

Ay 105 Lab Experiment #5: CCD Characteristics and Operation

This week you'll study the characteristics of Charge Coupled Device (CCD) imaging array detectors. These detectors are the most common type used in astronomy for optical wavelengths. In particular, you will learn about CCD quantum efficiency (QE) and the solid state architecture of CCDs in Part A, and in Part B you will investigate the electronic circuitry that operates a CCD. There are two sub-parts to Part A, and one group can work on the experimental part while the other works on the computer. You can then join forces with the other group to work on the electronics. Part B can be continued into the second lab, when we will operate an example set of readout electronics of a CCD detector.

A. Insights into CCD Quantum Efficiency

CCD quantum efficiency is usually measured with a setup including a QTH lamp, interference filters or a monochromator, lenses or an integrating sphere, and a calibrated photodiode reference detector. Based on week one's experience in the Ay105, you could probably create such a test setup and use it to measure QE as a function of wavelength. After several hours spent making CCD exposures and photodiode current measurements, you might derive a graph like one of those shown in Figure 1.

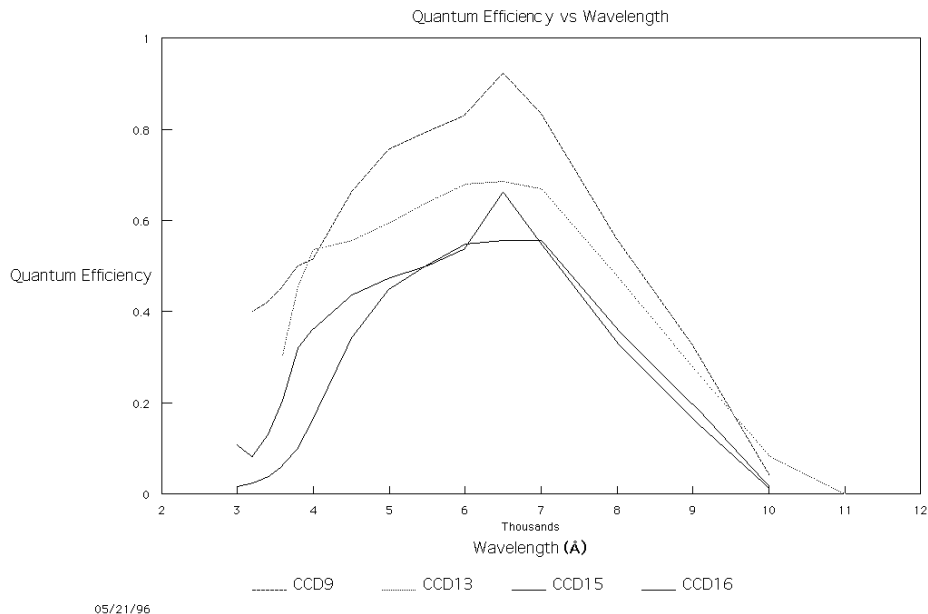


Figure 1: QE vs. λ for Palomar CCDs.

In this lab, however, you will be investigating a slightly different and more interesting question: *why* do CCDs have the quantum efficiency variations with wavelength

shown in Figure 1? In particular, take special note of three characteristics of Figure 1 (before and during lab, as well as in the discussion section of your report). (1) The peak response at wavelengths $5500 \text{ \AA} \lesssim \lambda \lesssim 6000 \text{ \AA}$. (2) The decrease in QE for wavelengths $\lambda > 6500 \text{ \AA}$. (3) The decrease in QE (and the wide variation from device to device) at wavelengths $\lambda < 5000 \text{ \AA}$.

Configuration. The instructor will give you a Texas Instruments 800×800 , three-phase thinned CCD glued between a pair of fused quartz plates. Take a moment to examine the CCD (both front and back), visually and through the low-power microscope. Describe some of what you see (and find most interesting) in your lab notebook. Then install the CCD sandwich in the filter holder, and place it on a rotating stage. Arrange the green HeNe ($\lambda = 5348 \text{ \AA}$) laser to shine on the CCD array.

The experimental objective here will be to measure the fractions of incident laser light that are transmitted and reflected upon striking the CCD array. As mentioned during lecture, “thinned” CCDs are only about $8 - 10 \mu\text{m}$ thick; the Si back surface also appears shiny, so it is not unreasonable to suppose that some photons reflect off of the surface and some photons go right through without being absorbed. Of course, photons in both of these categories stand no chance of being detected. The remaining fraction is thus an upper limit on the quantum efficiency of the CCD:

$$QE(\lambda) \leq 1 - R(\lambda) - T(\lambda),$$

and so your first task will be to determine how closely the actual CCDs in Figure 1 approach the ideal case of detecting all photons which are not reflected or transmitted.

