping clock. The first three clocks operate at the pixel rate, and move the CCD imager output through to the ADC. The black clamp clock operates at the horizontal line rate, and is only used if the black level offset is added before the ADC.

CCD cameras with analog color processing employ a number of additional clocks to drive the color separation and processing operations. Digital CCD cameras will not require these extra clocks. This means that the two or three extra crystal oscillators (required for analog color processors) will not be needed on a synchronous digital CCD camera. These extra clocks often corrupt the CCD output signal, and force the camera designer to employ clock skewing or tweaking to place the critical clock edges away from system anomalies. Cameras employing digital color processing will tend to only need one pixel-rate clock, from which all other processing proceeds. These synchronous systems are not susceptible to the interference issues previously seen with analog processed CCD cameras. Additionally, digital color processed CCD cameras do not need the critical clock edges moved for proper operation.

Digital Limiter Formatter

The output of the analog processor is the mosaic data stream. This data will contain the image data, and the timing and queuing information. The digital limiter is designed to keep the image portions of the video stream within their predefined code limits, as well as to add the timing information into the stream if necessary. The limiter/output formatter will need a matching decoder unit inside the digital processor to separate these streams once they are inside the digital processor.

COLOR SPACE PROCESSING

The color space processor in a digital camera converts the mosaic data stream from the analog processor, into a standards-compliant digital component video stream. The most common output format is the YCrCb 4:2:2 specified in the ITU-601 recommendation. There are several variants covered by the ITU-601 recommendation, including an RGB version (this is a 4:4:4 format). In addition to generating this standards-compliant data stream from the analog mosaic data, the digital unit might also perform several image processing functions to improve, or modify the digital image. These can include white balance, gamma correction, saturation control, etc..

Color Space

A color space is a three axis system used to describe a given hue, saturation and intensity. The human eye is generally capable of seeing color from the violet wavelengths of 400 nm to the red wavelengths of 760 nm. In 1931 the Commission Internationale de L'Eclairage (CIE) established a set of standard curves to specify how a given color is translated into three components. These *color spaces* perform much the same function as Cartesian and polar coordinate systems perform in physical space. Some examples of common color spaces are YUV, YCrCb, RGB, YIC, as well as the mosaic color space used by four-color CCD imagers.

Color Space Conversions

Moving from one color space to another is a matter of multiplying the original image data values by a set of conversion matrix coefficients. For a color described by two color spaces, the conversion is lossless and reversible. Some colors might be described by one color space system, but are not within the description range of the target color space; these colors can not be simply converted back and forth. For example, there are several colors clearly described by YCrCb, but not correctly described by an RGB color space system. Fortunately, the range of YCrCb, mosaic, and RGB color spaces overlap each other very closely, and few colors fall into the fringe area. This means that a camera using a typical four-color mosaic-based CCD imager can easily use YCrCb (often called YUV) and RGB color-spaces without perceivably affecting the image quality.

Color space conversion coefficients are usually described in terms of the conversion to a pure RGB color space. A conversion matrix can be transposed for the reverse operation. When moving from one color space to another (neither being RGB), the various coefficients can be multiplied together to create a new conversion matrix relating the two color spaces directly. The ITU-601 specified R'G'B' to Y'CrCb color space conversion coefficients are listed below (the primes denote gamma corrected signals to be discussed later).

$$\begin{matrix} Y' \\ Cb \\ Cr \end{matrix} = \begin{bmatrix} 16 \\ 128 \\ 128 \end{bmatrix} + \begin{bmatrix} \frac{77}{256} & \frac{150}{256} & \frac{29}{256} \\ \frac{-44}{256} & \frac{-87}{256} & \frac{131}{256} \\ \frac{131}{256} & \frac{-110}{256} & \frac{-21}{256} \end{bmatrix} \times \begin{bmatrix} R' \\ G' \\ B' \end{bmatrix}$$

Note: R',G',B' values are limited between 16 and 235

Separating the Mosaic Imager Data

Most CCD imager output data is in the form of vertically aligned pixel pairs, Magenta+Yellow,

Green+Cyan, etc. (see Figure 9). Within the mosaic data stream is coded the brightness and red/blue color difference content. The process of converting the raw mosaic data into the individual brightness and color difference signals is called color separation. This color separated intermediate mosaic brightness and color difference information should not be confused with the Luminance (Y) and Chrominance (Cr and Cb) of the YCrCb color space. The equations below illustrate the relationship between the individual pixel groupings and the RGB color space components.