Method. You’ll make an optical path to collect both the light transmitted straight through the array, plus another optical lens assembly off-axis to reimage the reflected light onto the photodiode. Start with the laser on the left side of the table, illuminating the CCD in the middle of the bench; you will want to use the iris between the laser and CCD, stopped down, to block extraneous laser light from striking the CCD. Place an optical assembly to collect reflected light back from the CCD at about 135 degrees from the laser path. You’ll need to rotate the CCD slightly away from the laser beam’s normal incidence in order to collect the reflected beam; use the 63mm diameter, 356 mm EFL achromatic lens to collect and collimate the reflected light over some non-zero solid angle. (Notice that the $10 \mu\text{m}$ thick CCD “membrane” is not a precisely flat surface, so the reflected beam will in general not be very well collimated). Then use the 80mm $f/2.8$ triplet lens to focus the collimated light, forming an image of the point on the CCD from which the light was reflected. Locating this image and placing your photodiode at this point will enable you to measure the amount of laser power reflected by the CCD. Determine whether or not you are likely to collect all of the reflected light by calculating the ratio of the image size (on the photodiode) to the object size (the reflecting spot on the CCD), then estimate the latter and derive an estimate of the former for comparison with the 1 mm^2 photodiode area. Record the photodiode current generated by the laser light reflecting from the CCD. (Note that the laser intensity varies, so monitor the photodiode current for a while and estimate

a likely uncertainty in all of your measurements.)

Next, slide the CCD sandwich within the filter holder so that the laser strikes only the clear fused quartz plates, and measure the power reflected by the plates alone. Finally, substitute the diagonal mirror's post for the CCD's filter holder post and align the mirror so that the light strikes the photodiode, and measure the power reflected by the aluminum mirror. Bare Aluminum has a reflectivity of 89% (could you determine this yourself, and if so, how accurately?), so from this you can determine the power in the laser beam incident on the CCD sandwich. (Why is using the mirror a better approach than moving the photodiode so that it intercepts the incident laser beam directly?) Subtracting from this the fraction reflected by the cover glass alone will enable you to determine the fraction that was reflected by the CCD itself. This is the fraction $R(\lambda)$ in equation (1) at the wavelength of the green laser. Now substitute the red laser for the green laser, and repeat these measurements for $R(\lambda)$ at the wavelength of the red laser.

Next, use either the 50mm or 180mm EFL Nikon lens on-axis behind the CCD to focus the transmitted light, and place the photodiode at the position of the image of the light spot on the CCD. Do you notice anything out of the ordinary in the transmitted laser light, using a paper screen between the CCD and the Nikon lens? Can you explain the origins of what you see? (The effect seen here is not important for the purposes of this lab, but if it intrigues you, feel free to remove the Nikon lens and measure the separation of points at a known distance from the CCD to solve for the relevant angle, then take this angle and use the known wavelength of the laser to infer the characteristic size within the CCD causing the effect observed. Put the Nikon lens back when you're finished.) Repeat the calculation for the ratio of object to image sizes, and measure the transmitted laser power. Decide what reference or calibration measurement(s) you will need to derive the transmitted fraction, and make these supplementary measurements.

Finally, replace the red laser with the green laser once again and measure the fraction of green laser light (if any) transmitted through the CCD. If you can think of a way to make a meaningful measurement of this light (or establish a meaningful upper limit on the fraction of laser light transmitted), take the necessary measurements. This ends the experimental phase of Part A; trade with the other group to move to the theoretical investigation.

B. Theoretical Investigation

For this portion of the lab you will be using the program "CCD Fringe," which you will find in a folder on the desktop of the PCs. Attached to this lab manual is an appendix which describes the theory of thin films on which the CCD Fringe program is based; you have all no doubt seen interference "fringes" from thin films (oil on wet pavement or water; soap bubbles, etc.) before... this program treats the CCD device structure as a thin film stack and calculates the theoretical maximum

$QE(\lambda)$ from the fraction of incident energy absorbed in one particular device layer which you designate the “photosensitive layer” when entering the device structure. It is hoped that you will gain insights into CCD quantum efficiency by comparing the $QE(\lambda)$ performance in Figure 1, your experimental results using equation (1) at two wavelengths, and theoretical results you derive from the CCD Fringe program. Read through the appendix to familiarize yourself with the theoretical basis of the CCD Fringe program. At the first menu, “How do you wish to enter the device dimensions?”, you should select the interactive choice by typing: 1 <cr> at the prompt.

For the surrounding medium, select ‘Vacuum or Air’ by typing: V<cr> at the prompt (even though your CCD was sandwiched for protection between two fused quartz [silicon dioxide] plates, there is not optical contact between the plates and the CCD so in fact air surrounds the CCD chip).

For the thinned CCD device you’re using, there is only one layer before the photosensitive silicon layer, and this is a silicon dioxide [Oxide] layer 10 or 20 Angstroms thick which grew naturally on the back side after the CCD device was thinned. Adopt a thickness of $0.0015 \mu\text{m}$ for the backside surface oxide layer.

The photosensitive layer is the Si wafer on which the CCD array structure was fabricated. The back surface of the wafer was then thinned away using a liquid chemical etch. The target thickness was 8 to 10 microns; if you saw green laser light transmitted, adopt 8 or fewer microns as the thickness of the photosensitive layer; otherwise, adopt 9 or 10 microns.

After the photosensitive Si layer comes the “frontside” CCD gate structure, which consists of three layers:

The deepest of these (first encountered by any photons which reach the front side of the photosensitive Si layer) is an oxide layer which acts as an insulator between the polysilicon gates and the photosensitive Si layer. The thickness to adopt for this insulating oxide layer is 100 nm, or $0.1 \mu\text{m}$.

The middle layer on the frontside are the gates which provide clocking to transfer charge while reading out the CCD; these gates are made of a highly doped [conducting] silicon called “polysilicon”, and are roughly $0.5 \mu\text{m}$ thick on average.

Lastly, the outermost frontside layer is a protective layer of silicon dioxide [Oxide] 2 to 4 microns thickness, designed to protect the polysilicon gates from damage [open connections or shorts] through contact with dust particles, moisture, etc., once the chip leaves the cleanroom-like conditions in which it was fabricated. Adopt $3.0 \mu\text{m}$ for the protective oxide layer thickness in your model.

The CCD Fringe program will now prompt you for a filename to store the CCD structure information you have given it. You must supply the filename complete with

extension: the extension “.CCD” is a convenient one to use.

Answer: YES<cr> when asked whether you want to plot the device cross section. A Plot Window should appear having an arbitrary horizontal axis, but a vertical axis whose units are microns with 0.0 corresponding to the surface of the chip on which light is incident (the backside in our case). Depth into the CCD increases upwards from this line; confirm that the order and approximate thicknesses of the layers shown corresponds to what you expected based on the input data. Note that in the current version of the CCD Fringe program, the plot lines are ‘remembered’ by Windows should you iconify, overlay another window on top of, move, or resize the Plot Window, but the character text is not retained (so study the window when it appears without doing any of these mouse operations, or the text will be lost!). Close the Plot Window to continue with CCD Fringe by holding down the Alt key and pressing the F4 key, or instead by clicking with the mouse on the ‘-’ box in the upper left hand corner of the window, and clicking secondly on ‘Close’ in the drop down menu with the mouse. Type ‘N’ or use the mouse to click on ‘No’ when asked whether to save the plot (the save plot option has yet to be implemented, so the ‘Yes’ choice will do you no good).

The next menu selections assign “collection probabilities” to the photosensitive layer: that is, when a photon is absorbed by the silicon photosensitive layer creating an electron-hole pair, what is the probability that the electron will be collected and eventually detected? Later on you will experiment with these collection probabilities, but for now select the first choice (set all probabilities to 1.0) by typing: 1<cr> at the menu prompt.

The final input required is to supply a wavelength range. The CCD Fringe program uses input files from the disk containing data equivalent to the complex refractive index \tilde{n} for each material (Silicon, Polysilicon, Silicon Dioxide, etc.) as a function of wavelength. In practice the index of refraction [real part] is taken from the file: MATERIAL.NDX, and the absorption coefficient [related to the imaginary part of \tilde{n}] is taken from the file: MATERIAL.ABS on the disk. The reason for bringing this to your attention now concerns your choice of wavelength range: if you ask CCD Fringe to calculate $QE(\lambda)$ for some λ outside the range over which it has input data for the \tilde{n} of the materials being used, unpredictable results will occur. Experience has shown that such problems are most severe at the short wavelength limit, so (i) you should not select a minimum wavelength less than 3200 Å, and (ii) the program will actually begin calculating $QE(\lambda)$ at the maximum wavelength and work its way towards shorter wavelengths, so that if some error occurs you will be able to see what minimum wavelength was last successful and make future choices accordingly. As a first try, enter 0 0 <cr> for the default wavelength range.