Mg	G	Mg	G
Ye	Cy	Ye	Cy
G	Mg	G	Mg
Ye	Cy	Ye	Cy
Mg	G	Mg	G

Figure 9. Typical Mosaic Filter Organization

Red Line Components:

$$\text{Brightness} = (\text{Mg} + \text{Ye}) + (\text{G} + \text{Cy}) = 3\text{G} + 2\text{R} + 2\text{B}$$

$$\text{Red Color Difference} = (\text{Mg} + \text{Ye}) - (\text{G} + \text{Cy}) = 2\text{R} - \text{G}$$

Blue Line Components:

$$\text{Brightness} = (\text{Ye} + \text{G}) + (\text{Cy} + \text{Mg}) = 3\text{G} + 2\text{R} + 2\text{B}$$

$$\text{Blue Color Difference} = -(\text{Ye} + \text{G}) + (\text{Cy} + \text{Mg}) = 2\text{B} - \text{G}$$

The color separated mosaic data is converted into either RGB color space or YCrCb color space using a matrixing operation. The matrix multiply operation is simply the three color separated mosaic components multiplied by a 3x3 matrix of color separation/conversion coefficients. These coefficients convert the mosaic data into RGB or YCrCb color space, but also adjust for the specific color filter characteristics of the CCD imager. Each manufacturer of CCD imagers employs their own color filter manufacturing process, which results in slight variations in the precise filtering characteristics of the mosaic pattern. This mosaic variation is corrected in the color separation/conversion matrix co-

efficients. Once established, these color separation/conversion coefficients should remain constant for a given color filter process and pixel sensitivity curve.

Gamma Correction

The light emitted from a Cathode Ray Tube (CRT) is not linearly proportional to the input voltage level. Gamma correction is the process of correcting the CRT input voltage such that the image produced is linear in color and intensity. Basic gamma correction is a power-law function applied to the image data. Without gamma correction the dark areas of an image will appear too dark.

The intensity of the CRT display is approximately the input voltage raised to the 2.5 power. The exponent value (2.5 in this case) is generally known as gamma. The correction function applies an “inverse” 1/2.5 power function to create a gamma corrected signal to be applied to the CRT.

$$\text{CRT Display Intensity} = (\text{Input Level})^{1/2.5}$$

The RGB and YCrCb type image data streams are generally gamma corrected in the camera before being sent on to the computer or monitor. Once RGB data is gamma corrected, it is termed R'G'B' data with primes to indicate the non-linear gamma processing has been performed. Similarly, YCrCb is termed Y'CrCb once gamma correction has been performed. A video stream should only be gamma corrected once. A signal that has been gamma corrected twice will have mid-tones that appear to be too light. Since a gamma corrected signal is not easily processed, the gamma correction should be reversed before any additional processing is performed.

In practice, the exponent of 2.5 yields an image which appears to have poor contrast in dimly lit environments. Most broadcast systems compensate for this effect by assuming an exponent value of 2.2 when performing gamma correction. The actual

gamma correction equation for each of the image components is:

$$CRT\ Display\ Intensity = (Input\ Level)^{1/2.2}$$

or

$$CRT\ Display\ Intensity = (Input\ Level)^{0.45}$$

Gamma correction is reversed using an inverted exponent. For example, reversing gamma correction on the gamma corrected R' red component is as follows:

$$R = (R')^{2.2} = (R')^{1/0.45}$$

The ITU BT.709 recommendation specifies a slightly more complicated gamma correction methodology. In the BT.709 approach input levels less than 1.8% of full scale are corrected with a straight line approximation of the gamma power function. Over this very limited input range, the input is multiplied by 4.5 rather than being raised to the power of (0.45).