The program will now begin performing a cubic spline interpolation of the refractive index and absorption coefficient data, and then a new Plot Window will be opened in which the resulting values of $QE(\lambda)$ will be plotted. How does this first attempt at $QE(\lambda)$ compare with Figure 1 and your experimental result using equation

(1)? In particular, what can you say about the collection probabilities for $\lambda \simeq 5500 \text{ \AA}$? Furthermore, what do you find most striking about the $QE(\lambda)$ results for long wavelengths, and can you explain this on the basis of your experience with the red versus green lasers? (Hint: read the appendix carefully, and recall what $T(\lambda)$ was for each laser wavelength). Close this plot window as before, and then close the main CCD Fringe text window.

Next, go to the CCD Fringe program directory; here you will find the filename ending in “.CCD” which you wrote when finished with data input. Type this file and look at its format: the first line is the total number of layers, then each layer is listed followed by its thickness in microns (the surrounding medium appears first and last, and the program ignores the first and last thicknesses). Lastly the file contains an integer which signifies which layer in the list is the photosensitive layer. If you prefer, you can edit this file directly (saving it with the same or a different name) to make changes and read it back into CCD Fringe instead of answering the interactive prompts for device structure.

The first exercise you should do is confirm your hypothesis about the cause of the long wavelength $QE(\lambda)$ features. Can you predict what should happen when the thickness of the photosensitive layer is increased or decreased by a few microns? Write down this prediction in your lab notebook, and then run CCD Fringe a second time to determine whether your hypothesis is correct. Record the old and new thicknesses used, and what changes you observed.

Next, type the file: `... \CCD \SILICON.NDX` (the two columns are λ in μm and refractive index [real part]). Compute the reflection coefficient $R = (n - 1)^2 / (n + 1)^2$ for $n = 4.00$ as appropriate for 5600 \AA . Do the same for $\lambda = 3860 \text{ \AA}$ and comment on the shape of the blue side of $QE(\lambda)$ calculated by the CCD Fringe program. Now type the file: `... \CCD \SILICON.ABS` (the first column is again wavelength in microns; the second column is the absorption coefficient¹ α in cm^{-1} , defined such that light is attenuated by the factor $e^{-\alpha x}$ after travelling a distance x through a medium whose absorption coefficient is α). What fraction of the incident light remains after $10 \mu\text{m}$ through silicon at $\lambda = 7000 \text{ \AA}$? At $\lambda = 5000 \text{ \AA}$? At $\lambda = 1 \mu\text{m}$? Do these results confirm your earlier ideas about the red and near-infrared behavior of the $QE(\lambda)$ curve?

Of the CCD characteristics listed at the bottom of page 1, the only one as yet unaddressed is the device-to-device variation in blue quantum efficiency. Looking at the SILICON.ABS file, what can you say about the absorption of photons with $\lambda < 5000 \text{ \AA}$? Specifically, determine at what depth in from the backside of the photosensitive Si layer the incident light energy has dropped by the factor $1/e$ for $\lambda = 5000 \text{ \AA}$, 4000 \AA , and 3500 \AA . Would you expect the typical 3500 \AA photon to be absorbed closer to the back surface of the CCD or farther away compared to the typical 5000 \AA photon? For which wavelength would the $QE(\lambda)$ suffer most if the collection probability very near

¹For the curious, the imaginary part of the complex refractive index \tilde{n} , denoted k , is related to α as follows: $k = \alpha\lambda/4\pi$ where α and λ are in reciprocal units so that k is dimensionless.

the backside were reduced? Such a reduction does in fact happen, due to the tendency of the back surface to acquire a positive charge which creates a “backside potential well” in which photo-generated electrons can become trapped. Run the CCD Fringe program once more, and this time set the collection probabilities to mimic a backside potential well in the first 5% depth into the photosensitive Si layer. Compare the blue and near-UV quantum efficiency to the ideal case (collection probabilities set to 1.0) run earlier. Do you find similar curves among those in Figure 1 to the $QE(\lambda)$ curve generated by the CCD Fringe program for a device with a backside potential well?