Color Aliasing

Because the brightness and color information is derived from a number of pixels, it is possible for some types of scenes to distort the color separation process. The classic example of this distortion process occurs when only half of a given mosaic block of pixels is illuminated. Under these conditions, the brightness will be correctly calculated, but the calculated color difference signals will often be quite skewed. Since the color separation relies on accurate color difference inputs, the overall calculated color will be wrong; this is known as color aliasing.

Aliasing effects which occur at major transitions from dark to light (where mosaic block are commonly only half-illuminated) are not very visible to the eye. Repeating patterns of dark and light scene areas are capable of generating repeating half-illuminated mosaic blocks which gives rise to larger areas with aliased color. This effect is especially noticeable when the dark and light areas are closely spaced and span an integer number of mosaic

blocks between transitions. Striped shirts with a large content of vertical dark to light to dark transitions are very effective at demonstrating the effect. Simply move away from the camera until the spatial density of the shirt's alternating pattern approaches the mosaic block density on the imager and the shirt will take on a shimmering "oil slick" type look.

Color aliasing is fundamental to all single imager cameras which use a color dot filtering block to define both brightness and color. The three basic compensations for single imager color aliasing are to (a) slightly defocus the camera, (b) place a blurring filter in front of the imager, or (c) filter the chrominance signal inside the camera. Both the defocusing and blurring filter options reduce image sharpness when used alone. The chrominance filtering can severely reduce the image's color saturation, which makes the scene look gray and flat.

A combination of anti-aliasing schemes yields some of the most pleasing results. By selectively filtering the chrominance data when aliasing is detected, the image's saturation is unaffected in areas of broad color where aliasing is not a problem. Additionally, the image can often be improved by using a luminance peaking filter which sharpens the edges within the image. Basic chrominance filtering and color decimation can soften the image; the luminance peaking filter helps to return the image back to its proper appearance.

Chroma Filtering

In addition to the chrominance filtering used to combat color aliasing, further chrominance filtering may be needed to reduce "color noise." Color noise is a low level variation in the pixel-to-pixel color over an area of the scene. The human eye is very effective at integrating these color variations; an area of the scene such as the background might have a very high pixel-to-pixel color noise content and yet still look like one solid color to the person viewing it.

Color noise is an acute problem if the image is to be compressed for transmission to another location. Most video codecs (coder-decoder) employ some level of motion detection which only transmits portions of the scene which have changed from the last image. This dramatically improves the transmission efficiency by not re-sending the unchanging redundant data. Areas of the image with high color noise will appear as motion to most video codecs. This means that the static background in a video-conferencing application will constantly be sent to the receiver. This reduces the portion of the available bandwidth that could otherwise be used for sending data, voice, or improved image quality. Appropriate filtering within the camera can reduce color noise below the motion detect threshold of the video codec.

Luma Filtering

The image can often be improved by using a luminance peaking filter which sharpens the edges within the image. The fact that four mosaic block pixels are needed to determine each brightness value causes the luminance data to be somewhat filtered. The effect is slight, but when combined with chrominance filtering and color decimation the overall effect can be to soften the image. The luminance peaking filter helps to return the image back to its proper appearance.

Saturation Control

The saturation of a color describes its vividness. Images with very low saturation appear very gray and flat. Images with high saturation seem to have too much color content. The amount of color within an image can be controlled by adjusting the chrominance data gain within the camera. Saturation control is usually affected in YCrCb color space by adjusting the red (Cr) and blue (Cb) gain coefficients.

White Balance Control

All light sources can be described by their temperature or color quality. The camera's light source affects the image's apparent color. The White Balance control adjusts the camera image such that white objects appear white when displayed on a monitor. The white balance is accomplished by adjusting the red and blue color components such that the red and blue chrominance (Cr and Cb) components are both minimized. The red and blue color components are typically adjusted in RGB color space.

Final Output Formatting

The camera image must be properly formatted before being output. Internally the camera is maintaining three separate data streams; typically red, green, and blue or Y, Cr, and Cb depending on the particular color space. For a full RGB output, these three data streams simply need to be output in parallel. The standard 4:2:2 YCrCb component digital video involves removing some of the picture information, and then interleaving the data into a single stream. Additionally, video timing queues must be inserted into the image data to provide a means of synchronizing the display monitor with the camera image. These timing queues effectively identify the top of a frame or field, and the start and end of each active video line (SAV and EAV).

The ITU-601/656 Recommendations

The International Telecommunications Union is a worldwide body with a charter to encourage standardization for all aspects of telecommunications. The ITU-601 recommendation describes a component digital video format which is compatible with NTSC and PAL systems. A number of digital video color formats are specified, including 4:2:2 YCrCb, and both RGB and YCrCb 4:4:4 formats. All 601 formats include embedded timing queues for EAV and SAV. The ITU-601 recommendation does not discuss physical transport or connector issues.

The ITU-656 recommendation describes the physical transport and connectors to be used to move ITU-601 images from one place to another. The ITU-656 recommendation is basically an 8-bit standard, with two optional bits reserved for compatibility with the 10-bit standard. Many of today's digital video components list 656 compatibility which implies that the embedded EAV and SAV signals are present in the video data stream. This means that a 656-compliant receiver does not need external HSYNC and VSYNC timing signals since it decodes the timing information from the video data stream itself. There are a number of "656-like" components which accept the interleaved YCrCb data stream, but that do not decode the timing data. These "656-like" devices will require the external HSYNC and VSYNC timing signals for proper operation.

Square vs. Rectangular Pixels

Pixel shape is a function of image line count and the number of pixels per line. For broadcast television the line count is generally fixed at 525 or 625 depending on the frequency of the country's power system. The number of pixels per line is dependent on the system pixel clock since each line is roughly 62 μ sec long (52 μ sec active video and 10 μ sec timing information). The pixel clock rate is calculated as $1/(52 \mu\text{sec}) \times \text{number of active pixels}$.

For broadcast television, the number of pixels can vary from 352 active pixels up to 768 active pixels. At 352 pixels per line, individual pixels are short and flat; at 768 active pixels per line, the pixels are slightly taller than they are wide. At 640 active pixel (with 480 active lines on a 525 total line system) the pixels are square, equally wide as they are tall. Computer monitors are almost universally based on square pixel organizations.

For a digital video image to be displayed without distortion, it should be displayed at the same clock rate and number of horizontal pixels at which it was created. A 10 MHz, 512 pixel/line image will seem

too narrow when displayed by a 12.24 MHz 640 pixel/line computer monitor. Video scalers are available to convert from one pixel rate and image size to another. Once correctly scaled for the display, the image will have the correct aspect ratio.

POWER SUPPLY DESIGN

CCD Power Requirements

Commercial CCD imagers require several unique power supply voltages for proper operation. Generally there will be a +15 V to +18 V (at roughly 5 mA) supply for basic CCD biasing, electronic exposure control, and output buffers. CCD imagers also require a negative voltage in the 0 V to -9 V range at roughly 3 mA. There will also be a mid-level voltage in the -1 V to +4 V range at minimal current draw. Since several unique voltages are required, cameras typically employ on-board switching power supplies which are fed from one or two standard voltages.

Synchronized vs Filtered switching supplies

Any supply perturbations during the active video line stand a good chance of being visible in the image. The two basic approaches to avoid visible power supply artifacts are (1) to restrict supply perturbations to occur only during the blanking periods, or (2) to filter the supplies to the point where the perturbations are no longer visible. This allows the clocking to occur at any point within the horizontal line, and perhaps to be asynchronous without creating any visible artifacts.

The restricted clocking approach assumes that the power supply droop which occurs during the horizontal line is not visible in the image. This is a fair assumption since the human eye is much more sensitive to sharp perturbations in the image than it is to slowly changing variations like those caused by supply droop. Each approach has its benefits and its limitations; the restricted switching versions require more design effort, while the heavily filtered designs increase the BOM cost.

Clock Rate Overview

The three basic clock periods within a CCD camera are the 50/60 Hz frame or field rate, the 15 kHz horizontal line rate, and the 1 MHz to 13.5 MHz pixel rate. The field/frame rate is too slow for effective power and voltage conversion. The high-speed nature of the pixel clock is too fast for most switching power supply circuits to be effective. This leaves the 15 kHz horizontal line rate as the most workable clock source if a synchronized switching power supply is used.