Discussion. Summarize what you’ve learned about interpreting the $QE(\lambda)$ curve for CCDs. Based on what you know, do you think $QE=1.0$ (perfect quantum efficiency) can be achieved, or nearly so, for some wavelengths? Suggest how $QE(\lambda)$ might be improved at $\lambda = 5500 \text{ \AA}$. How might you improve it at $\lambda = 1 \mu\text{m}$?

C. CCD Electronics

Set up on the optical table is a Loral CCD Development [circuit] Board containing at its center a socket for a Loral CCD412 chip (512 x 512 CCD detector array). You will use an oscilloscope to observe and measure the CCD clocking waveforms which accomplish charge transfer in the vertical column direction inside the CCD array and the horizontal row direction inside the CCD serial register.

Setup. Sketch the initial configuration you find at the start of lab. You will notice that two benchtop power supplies provide input DC power to the CCD circuit board via a DB25 25-pin connector, and that BNC connectors for coaxial cable are provided for the video signal(s) from the CCD. The CCD412 chip has one output amplifier, so only the VIDEO B connector will be used at the end of this lab (time permitting). Notice also how the board contains numerous identical clock-driver circuits, each of which is provided with “test points” (metal turret lugs rising up from the board) for attaching leads from the digital multimeter or oscilloscope. Finally, note that two of the wires from the DB25 connector go to a momentary contact “reset” switch for the Altera programmable SAM (Stand Alone Microprocessor) chip which controls the CCD clocking logic.

Method. Decide which of the two oscilloscopes you and your partner wish to use. If you have little or no experience with an oscilloscope, chances are you will find the 2252 analog scope more intuitive and easier to use. On the other hand, if you are very familiar already with analog oscilloscopes, you are encouraged to take advantage of this opportunity to learn to use the TDS420 digital scope. Place the oscilloscope selected on the optical table near the CCD development board and power supplies, plug it in, and turn on the oscilloscope [only]. Locate the instruction manual for the scope and familiarize yourself with its basic functions. Unpack the digital multimeter and do likewise.

Using the attached CCD412 data sheet “Pin Connections” chart as a guide, attach

scope probes to vertical phases $\phi V1$, $\phi V2$, $\phi V3$, and the transfer gate ϕVTG . You should insert short pieces of wire into the CCD socket at the appropriate pins and attach the scope probes to the other ends of the wire. BE VERY CAREFUL NOT TO ALLOW THESE WIRES OR THE METAL PARTS OF THE SCOPE PROBES TO TOUCH EACH OTHER OR THE CIRCUIT BOARD. Note that pin VSS is the CCD substrate, and as such is to be used as a common ground for all the scope probes. Once you have all the leads connected and have verified that no wires are touching (ask your instructor or the T.A. if you are at all unsure about something), turn on the two DC power supplies and press the “reset” switch for the SAM chip. Adjust the oscilloscope so that it triggers on the rising edge of the $\phi V3$ signal, and adjust the scope vertical volts/division and horizontal sweep time so that you can see the four waveforms from the CCD vertical clocks. Sketch their relative amplitude and phase in your lab notebook, along with the following measurements: the positive and negative voltage levels of each clock line, the total time $\phi V1$ is low and $\phi V2$, $\phi V3$, and ϕVTG are high, and the time during which any two of $\phi V1$, $\phi V2$, and $\phi V3$ overlap in the high state. Lastly, measure the total time from the rising edge of $\phi V3$ to the falling edge of $\phi V2$ and the time elapsed between two successive rising edges of $\phi V3$. Turn off the DC power supplies when you are finished making these measurements.

Next locate pins on the CCD412 socket for the horizontal clock phases $\phi H1$, $\phi H2$, $\phi H3$, and the reset gate ϕR . Attach the oscilloscope probes to wires inserted in these four sockets, using VSS once again as a common ground for the scope leads. CHECK THAT NO WIRES ARE TOUCHING, then turn on the DC power supplies and press the reset button for the SAM chip. Readjust the oscilloscope, triggering on the rising edge of ϕR . Sketch the waveforms in your lab notebook and measure the high and low voltage levels, pulse widths, and high state overlap for $\phi H1$, $\phi H2$, and $\phi H3$, in a manner similar to that used for the vertical clock phases above. Measure the total time interval between successive ϕR reset pulses. Turn off the DC power supplies and remove the wires from the CCD socket. If less than 45 minutes remain in lab, you may stop at this point and continue with the analysis and discussion sections at home. Time permitting, one partner should put the elastic wrist strap over their (right or left) hand, then slide it up the forearm high and tight enough that the strap is in good contact with your skin. Press the palms of both hands flat against the conductive plastic mat on top of the optical table (CCDs are very sensitive to electrostatic discharge; the sense node capacitor especially can be destroyed by the slightest static charge: hence these precautions that you be well grounded before handling the CCD chip). You may now open the conductive black plastic box containing the CCD, and—using whichever hand (right or left) has the grounding wrist strap—lift the CCD chip out of the conductive foam by the two (smaller) edges of the package which do not have metal pins on them. Note which end has marks indicating pin #1, and insert the CCD into the socket on the development board (the socket is a Zero Insertion Force [ZIF] socket, which means that the lever should be raised before inserting the CCD and lowered afterwards to hold the CCD in place. Doublecheck that pin #1 on the CCD corresponds to pin #1 on the socket.