Horizontal Sync Driven DC-DC Converter

The synchronized dc-dc power supply must restrict all switching events to within the horizontal blanking period (HREF low) which is 10 μ sec in duration and occurs at the 15 kHz horizontal line rate. The 16% HREF duty cycle means that an entire line's worth of power must be transferred within the short 10 μ sec period, while the load continues to drain this energy during the remainder of the horizontal line. Peak power transfer during HREF is at least six times the average power consumption; this ratio is even larger when power supply efficiency is considered. NOTE: Many integrated synchronized dc-dc converter are PWM based and do not restrict all switching events within the HREF period.

DC-DC Converter Architectures

Two of the more common cost-effective dc-dc converters implementations are the charge pump and the buck-boost designs. A simple charge pump can generate the required +15 V supply from a +12 V source with relatively good efficiency and low cost. The buck-boost design is needed to efficiently generate the required +15 V supply from a +5 V or +3.3 V source.

Efficiency

Most power conditioning circuits (and their inductors) are more efficient when switching in the hun-

dreds of kilohertz range. This type of clock is not readily available in most cameras, which greatly undermines dc-dc power supply efficiency. A +12 V to +15 V charge pump design is only about 60% efficient when clocked by HCLK.

Efficiency and architecture options might be improved by using a higher speed clock. This type of clock can be generated by dividing down from the pixel clock. If the division results in a simple integer number of cycles per line, the switching events will be fully synchronized to the image, i.e. the switching events will not appear to walk relative to the horizontal line (walking will tend to create Herringbone patterns in the image). Because the more frequent perturbations are smaller in magnitude than the HREF clocked counterparts, less power supply filtering is needed. This approach would offer a compromise between poor efficiency due to clock rate, and expensive filtering.

The standard HREF clocked circuits might be improved by gating a higher speed clock with HREF to contain all switching during the blanking period. In this case very little power supply filtering should be needed, while enjoying the benefits of higher speed clocking.

Component Power Ratings

Component power ratings have become especially important with the increased use of small-format surface mount components. Table 2 indicates the basic power handling capability of several common surface mount resistors.

Surface Mount Size	Power Rating
1210	1/8 W
805	1/10 W
603	1/12 W

Table 2. Power Ratings of Common Resistor Packages

Key power supply inductors are additional component where power ratings can be easily exceeded. The relatively low-frequency HCLK-type clocks with low duty cycles used in a buck-boost voltage

converter allow currents to build into the 100 mA to 300 mA range, possibly exceeding the inductor's power rating. Additionally, the buck-boost architecture can generate very high voltages across the inductor (sometimes in the 50 V+ range). The traditional $P_{RMS} = I \times V_{RMS}$ power calculation should be used to calculate inductor power. Inductor current ratings generally define the maximum flux permissible without saturating the inductor.

Component Size and Cost

The need for smaller and smaller cameras continues to drive camera design methodology. The very-large-scale-integration offered by today's advanced camera chip-sets goes a long way toward reducing overall camera size and cost. Many of the discrete components still present inside the camera are also being reduced in size. Handling 1206 and 805 format surface mount components is relatively easy, both in terms of physical size and the fact that the component value is generally printed on the device. The 603 format surface mount devices offer a substantially smaller package (also lower power handling capability, see above) but usually at a slight price premium. The smaller size is somewhat more difficult to deal with, but the real problem with 603 components is that the device value is not generally printed on the package itself. This makes the initial design more tedious, and perhaps opens small scale manufacturing up to more construction defects.

Determining the System Power Budget

Determining an overall camera power budget is complicated by the unique CCD related power supplies that tend to be generated on-board the camera. Sub 50% conversion efficiency increases the impact of these supplies on the overall camera power drain. The +15 V type supply is perhaps the worst of these supplies and should be given the most attention. The first step is to estimate all potential drains on the +15 V supply. HREF-and HSYNC-driven dc-dc converters are especially prone to sup-

ply droop. Any potential supply droop should be carefully considered in terms of overall camera operation. The required dc-dc converter input current is calculated while taking the dc-dc converter's efficiency and duty cycle into consideration.