Attach a BNC cable from the VIDEO B connector on the development board to channel 1 on the oscilloscope. Attach a probe from channel 2 to the LINE SYNC test point on the top of the circuit board, using the adjacent GND test point to ground the scope probe. Then turn on the DC power supplies and push the SAM reset switch. Adjust the oscilloscope so that it triggers on channel 2 (the line sync) and displays channel 1: this is the amplified signal from each pixel from the CCD as it is being read out continuously. Adjust the horizontal sweep time of the scope so that you see one complete line of pixels on the display, as evidenced by the pixel gap between successive lines (this is easier to demonstrate than explain in words, so ask the instructor or T.A. if it's not clear what this instruction means). The video signal is inverted, meaning that light will move the pixel voltages downward on the scope while darkness will move them upwards. The CCD is sensitive enough that room light during continuous readout is sufficient to saturate the pixels; turning out the lights (or conversely, shining a flashlight on the CCD) should reveal a change in the video level on the oscilloscope. Record what happens in your lab notebook. Is there a delay between the time when the CCD is darkened and when the video signal decreases as seen on the oscilloscope?

Next expand the horizontal axis of the oscilloscope so that you can see individual pixels rather than the complete line of 512 pixels. Sketch in your notebook the repeated pixel waveform you observe, and which portion(s) of it change when light is applied. Use the attached board layout drawing to locate resistor R3 in the reset gate ORG driver section, and touch the channel 3 scope probe to the R3 lead farthest from the CCD socket on the board (clip the ground lead from the probe to the same ground lug as for channel 2). Using the logical reset signal as the time reference, measure the pixel video waveform transition times and record the results in your lab notebook. Turn off the DC power supplies and oscilloscope, and once again taking grounding precautions against static charge, carefully return the CCD to its conductive container.

Analysis. Assuming the vertical pixel gates are arranged top to bottom in the order ...-1-2-3-1-2-3-1-2-3-..., determine from your study of the vertical clock waveforms in which direction signal charge is being transferred inside this CCD. Keep in mind that a positive gate voltage creates a potential well for negatively charged electrons. If the charge transfer direction is not obvious to you from the vertical clock waveforms, draw the pixel wells for a series of time steps as done as an example in lecture. Judging from its clock waveform, the vertical transfer gate takes the place of which vertical clock phase? The transfer gate is therefore adjacent to which vertical pixel gate? How long does it take to transfer charge vertically from one row of pixels to the next? How long a time period elapses between vertical transfer cycles?

Assuming the horizontal serial register pixel gates are arranged left to right in the order ...-1-2-3-1-2-3-1-2-3-..., determine from your study of the horizontal clock waveforms in which direction charge is being transferred inside this CCD. How long does it take to transfer charge horizontally from one serial pixel to the next? How long a time period elapses between serial transfer cycles? How many serial transfer cycles can occur during the time the vertical clocks are fixed between vertical transfers?

If you had the opportunity to examine the amplified video signal from the CCD development board, comment on the different levels of the video waveform in the CDS context (Correlated Double Sampling) discussed in class. Can you identify aspect(s) of the video waveform with the action of the reset gate? Comparing the video signal with the horizontal clock timing, can you determine which horizontal clock phase is responsible for transferring charge into the output well?

Discussion. From the serial transfer cycle time, calculate how many pixels are being read out per second. How long therefore does it take these electronics to readout a full 512 x 512 CCD array? Can you safely ignore the vertical transfer time in this calculation? What factors do you think limit the maximum readout speed of CCDs? Include both practical limits (what components in a complete [computer-interfaced] CCD electronics system do you suspect would be slowest?) and more fundamental limits (what is going on inside the CCD which forces the electronics circuitry to wait?). How could you alter the design of the CCD itself to achieve a faster readout speed?