Generating a substantial amount of +15 V supply current (especially from a +5 V or a +3.3 V supply) requires a great deal of input power. If the input supply current capacity is at its limit, the load on the generated +15 V may need to be reduced. Certainly the CCD burden cannot be reduced, but the CCD output buffer amplifier could be designed for greater efficiency, and a few other changes made to help bring the input supply current into line with system goals. The best way to address the dc-dc converter power efficiency is to eliminate the need for the converters by employing one of the "+5 V only" CCD imagers which should be available in the near future. These single-supply CCD imagers should not be confused with the new CMOS imagers.

DIGITAL VIDEO PROCESSING

Video Encoders

Once an image has been converted to a digital video format, it must be re-encoded in some fashion to be viewed and used. This task is performed by video encoders which convert the digital video stream into a monitor-ready signal including timing information. There are numerous analog encoders available which convert component digital video into PAL and NTSC compatible analog output signals for immediate display. Additionally, the component digital video can be digitally "encoded" inside a computer for display within the graphics subsystem.

Analog Video Encoders

Analog encoders such as those available from Crystal, convert the component digital video into analog composite S-Video or RGB/YUV. This is the familiar video signal present at the back of most

televisions and VCRs. The composite video is decoded inside the receiver and converted into luminance and chrominance signals. The S-Video variant is already in the luminance/chrominance format, and does not require the initial decode. The television may perform some image processing such as color, brightness, or tint adjustment before ultimately sending RGB drive signals to the display tube.

Computer Graphics Adapter Encoders

Although graphics adapters for computer systems employ slightly different internal processing than television based systems, they take component digital video inputs directly. The majority of graphics adapter chips (including all of the Cirrus Logic line of parts) accept YUV type input data through the feature connector. The older designs take 16-bit 4:2:2 formatted video data plus external synchronization, while the newer designs add 8-bit interleaved YUV video data input with embedded EAV and SAV synchronization, as specified in the ITU 656 recommendation.

Embedded Sync vs. non-embedded sync. Format

Many of the older encoders and graphics adapter chips employed the non-embedded approach which provided video data through an 8-bit or 16-bit channel, while keeping synchronization signals like HSYNC and VSYNC on separate lines. The newest devices have been upgraded to accept com-

ponent digital video over an 8-bit channel with embedded EAV and SAV timing cues as specified in the ITU 656 recommendation. Many encoders and graphics adapter chips of slightly less recent design can be found with a combination of these features. Some allow for the interleaved video data but do not decode EAV and SAV timing data. These chips will require slightly greater effort to work with than those designs which are fully ITU 656 compliant.

ZV-Port Standard

The ZV-Port (*Zoom Video*) specification has been incorporated into many of the portable computers built after mid-1996. This specification defines a means of delivering digital component video into the graphics adapter via the PC-Card slots on the portable computer. Once the computer has determined that a ZV-Port enable source is present, the internal host adapter chip re-routes the data stream to the graphics adapter. The video data is then placed in a window within the display area without any further intervention from the host CPU. Once in the graphics adapter memory buffer area, the video data can be moved via the PCI bus to the host CPU for compression or storage.

The ZV-Port specification employs a 16-bit wide non-interleaved YUV type input channel with external HSYNC and VSYNC signals. There has been some discussion of incorporating fully compliant ITU-656 EAV and SAV decoding in a future version of ZV-Port.

COMPRESSION STANDARDS

DCT Based Compression

Wavelet Based Compression

VIDEOCONFERENCING STANDARDS

ITU H.320

ITU H.324

ITU H.323

OTHER INTERCONNECT OPTIONS

P1394

Parallel Port

Serial Port

Universal Serial Bus

Still Image Photography

A CAMERA DESIGN EXAMPLE

Lens

CCD and Vertical Drive Section

Design Partitioning

Analog and Timing Generator Programming

Digital and Color Separation Coefficients

The DC-DC Converter

Audio and Microphones

Connectors

Shielding

Final Camera Timing

Bill of Materials and Costs

• Notes •

SMART
Analog™